

Part No. 060166-10, Rev. C
March 2005

Omni Switch/Router™ User Manual

Release 4.5



www.alcatel.com

An Alcatel service agreement brings your company the assurance of 7x24 no-excuses technical support. You'll also receive regular software updates to maintain and maximize your Alcatel product's features and functionality and on-site hardware replacement through our global network of highly qualified service delivery partners. Additionally, with 24-hour-a-day access to Alcatel's Service and Support web page, you'll be able to view and update any case (open or closed) that you have reported to Alcatel's technical support, open a new case or access helpful release notes, technical bulletins, and manuals. For more information on Alcatel's Service Programs, see our web page at www.ind.alcatel.com, call us at 1-800-995-2696, or email us at support@ind.alcatel.com.

**This Manual documents Release 4.5 Omni Switch/Router hardware and software.
The functionality described in this Manual is subject to change without notice.**

Copyright© 2005 by Alcatel Internetworking, Inc. All rights reserved. This document may not be reproduced in whole or in part without the express written permission of Alcatel Internetworking, Inc.

Alcatel® and the Alcatel logo are registered trademarks of Alcatel. Xylan®, OmniSwitch®, PizzaSwitch® and OmniStack® are registered trademarks of Alcatel Internetworking, Inc.

AutoTracker™, OmniAccess™, OmniCore™, Omni Switch/Router™, OmniVista™, PizzaPort™, PolicyView™, RouterView™, SwitchManager™, SwitchStart™, VoiceView™, WANView™, WebView™, X-Cell™, X-Vision™ and the Xylan logo are trademarks of Alcatel Internetworking, Inc.

All-In-OneSM is a service mark of Alcatel Internetworking, Inc. All other brand and product names are trademarks of their respective companies.



A L C A T E L

26801 West Agoura Road
Calabasas, CA 91301
(818) 880-3500 FAX (818) 880-3505
info@ind.alcatel.com

US Customer Support—(800) 995-2696
International Customer Support—(818) 878-4507
Internet—<http://eservice.ind.alcatel.com>

Cautions

FCC Compliance: This equipment has been tested and found to comply with the limits for Class A digital device pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions in this guide, may cause interference to radio communications. Operation of this equipment in a residential area is likely to cause interference, in which case the user will be required to correct the interference at his own expense.

The user is cautioned that changes and modifications made to the equipment without approval of the manufacturer could void the user's authority to operate this equipment. It is suggested that the user use only shielded and grounded cables to ensure compliance with FCC Rules.

This equipment does not exceed Class A limits per radio noise emissions for digital apparatus, set out in the Radio Interference Regulation of the Canadian Department of Communications.

Avis de conformité aux normes du ministère des Communications du Canada

Cet équipement ne dépasse pas les limites de Classe A d'émission de bruits radioélectriques pour les appareils numériques, telles que prescrites par le Règlement sur le brouillage radioélectrique établi par le ministère des Communications du Canada.

Lithium Batteries Caution: There is a danger of explosion if the Lithium battery in your chassis is incorrectly replaced. Replace the battery only with the same or equivalent type of battery recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions. The manufacturer's instructions are as follows:

Return the module with the Lithium battery to Alcatel. The Lithium battery will be replaced at Alcatel's factory.

Table of Contents

1 Omni Switch/Router Chassis and Power Supplies	1-1
Omni Switch/Router User Interface (UI) Software	1-2
Omni Switch/Router Network Management Software (NMS)	1-2
Omni Switch/Router Distributed Switching Fabric	1-3
Omni Switch/Router Fabric Capacity	1-4
Omni Switch/Router Applications and Configurations	1-5
Omni Switch/Router as the Backbone Connecting Several Networks	1-5
Omni Switch/Router as the Central Backbone Switch/Router and in the Wiring Closet	1-6
Omni Switch/Router Chassis and Power Supplies	1-7
OmniS/R-3	1-8
OmniS/R-3 Chassis Technical Specifications	1-9
OmniS/R-5	1-10
OmniS/R-5 Technical Specifications	1-12
OmniS/R-9 and OmniS/R-9P	1-13
OmniS/R-9 Technical Specifications	1-15
OmniS/R-9P Technical Specifications	1-16
OmniS/R-9P-48V Technical Specifications	1-17
Omni Switch/Router Power Requirements	1-18
Grounding a Chassis	1-21
The Omni Switch/Router Hardware Routing Engine (HRE-X)	1-22
Valid HRE-X Configurations	1-23
HRE-X Router Registers versus Feature Limitations	1-23
Connecting a DC Power Source to an OmniS/R-PS5-DC375	1-24
Installing DC Power Source Wire Leads	1-24
Connecting a DC Power Source to an OmniS/R-PS9-DC725	1-27
Installation Requirements	1-27
Installing DC Power Source Wire Leads	1-28
Replacing Power Supplies (9-Slot Chassis)	1-30
2 The Omni Switch/Router MPX	2-1
Omni Switch/Router Management Processor Module (MPX) Features	2-1
MPX Technical Specifications	2-1
MPX Serial and Ethernet Management Ports	2-4
Ethernet Management Port	2-5
Configuring MPX Serial Ports	2-6
Flash Memory and Omni Switch/Router Software	2-7
Flash Memory Guidelines	2-8

MPX Redundancy	2-9
Change-Over Procedure	2-9
MPX Redundancy Commands	2-10
3 Omni Switch/Router Switching Modules	3-1
Required Image Files	3-3
Installing a Switching Module	3-4
Removing a Switching Module	3-6
Hot Swapping a Switching Module	3-7
Diagnostic Tests	3-9
Handling Fiber and Fiber Optic Connectors	3-10
Gigabit Ethernet Modules	3-12
GSX-K-FM/FS/FH-2W	3-12
GSX-K-FM/FS/FH-2W Technical Specifications	3-13
Auto-Sensing 10/100 Ethernet Modules	3-15
Ethernet RJ-45 Pinouts	3-15
Ethernet RJ-45 Specifications	3-15
ESX-K-100C-32W	3-15
ESX-K-100C-32W Technical Specifications	3-17
Fast (100 Mbps) Ethernet Modules	3-19
ESX-K-100FM/FS-16W	3-19
ESX-K-100FM/FS-16W Technical Specifications	3-20
WAN Modules	3-22
WAN Pinouts	3-22
WAN BRI Port Specifications (S/T Interface)	3-23
WAN BRI Port Specifications (U Interface)	3-23
WAN T1/E1 Port Specifications	3-24
WAN Serial Port Specifications	3-25
WSX-S-2W	3-27
WSX-S-2W Technical Specifications	3-27
WSX-SC	3-29
WSX-SC Technical Specifications	3-30
WSX-FT1/E1-SC	3-32
WSX-FT1/E1-SC Technical Specifications	3-33
WSX-FE1-SC Cabling/Jumper Settings	3-35
WSX-BRI-SC	3-36
WSX-BRI-SC Technical Specifications	3-37

4 The User Interface	4-1
Overview of Command Interfaces	4-1
Changing Between the CLI and UI Modes	4-2
Exit the Command Interface	4-3
UI to CLI Command Cross Reference	4-4
Hardware Commands	4-4
Hardware Table	4-4
Basic Switch Management Commands	4-4
Basic Switch Management Table	4-5
Network Management Commands	4-6
Network Management Table	4-6
Layer II Switching Commands	4-7
Layer II Switching Table	4-7
Groups, VLANs, Policies Commands	4-8
Groups, VLANs, Policies Table	4-8
Routing Commands	4-10
Routing Table	4-10
WAN Access Commands	4-11
WAN Access Table	4-11
Troubleshooting Diagnostics Commands	4-13
Troubleshooting/Diagnostics Table	4-13
User Interface Menu	4-14
Main Menu Summary	4-15
General User Interface Guidelines	4-16
Entering Command Names	4-16
Quitting a Command	4-17
Scrolling	4-17
The UI Configuration Menu	4-17
Configuring the System Prompt	4-18
Configuring More Mode for the User Interface	4-19
Setting Verbose/Terse Mode for the User Interface	4-22
Configuring the Auto Logout Time	4-24
Viewing Commands	4-25
Changing Passwords	4-25
Command History and Re-Executing Commands	4-26
Abbreviating IP Addresses	4-28
User Interface Display Options	4-30
Setting Echo/NoEcho for User Entry	4-31
Setting the Login Banner	4-31
Creating a new Banner	4-32
Permanent Banner	4-32
Banners for Different Access Methods	4-32
Login Accounts	4-33

Multiple User Sessions	4-33
Listing Other Users	4-34
Communicating with Other Users	4-35
Deleting Other Sessions	4-35
Advanced Kill Command Options	4-37
UI Table Filtering (Using Search and Filter Commands)	4-38
The Search Command	4-39
Renewing a Search	4-40
The Filter Command	4-41
Combining Search and Filter Commands	4-42
Using Wildcards with Search and Filter Commands	4-44
Wildcard Command Options	4-44
5 Installing Switch Software	5-1
Using FTP Server	5-2
Using FTP Client	5-3
Using ZMODEM	5-4
Using ZMODEM with the load Command	5-4
Using ZMODEM With the Boot Line Prompt	5-5
6 Configuring Management Processor Modules	6-1
Changing Serial Port Communication Parameters	6-2
Changing Port Speed When Communication With The Switch Lost	6-3
Configuring the Modem Port	6-3
Modem Port Mode	6-3
Configuring SLIP	6-4
Configuring the Ethernet Management Port	6-5
Ethernet Management Ports and Redundant Management Processor Modules	6-7
The MPM Command/Menu	6-9
Displaying MPX Redundancy	6-9
MPM Menu Commands	6-9
Using MPM Commands with Software Release 3.2 and Later	6-10
Listing the Secondary MPX Files	6-11
Transferring a File to the Secondary MPX	6-11
Replacing a File on the Secondary MPX	6-12
Loading a File from the Secondary MPX	6-12
Removing a File from the Secondary MPX	6-13
Giving Up Control to the Secondary MPX	6-14
Setting the Load Suffix	6-14
Setting Automatic Config Synchronization	6-15
Enabling Automatic Config Synchronization	6-15
Disabling Automatic Config Synchronization	6-15

Synchronizing Configuration Data	6-16
Synchronizing Image Files	6-16
Loading a File From the Primary MPX	6-17
Gaining Control from the Primary MPX	6-18
Resetting a Secondary MPX	6-19
Displaying and Setting the Swap State	6-20
Displaying the Swap State	6-20
Enabling the Swap Mode	6-20
Disabling the Swap Mode	6-21
7 Managing Files	7-1
File Menu	7-1
Displaying the Current Directory	7-2
Configuration and Log File Generation	7-2
Changing Directories	7-2
Listing Switch Files	7-3
Deleting Switch Files	7-4
Deleting Multiple Files	7-4
Deleting All Image Files	7-5
Copying System Files	7-6
Displaying Text Files	7-6
Editing Text Files	7-7
Clearing the Text Buffer	7-7
Loading an ASCII File into the Text Buffer	7-8
Listing the Contents of the Text Buffer	7-8
Adding Lines of Text to the Text Buffer	7-8
Deleting a Line of Text from the Text Buffer	7-9
Inserting a Line of Text into the Text Buffer	7-9
Editing a Line Name of Text in the Text Buffer	7-9
Creating a File Name for the Text Buffer	7-10
Creating a Text File from the Text Buffer	7-10
Real-World Examples	7-11
Real-World Example 1	7-11
Real-World Example 2	7-12
System Menu	7-13
Checking the Flash File System	7-14
Creating a New File System	7-15
8 Switch Security	8-1
Changing Passwords	8-2
Rebooting the Switch	8-3

Secure Switch Access	8-4
Configuring the Secure Switch Access Filter Database	8-4
Configuring Secure Access Filter Points	8-7
Enabling/Disabling Security Parameters	8-9
Adding Filters	8-9
Deleting Filters	8-9
Viewing Secure Access Violations Log	8-10
Managing User Login Accounts	8-11
Partition Management Requirements	8-11
Default Accounts	8-12
Adding a User Account Using the UI Command Mode	8-12
Adding a User Account Using the CLI Command Mode	8-13
Assigning Account Privileges Using the CLI Command Mode	8-13
Assigning Account Privileges Using the UI Command Mode	8-16
Command Family Table	8-18
Global Family Table	8-19
Modifying a User Account	8-20
Deleting a User	8-20
9 Configuring Switch-Wide Parameters	9-1
Summary Menu	9-1
Displaying the MIB-II System Group Variables	9-2
Displaying the Chassis Summary	9-3
Displaying Current Router Interface Status	9-4
System Menu	9-5
Displaying Basic System Information	9-6
Setting the System Date and Time	9-8
Viewing Slot Data	9-14
Viewing System Statistics	9-15
Clearing System Statistics	9-16
Viewing Task Utilization Statistics	9-17
Viewing Memory Utilization	9-19
Viewing MPX Memory Statistics	9-20
Checking the Flash File System	9-21
Checking the SIMM Files	9-21
Creating a New File System	9-22
Creating a SIMM File System	9-22
Configuring System Information	9-23
Viewing CAM Information	9-24
Configuring CAM Distribution	9-25
Configuring the HRE-X Router Port	9-27
Configuring and Displaying the HRE-X Hash Table	9-29

Duplicate MAC Address Support	9-30
Multicast Claiming	9-32
Disabling Flood Limits	9-32
Saving Configurations	9-33
10 Switch Logging	10-1
Logging Overview	10-1
Configuring the Syslog Parameters	10-2
Configuring Switch Logging	10-6
Displaying the Command History Entries in the MPM Log	10-9
Displaying the Connection Entries in the MPM Log	10-10
Displaying Screen (Console) Capture Entries in the MPM Log	10-11
Displaying Debug Entries in the MPM Log	10-13
Displaying Secure Access Entries in the MPM Log	10-13
11 Health Statistics	11-1
The Health Statistics Management Menu	11-1
Setting Resource Thresholds	11-2
Setting Bandwidth Thresholds	11-3
Setting Miscellaneous Thresholds	11-4
Setting the Sampling Interval	11-6
View Switch-Level Statistics	11-6
View Module-Level Statistics	11-7
View Port-Level Statistics	11-8
Reset Health Statistics	11-8
12 Network Time Protocol	12-1
Introduction	12-1
Stratum	12-2
Using NTP in a Network	12-2
NTP and Authentication	12-4
Network Time Protocol Management Menu	12-5
NTP Configuration Menu	12-6
Configuring an NTP Client	12-6
Configuring an NTP Client/Server	12-8
Configuring Client/Server Authentication	12-9
Configuring a New Peer Association	12-12
Configuring a New Server	12-13
Configuring a Broadcast Time Service	12-13
Unconfigure Existing Peer Associations	12-14
Set the Server's Advertised Precision	12-14

NTP Information Menu	12-15
Display List of Peers the Server Knows About	12-15
Display Peer Summary Information	12-16
Display Alternate Peer Summary Information	12-17
Display Detailed Information for One or More Peers	12-18
Print Version Number	12-20
Display Local Server Information	12-21
NTP Statistics Menu	12-23
Display Local Server Statistics	12-23
Display Server Statistics Associated with Particular Peer(s)	12-24
Display Loop Filter Information	12-26
Display Peer Memory Usage Statistics	12-26
Display I/O Subsystem Statistics	12-27
Display Event Timer Subsystem Statistics	12-28
Reset Various Subsystem Statistics Counters	12-28
Reset Stat Counters Associated With Particular Peer(s)	12-28
Display Packet Count Statistics from the Control Module	12-29
Display the Current Leap Second State	12-30
Turn the Server's Monitoring Facility On or Off	12-31
Display Data The Server's Monitor Routines Have Collected	12-31
NTP Administration Menu	12-33
Set the Primary Receive Timeout	12-33
Set the Delay Added to Encryption Time Stamps	12-33
Specify the Host Whose NTP Server We Talk To	12-34
Specify a Password to Use for Authenticated Requests	12-34
Set Key ID to Use for Authenticated Requests	12-34
Set Key Type to Use for Authenticated Requests (DES MD5)	12-35
Set a System Flag (Auth, Bclient, Monitor, Stats)	12-35
Clear a System Flag (Auth, Bclient, Monitor, Stats)	12-35
NTP Access Control Menu	12-36
Change the Request Message Authentication Key ID	12-36
Change the Control Message Authentication Key ID	12-37
Add One or More Key ID's to the Trusted List	12-37
Display the Trusted Key ID List	12-37
Remove One or More Key ID's from the Trusted List	12-38
Display the State of the Authentication Code	12-38
Create Restrict Entry/Add Flags to Entry	12-39
View the Server's Restrict List	12-40
Remove Flags from a Restrict Entry	12-41
Delete a Restrict Entry	12-41
Configure a Trap in the Server	12-41
Display the Traps Set in the Server	12-42
Remove a Trap (Configured or Otherwise) from the Server	12-42

13	SNMP (Simple Network Management Protocol)	13-1
	Introduction	13-1
	Configuring SNMP Parameters and Traps	13-2
	Configuring a New Network Management Station	13-4
	Viewing SNMP Statistics	13-8
	Trap Tables	13-11
	SNMP Standard Traps	13-15
	Extended Traps	13-27
14	DNS Resolver and RMON	14-1
	Introduction	14-1
	Configuring the DNS Resolver	14-1
	The Names Submenu	14-1
	Remote Network Monitoring (RMON)	14-3
	Probes and Events	14-3
	Ethernet Probes	14-3
	History Probes	14-3
	Alarm Probes	14-3
	Monitoring Probes	14-4
	Monitoring Events	14-5
	Configuring Router Port MAC Addresses	14-6
	Restoring Router Port Mac Addresses	14-6
15	Managing Ethernet Modules	15-1
	Overview of Omni Switch/Router Ethernet Modules	15-1
	Kodiak Ethernet Modules	15-3
	The Ethernet Management Menus	15-4
	Configuring 10/100 Auto-Sensing Ports	15-5
	Connecting Kodiak Modules to Non-Auto-Negotiating Links	15-6
	Configuring Kodiak Ethernet Ports	15-7
	Viewing Configurations for 10/100 Ethernet Modules	15-8
	OmniChannel	15-9
	The Server Channel Feature	15-10
	Server Channel Limitations	15-11
	Creating an OmniChannel	15-11
	Adding Ports to an OmniChannel	15-13
	Deleting an OmniChannel	15-13
	Deleting Ports from an OmniChannel	15-14
	Viewing OmniChannel Parameters	15-14

16	Managing 802.1Q Groups	16-1
	IEEE 802.1Q Sections Not Implemented	16-2
	Application Example	16-3
	Single vs. Multiple Spanning Tree	16-4
	Assigning an 802.1Q Group to a Port	16-7
	Configuring 802.1Q on 10/100 Ethernet Ports	16-8
	Configuring 802.1Q on Gigabit Ethernet Ports	16-11
	Modifying 802.1Q Groups	16-12
	Modifying 802.1Q Groups for 10/100 Ports	16-12
	Modifying 802.1Q Groups for Gigabit Ethernet Ports	16-14
	Viewing 802.1Q Groups in a Port	16-16
	Viewing 802.1Q Statistics for 10/100 Ports	16-17
	Deleting 802.1Q Groups from a Port	16-18
17	Configuring Bridging Parameters	17-1
	Configuration Overview	17-3
	Bridge Management Menu	17-4
	Selecting a Default Group	17-7
	Using the + or - to Change Groups	17-7
	Bridging Commands	17-8
	Displaying Bridge Forwarding Table	17-8
	Configuring a Static Bridge Address	17-10
	Modifying a Static Bridge Address	17-11
	Deleting a Static Bridge Address	17-12
	Displaying Static Bridge Addresses	17-13
	Displaying Bridge Port Statistics	17-14
	Displaying Media Access Control (MAC) Information for a Specific MAC address	17-16
	Displaying Media Access Control (MAC) Information for all MAC addresses	17-17
	Display Statistics of Bridge MAC Addresses	17-17
	Clear Statistics of Bridge MAC Addresses	17-18
	Display Remote Trunking Stations	17-18
	View the Domain Bridge Mapping Table	17-19
	Setting Flood Limits	17-21
	Setting Flood Limits for a Group	17-21
	Displaying Group Flood Limits	17-22
	Configuring Spanning Tree	17-23
	Configuring Spanning Tree Parameters	17-25
	Display Spanning Tree Bridge Parameters	17-28
	Configuring Spanning Tree Port Parameters	17-30
	Displaying Spanning Tree Port Parameters	17-32

Configuring Fast Spanning Tree	17-34
Truncating Tree Timing & Speedy Tree Protocol	17-35
Truncating Tree Timing	17-35
Speedy Tree Protocol	17-35
Configuring Truncating Tree Timing & Speedy Tree Protocol	17-35
Displaying Fast Spanning Tree Port Parameters	17-36
Enabling Fast Spanning Tree Port Parameters	17-38
Disabling Fast Spanning Tree Port Parameters	17-39
Configuring Source Routing	17-40
SAP Filtering	17-40
Enabling SAP Filtering	17-40
Disabling SAP filtering	17-41
Configuring SAP Filtering	17-41
Viewing SAP Filtering	17-42
Configuring Source Route to Transparent Bridging	17-43
Enabling SRTB for a Group	17-44
Disabling SRTB for a Group	17-45
Viewing the RIF Table	17-46
Clearing the RIF Table	17-47
18 Configuring Frame Translations	18-1
Any-to-Any Switching	18-1
Translating the Frame	18-3
The MAC Header	18-4
Canonical versus Non-Canonical	18-4
Abbreviated Addresses	18-4
Functional Addresses and Multicasts	18-4
The RIF Field	18-5
Source Route Termination by Proxy Not Supported	18-5
Encapsulation	18-6
Protocols other than IP and IPX	18-6
The SNAP Conversion	18-7
Other Conversions	18-7
Summary of Non-IPX Encapsulation Transformation Rules	18-7
IPX Encapsulation Transformation Rules	18-8
The Network Header	18-9
Address Mapping	18-9
Address Mapping in IP: ARP	18-9
Address Mapping in IPX	18-10
Frame Size Requirements	18-11
Insertion of Frame Padding	18-11
Stripping of Padding for all IEEE 802.3 Frames.	18-11
No stripping of non-IPX Ethertype Frames	18-11
IPX Specific Stripping	18-11

MTU Handling	18-12
IP Fragmentation	18-12
ICMP Based MTU Discovery	18-12
IPX Packet Size Negotiation	18-12
Other Protocols	18-12
Banyan Vines	18-13
Configuring Encapsulation Options	18-14
Forwarding versus Flooding	18-14
Port Based Translation Options	18-14
MAC Address Based Translation Options	18-14
“Native” versus “Non-Native” on Ethernet	18-15
“Native” versus “Non-Native” on FDDI and Token Ring	18-15
No Translation on Trunk or PTOP ports	18-15
The Proprietary Token Ring IPX Option	18-15
The User Interface	18-16
The addvp, modvp and crgp Commands	18-17
The Default Translation Option	18-17
Ethernet Factory Default Translations	18-18
FDDI Factory Default Translations	18-18
Token Ring Factory Default Translations	18-19
ATM LANE Factory Default Translations	18-19
The Ethertype Option	18-20
The SNAP Option	18-21
The LLC Option	18-23
Interaction with the new interface	18-24
The “vi” Command	18-24
The Switch Menu	18-25
Proprietary IPX Token Ring	18-25
Factory Defaults	18-25
Default Ethernet Translations	18-26
Default FDDI Translations	18-27
Default Token Ring Translations	18-28
Port Translations	18-30
Configuring Additional Ports	18-31
Displaying Ethernet Switch Statistics	18-31
Displaying Token Ring Switch Statistics	18-35
Any to Any MAC Translations	18-39
Default Autoencapsulation	18-40
Translational Bridging	18-41
Learning	18-41
Translations across Trunks	18-41
Dissimilar LAN Switching Capabilities	18-42
Switching Between Similar LANs	18-42
Switching Between Ethernet LANs Across a Trunked Backbone	18-43
Switching Between Similar LANs across a Native Backbone	18-44

19 Managing Groups and Ports	19-1
How Ports Are Assigned to Groups	19-2
Static Port Assignment	19-2
Dynamic Port Assignment (Group Mobility)	19-2
How Dynamic Port Assignment Works	19-3
Mobile Groups	19-5
Configuring Mobile Groups	19-5
Turning Group Mobility On or Off	19-6
Understanding Port Membership in Mobile Groups	19-7
How a Device Is Dropped from the Default Mobile Group (def_group) ..	19-9
How a Port's Primary Mobile Group Changes (move_from_def)	19-10
How a Port Ages Out of a Mobile Group (move_to_def)	19-11
Configuring Switch-Wide Group Mobility Variables	19-12
Viewing Ports in a Mobile Group	19-14
Viewing a Port's Mobile Group Affiliations	19-14
Non-Mobile Groups and AutoTracker VLANs	19-15
Routing in a Non-Mobile Group	19-15
Spanning Tree and Non-Mobile Groups	19-16
Group and Port Software Commands	19-17
Creating a New Group	19-18
Step 1. Entering Basic Group Information	19-19
Step 2. Configuring the Virtual Router Port (Optional)	19-21
Step 3. Set Up Group Mobility and User Authentication	19-27
Step 4. Configuring Virtual Ports	19-28
Step 5. Configuring AutoTracker Policies (Mobile Groups Only)	19-34
Creating a WAN Routing Group	19-35
Viewing Current Groups	19-38
Modifying a Group or VLAN	19-40
Viewing Your Changes	19-41
Saving Your Changes	19-41
Canceling Your Changes	19-41
Changing the IP Address	19-41
Changing the IP Subnet Mask	19-41
Enabling IP or IPX Routing	19-42
Deleting a Group	19-43
Adding Virtual Ports	19-44
Modifying a Virtual Port	19-45
Deleting a Virtual Port	19-46
Viewing Information on Ports in a Group	19-47
Viewing Detailed Information on Ports	19-50
Viewing Port Statistics	19-53
Viewing Port Errors	19-55

Port Mirroring19-57

- How Port Mirroring Works19-57
- What Happens to the Mirroring Port19-57
- Using Port Mirroring With External RMON Probes19-58

Setting Up Port Mirroring19-60

Disabling Port Mirroring19-60

Port Monitoring19-61

- Port Monitoring Menu19-61
- RAM Disk System for Data Capture Files19-62
 - Configuring RAM Drive Resources (pmcfg)19-62
 - Changing the Default System Directory (cd)19-62
- Starting a Port Monitoring Session (pmon)19-63
 - If You Chose Dump to Screen19-64
 - If You Did Not Choose Dump to Screen19-64
- Ending a Port Monitoring Session19-65
- Viewing Port Monitoring Statistics (pmstat)19-65

Port Mapping19-66

- Groups/VLANs and Port Mapping19-66
- The Details of Port Mapping19-67
 - Who Can Talk to Whom?19-68
- Port Mapping Limitations19-68
- Creating a Port Mapping Set19-69
- Adding Ports to a Port Mapping Set19-70
- Removing Ports from a Port Mapping Set19-71
- Viewing a Port Mapping Set19-72
- Deleting a Port Mapping Set19-72

Priority VLANs19-73

- Mammoth vs. Kodiak Priority VLANs19-73
- Configuring VLAN Priority19-74
- Viewing VLAN Priority19-74

20 Configuring Group and VLAN Policies 20-1

AutoTracker Policy Types 20-2

Defining and Configuring AutoTracker Policies 20-4

- Where These Procedures Start 20-4
- Defining a Port Policy 20-5
- Defining a MAC Address Policy 20-6
- Defining a MAC Address Range Policy 20-7
- Defining a Protocol Policy 20-8
- Defining a Network Address Policy 20-11
- Defining Your Own Rules 20-13
- Defining a Port Binding Policy 20-15
- Defining a DHCP Port Policy 20-20
- Defining a DHCP MAC Address Policy 20-21
- Defining a DHCP MAC Address Range Policy 20-22

Viewing Mobile Groups and AutoTracker VLANs 20-23

Viewing Policy Configurations	20-24
Viewing Virtual Ports' Group/VLAN Membership	20-25
View VLAN Membership of MAC Devices	20-26
Application Example: DHCP Policies	20-27
The VLANs	20-27
DHCP Servers and Clients	20-28
DHCP Port and MAC Rules	20-29
21 Interswitch Protocols	21-1
Interswitch Protocol Commands	21-1
XMAP	21-2
XMAP Transmission States	21-3
Discovery Transmission State	21-3
Common Transmission State	21-4
Passive Reception State	21-4
Common Transmission and Remote Switches	21-4
Configuring XMAP	21-5
Enabling or Disabling XMAP	21-5
Viewing a List of Adjacent Switches	21-5
Configuring the Discovery Transmission Time	21-6
Configuring the Common Transmission Time	21-7
VLAN Advertisement Protocol (VAP)	21-8
VAP and Port Policies	21-9
Configuring VAP	21-9
GMAP	21-10
GMAP Updating Rules	21-10
Configuring GMAP	21-11
Enabling and Disabling GMAP	21-11
Configuring the Gap Time	21-11
Configuring the Interpacket Update Time	21-12
Configuring the Hold Time	21-12
Displaying GMAP Statistics by MAC Address	21-13
22 Managing AutoTracker VLANs	22-1
The AutoTracker Menu	22-2
AutoTracker VLANs	22-3
AutoTracker VLAN Policies	22-3
The Default VLAN	22-4
How Devices are Assigned to AutoTracker VLANs	22-5
The devvl Command	22-5
Devices that Generate a Secondary Traffic Type	22-6
Router Traffic in IP and IPX Network Address VLANs	22-7
Port Policy Functionality	22-9
Frame Flooding in AutoTracker VLANs	22-15

Routing Between AutoTracker VLANs	22-15
Creating AutoTracker VLANs	22-16
Step A. Entering Basic VLAN Information	22-16
Step B. Defining and Configuring VLAN Policies	22-18
Step C. Configuring the Virtual Router Port (Optional)	22-19
Modifying an AutoTracker VLAN	22-24
Deleting an AutoTracker VLAN	22-26
Viewing AutoTracker VLANs	22-27
Viewing Policy Configurations	22-28
Viewing Virtual Ports' VLAN Membership	22-29
View VLAN Membership of MAC Devices	22-30
Creating a VLAN for Banyan Vines Traffic	22-31
23 Multicast VLANs	23-1
How Devices are Assigned to Multicast VLANs	23-2
Multicast VLANs and Multicast Claiming	23-2
Frame Flooding in Multicast VLANs	23-3
Creating Multicast VLANs	23-4
Step A. Entering Basic Information	23-5
Step B. Defining the Multicast Address	23-6
Step C. Defining the Recipients of Multicast Traffic	23-7
Defining Recipients By Port	23-7
Defining Recipients By MAC Address	23-8
Modifying Multicast VLANs	23-9
Deleting a Multicast VLAN	23-11
Modifying a Multicast Address Policy	23-12
Viewing Multicast VLANs	23-13
Viewing Multicast VLAN Policies	23-14
Viewing the Virtual Interface of Multicast VLANs	23-15
24 AutoTracker VLAN Application Examples	24-1
Application Example 1	24-2
VLANs Based on Logical Policies	24-2
Application Example 2	24-4
VLANs in IPX Networks	24-4
IPX VLAN Assignment at Bootup	24-5
Application Example 3	24-7
IPX Network Address VLANs and Translated Frames	24-7
Application Example 4	24-8
Routing in IPX Networks	24-8
Application Example 5	24-10
Traversing a Backbone	24-10

25 IP Routing	25-1
Introduction	25-1
IP Routing Overview	25-2
Routing Protocols	25-2
Transport Protocols	25-3
Application-Layer Protocols	25-3
Additional IP Protocols	25-3
Setting Up IP Routing on the Switch	25-4
The Networking Menu	25-6
The IP Submenu	25-7
Viewing the Address Translation (ARP) Table	25-8
Displaying All Entries in the ARP Table	25-8
Adding Entries to the ARP Table	25-9
Deleting Entries from the ARP Table	25-10
Flushing Temporary Entries from the ARP Table	25-10
Finding a Specific IP Address in the ARP Table	25-10
Finding a Specific MAC Address in the ARP Table	25-11
Viewing IP Statistics and Errors	25-12
Viewing the IP Forwarding Table	25-15
Adding an IP Static Route	25-17
Removing an IP Static Route	25-19
Viewing ICMP Statistics and Errors	25-20
Using the PING Command	25-22
Viewing UDP Statistics and Errors	25-24
Viewing the UDP Listener Table	25-25
Viewing RIP Statistics and Errors	25-26
Viewing TCP Statistics	25-27
Viewing the TCP Connection Table	25-29
Using the TELNET Command	25-30
Cancelling a Telnet request	25-30
Tracing an IP Route	25-31
Flushing the RIP Routing Tables	25-32
Configuring IP RIP Filters	25-33
Adding a “Global” IP RIP Filter	25-33
Adding an IP RIP Filter For a Specific Group or VLAN	25-34
IP RIP Filter Precedence	25-35
Deleting IP RIP Filters	25-36

Displaying IP RIP Filters	25-37
Displaying a List of All IP RIP Filters	25-37
Displaying a List of “Global” IP RIP Filters	25-38
Displaying a List of Specific IP RIP Filters	25-38
Viewing the IP-to-MAC Address Table	25-39
Displaying All Entries in the IP-to-MAC Table	25-39
Displaying Information for a Specific IP Address	25-40
Flushing Entries from the Table	25-40
Enabling/Disabling Directed Broadcasts	25-41
Path MTU Discovery	25-42
26 UDP Forwarding	26-1
UDP Relay and RIF Stripping	26-1
UDP Relay Hardware/Software Support	26-2
UDP Relay Configuration Screen	26-3
BOOTP/DHCP Relay	26-4
Overview of DHCP	26-4
DHCP and the OmniS/R	26-4
BOOTP/DHCP Relay and Source Routing	26-5
BOOTP/DHCP Relay and Authentication	26-5
External BOOTP Relay	26-6
Internal BOOTP/DHCP Relay	26-7
Example 1	26-7
Example 2	26-8
Enabling BOOTP/DHCP Relay	26-9
Configuring BOOTP/DHCP Relay Parameters	26-10
NetBIOS Relays	26-11
Overview of NetBIOS	26-11
NetBIOS Relay Application	26-12
Configuring NBNS Relay	26-13
Next-Hop Addresses for NBNS	26-14
Forwarding VLANs for NBNS Relay	26-15
Configuring NBDD Relay	26-16
Next-Hop Addresses for NBDD	26-17
Forwarding VLANs for NBDD Relay	26-18
Generic Service UDP Relay	26-19
Generic Services Menu	26-19
Adding a Generic Service	26-19
Modifying a Generic Service	26-21
Deleting a Generic Service	26-22
Viewing UDP Relay Statistics	26-23

27 IPX Routing	27-1
Introduction	27-1
IPX Routing Overview	27-2
IPX Protocols	27-2
Setting Up IPX Routing on the Switch	27-3
The IPX Submenu	27-4
Viewing the IPX Routing Table	27-5
Displaying All Entries in the IPX Routing Table	27-5
Using IPXR with Frame Relay or ISDN Boards	27-6
Displaying a List of Specific IPX Routes	27-7
Viewing IPX Statistics	27-8
Viewing the IPX SAP Bindery	27-10
Using IPXSAP with Frame Relay or ISDN Boards	27-11
Displaying a List of Specific SAP Servers	27-11
Adding an IPX Static Route	27-12
Removing an IPX Static Route	27-13
Turning the IPX Router Complex On and Off	27-14
Flushing the IPX RIP/SAP Tables	27-15
Using the IXPING Command	27-16
Configuring IPX RIP/SAP Filtering	27-18
Adding a “Global” IPX RIP/SAP Filter	27-19
Adding an IPX RIP/SAP Filter for a Specific Group or VLAN	27-20
Deleting an IPX RIP/SAP Filter	27-22
Displaying IPX RIP/SAP Filters	27-23
Displaying a List of All IPX Filters	27-23
Displaying a List of “Global” IPX Filters	27-24
Displaying a List of Specific IPX Filters	27-24
IPX RIP/SAP Filter Precedence	27-25
Configuring IPX Serialization Packet Filtering	27-26
Enabling IPX Serialization Filtering	27-26
Disabling IPX Serialization Filtering	27-27
Configuring IPX Watchdog Spoofing	27-28
Enabling IPX Watchdog Spoofing	27-28
Disabling IPX Watchdog Spoofing	27-29
Configuring SPX Keepalive Spoofing	27-30
Enabling SPX Keepalive Spoofing	27-30
Disabling SPX Keepalive Spoofing	27-31
Controlling IPX Type 20 Packet Forwarding	27-32
Configuring NetWare to Minimize WAN Connections	27-33
Configuring RIP and SAP Timers	27-35
Adding a RIP and SAP Timer	27-35
Viewing RIP and SAP Timers	27-36

Configuring Extended RIP and SAP Packets	27-37
Enabling or Disabling Extended RIP and SAP Packets	27-37
Viewing the Current Status of Extended Packets	27-37
Configuring an IPX Default Route	27-38
Adding an IPX Default Route	27-38
Viewing the Status of an IPX Default Route	27-38
Disabling an IPX Default Route	27-38
28 Managing WAN Switching Modules	28-1
Introduction	28-1
Type of Service (ToS)	28-2
ToS and QoS Interaction	28-4
DTR Dial Backup	28-5
Supported Physical Interfaces	28-6
Universal Serial Port	28-6
ISDN Basic Rate Interface Port	28-6
Fractional T1 Port	28-6
Fractional E1 Port	28-6
Supported Protocols	28-7
Application Examples	28-7
Frame Relay WSX Using Serial Ports	28-7
Back-to-Back WSX Using T1 Ports	28-8
Combined Frame Relay with ISDN Backup	28-9
Omni Switch/Router WAN Modules	28-10
Cable Interfaces for Universal Serial Ports	28-11
DTE/DCE Type and Transmit/Receive Pins	28-11
Data Compression	28-12
Loopback Detection	28-13
The WAN Port Software Menu	28-14
Setting Configuration Parameters	28-14
Modifying a Port	28-14
Serial Port Example	28-15
ISDN-BRI Port Example	28-21
Fractional T1 Port Example	28-24
Viewing Configuration Parameters for the WSX	28-27
Viewing Parameters for all Submodules in the Chassis	28-27
Viewing Parameters for all Ports in a Single Submodule	28-28
Viewing Port Parameters	28-29
Deleting Ports	28-37
Obtaining Status and Statistical Information	28-38
Obtaining Information on All Boards in a Switch	28-38
Obtaining Information on the Ports for a Single WSX Board	28-40
Viewing Information on a Single Port	28-42
Configuring 31 Timeslots on a WAN E1 Port	28-45

29 Managing Frame Relay	29-1
Back-to-Back Frame Relay Configurations	29-3
Universal Serial Port Cable Interfaces	29-4
“Physical” and “Logical” Devices	29-4
Compression	29-5
Virtual Circuits and DLCIs	29-6
WSX Self-Configuration and Virtual Circuits	29-7
Congestion Control	29-8
Regulation Parameters	29-8
Discard Eligibility (DE) Flag	29-9
Interaction Among Congestion Parameters	29-9
Notification By BECN	29-11
Notification By FECN	29-12
Frame Formats Supported	29-13
Bridging Services	29-14
Frame Relay IP Routing	29-15
The Frame Relay Subnet and “Split Horizon”	29-16
Frame Relay IPX Routing	29-18
Trunking	29-19
Frame Relay Fragmentation Interleaving	29-20
The Frame Relay Software Menu	29-21
Setting Configuration Parameters	29-22
Modifying a Port	29-22
Modifying a Virtual Circuit	29-29
Adding a Virtual Circuit	29-32
Viewing Configuration Parameters for the WSX	29-33
Viewing Parameters for all WSXs in the Chassis	29-33
Viewing Port Parameters	29-34
Viewing Virtual Circuit Parameters	29-35
Deleting Ports and Virtual Circuits	29-36
Deleting a Virtual Circuit	29-36
Deleting a Port and Its Virtual Circuits	29-37
Obtaining Status and Statistical Information	29-38
Information on All Boards in a Switch	29-38
Information on the Ports for One WSX Board	29-42
Information on One Port	29-43
Information on One Virtual Circuit	29-51
Resetting Statistics Counters	29-54
Resetting Statistics for a WSX Board	29-54
Resetting Statistics for a WSX Port	29-54
Resetting Statistics for a Virtual Circuit (DLCI)	29-54

Managing Frame Relay Services	29-55
Configuring a Bridging Service	29-57
Configuring a WAN Routing Service	29-59
Step 1. Set Up a Frame Relay Routing Group	29-59
Step 2. Set Up a Frame Relay Routing Service	29-60
Configuring a Trunking Service	29-62
Viewing Frame Relay Services	29-64
Modifying a Frame Relay Service	29-65
Deleting a Frame Relay Service	29-66
30 Point-to-Point Protocol	30-1
PPP Connection Phases	30-1
Data Compression	30-2
Multi-Link PPP	30-2
Multilink Modes of Operation	30-3
PPP Fragmentation Interleaving	30-3
Overview of PPP Configuration Procedures	30-4
The PPP Submenu	30-6
PPP Configuration Overview	30-6
Setting Global PPP Parameters	30-7
Adding a PPP Entity	30-9
Modifying a PPP Entity	30-15
Viewing PPP Entity Configurations	30-16
Displaying the Configuration of All PPP Entities	30-16
Displaying the Configuration of a Specific PPP Entity	30-17
Displaying PPP Entity Status	30-18
Displaying the Status of All PPP Entities	30-18
Displaying the Status of a Specific PPP Entity	30-19
Deleting a PPP Entity	30-21
31 WAN Links	31-1
Introduction	31-1
Configuring WAN Interfaces	31-1
The Link Submenu	31-2
Adding a WAN Link	31-3
Adding WSX Port Links	31-3
Adding ISDN Call Links	31-4
Modifying a WAN Link	31-9
Modifying ISDN Links	31-9
Modifying WSX Links	31-10
Deleting WAN Links	31-11

Viewing WAN Links	31-12
Displaying All Existing WAN Links	31-12
Displaying Information for a Specific WAN Link	31-13
Displaying Link Status	31-15
Displaying Status for All WAN Links	31-15
Displaying Status for a Specific WAN Link	31-16
32 Managing ISDN Ports	32-1
Overview of ISDN	32-1
Basic Rate Interface (BRI) Versus Primary Rate Interface (PRI)	32-1
“U”, “S/T”, and “R” Interfaces	32-2
The “B,” “D,” and “H” Channels	32-2
The ISDN Submenu	32-3
Switch Configuration	32-3
Modifying an ISDN Configuration Entry	32-4
Deleting an ISDN Configuration Entry	32-5
Viewing an ISDN Configuration Entry	32-6
Displaying ISDN Configuration Entry Status	32-7
Displaying Status of All ISDN Ports	32-7
Displaying Status of a Specific ISDN Slot	32-8
Displaying Status of a Specific ISDN Port	32-9
33 Managing T1 and E1 Ports	33-1
T1 and E1 Overview	33-2
The T1/E1 Menu	33-3
Configuring a T1 Port	33-4
Configuring an E1 Port	33-8
Viewing T1/E1 Configuration and Alarm Information	33-11
Viewing Information for all T1/E1 Ports in the Switch	33-11
Viewing Information for T1/E1 Ports on One Module	33-12
Viewing Information For a T1 Port	33-13
Viewing Information For an E1 Port	33-15
Viewing T1/E1 Local Statistics	33-17
Viewing Total Local Statistics	33-17
Viewing Current Local Statistics	33-18
Viewing Local Historical Statistics	33-19
Viewing T1 Remote Statistics	33-20
Viewing Total Remote Statistics	33-20
Viewing Current Remote Statistics	33-21
Viewing Remote Historical Statistics	33-21
Clearing the Framers Statistics for a T1/E1 Port	33-22

34 Backup Services	34-1
Introduction	34-1
Backup Services Commands	34-2
Accessing the Backup Services Menu	34-2
Adding a Backup Service	34-3
Adding a backup for a Physical Port	34-3
Backing Up a Frame Relay PVC	34-6
Modifying a Backup Service	34-9
Modifying a backup for a Physical Port	34-9
Modifying a Frame Relay PVC Backup Service	34-10
Viewing Backup Service(s) Configurations	34-11
Viewing the Configurations of All Backup Services	34-11
Viewing the Configuration of a Single Backup Service (bsview Command)	34-11
Deleting a Backup Service	34-11
Viewing Backup Service Statistics	34-12
Clearing Backup Service Statistics	34-13
35 Troubleshooting	35-1
Detecting Problems	35-1
Reporting Problems	35-3
Report Hardware Details	35-3
Report Software Details	35-4
Understanding Problems	35-5
Software Installation Problems	35-5
Operational Problems	35-6
Deadlocked VLAN	35-6
Probable Cause	35-7
Solution	35-7
Problems with IP Applications	35-7
Probable Cause	35-7
Solution	35-7
Protocol Problems	35-8
Probable Cause	35-8
Solution	35-8
Hardware Problems	35-9
LEDs Do Not Light on All Modules	35-9
Probable Cause	35-9
Solution	35-9
Amber Color in LEDs	35-9
Probable Cause	35-9
Solution	35-9
Non-Blinking OK2 LED	35-9
Probable Cause	35-9
Solution	35-9

TEMP LED is Amber	35-10
Solution	35-10
STA LED Is Off	35-10
Probable Cause	35-10
Solution	35-10
Switch Does Not Boot When Flash File System Is Full and Trying To Create the mpm.cnf File	35-10
Probable Cause	35-10
Solution	35-10
Error Messages	35-11
Understanding Error Messages	35-11
Correcting Errors	35-11
Module Startup/Shutdown Error Messages	35-11
Serial Port Configuration Errors	35-12
Module Connection Errors	35-12
Chassis Error Messages	35-13
Chassis Error Messages Table	35-13
36 Running Hardware Diagnostics	36-1
Running Diagnostics	36-2
Login to Run Diagnostics	36-3
Resetting a Switching Module	36-4
Disabling a Switching Module	36-4
Temperature Masking	36-5
Running Hardware Diagnostics	36-6
Sample Command Lines	36-9
Halting Diagnostic Tests in Progress	36-9
Port Tests	36-9
Omni Switch/Router Port Test Wrap Cable/Plug Requirements	36-10
Sample Test Session: Ethernet Module	36-12
Displaying Available Diagnostic Tests	36-15
Configuring the Diagnostic Test Environment	36-16
Configuring Tests for Ethernet Modules	36-17
Running Frame Fabric Tests on Omni Switch/Routers	36-18
Running Diagnostics on an Entire Chassis	36-20
Diagnostic Test Cable Schematics	36-22

A	The Boot Line Prompt	A-1
	Entering the Boot Prompt	A-2
	Boot Prompt Basics	A-3
	Resuming Switch Boot (@)	A-3
	Displaying Current Configuration (p)	A-4
	Loading the Last Configured Boot File (l)	A-4
	Listing Available Files in the Flash Memory (L)	A-5
	Deleting All Files in the Flash Memory (P)	A-5
	Deleting Specific Files in the Flash Memory (R)	A-5
	Saving Configuration Changes (S)	A-6
	Viewing Version Number (V)	A-6
	Configuring a Switch with an MPX	A-7
B	Custom Cables	B-1
	V.35 DTE Cable (For WSX-to-DCE Device Connection)	B-2
	V.35 DCE Cable (For WSX-to-DTE Device Connection)	B-3
	RS232 DTE Cable (For WSX-to-DCE Device Connection)	B-4
	RS232 DCE Cable (For WSX-to-DTE Device Connection)	B-5
	RS530 DTE Cable (For WSX-to-DCE Device Connection)	B-6
	RS530 DCE Cable (For WSX-to-DTE Device Connection)	B-7
	X.21 DTE Cable (For WSX-to-DCE Device Connection)	B-8
	X.21 DCE Cable (For WSX-to-DTE Device Connection)	B-9
	RS449 DTE Cable (For WSX-to-DCE Device Connection)	B-10
	RS-449 DCE Cable Assembly (For WSX-to-DTE Device 75W Connection)	B-11
	RJ-45 to DB15F Cable Assembly (For T1/E1 Port 120W Connections)	B-12
	RJ-45 to BNC Cable Assembly (For E1 75W Port Connections)	B-13
	Index	I-1

1 Omni Switch/Router Chassis and Power Supplies

Alcatel's Omni Switch/Router (OmniS/R) is an advanced, multi-layer switching platform (Layer 2 and 3) that supports the most demanding switch requirements. With Omni Switch/Router, network administrators can replace aging FDDI or Fast Ethernet backbones with high capacity Gigabit Ethernet backbones.

◆ Important Notes ◆

Beginning with Release 4.4, FDDI is no longer supported. Beginning with Release 4.5, ATM, Token Ring, M013, and Mammoth-based Ethernet Modules are no longer supported.

Omni Switch/Router modules can be distinguished from older OmniSwitch modules by the **X** in the module name. For example, the ESM-100C-32W is an OmniSwitch module whereas the ES**X**-100C-32W is an Omni Switch/Router module.

Omni Switch/Router has a distributed switching fabric. In a 9-slot chassis operating at full duplex, Omni Switch/Router offers an aggregate 22 Gigabit per second (Gbps) distributed switching fabric. In addition, Omni Switch/Router offers new high density switching modules, including auto-sensing 10/100 Ethernet modules that offer high speed network connections to servers and desktops. (See *Omni Switch/Router Applications and Configurations* on page 1-5 for examples.)

The Omni Switch/Router Management Processor Module (MPX) module provides the core routing, VLAN MAC learning, SNMP, and file management functions for the entire Omni Switch/Router. In addition, the MPX has an Ethernet plug-in port for managing the switch. Only one MPX is required per Omni Switch/Router, but you can add another MPX for redundancy. See Chapter 2, "The Omni Switch/Router MPX," for more information on the MPX.

◆ Important Note ◆

Omni Switch/Router switching modules require an MPX. You cannot install any version of the MPM (i.e., MPM-C, MPM 1G, MPM II, or original MPM) in a chassis with an MPX.

An Omni Switch/Router Hardware Routing Engine (HRE-X). The HRE-X offers high-speed Layer 3 switching from 1.5 to 12.0 million packets per second (Mpps) in a fully loaded chassis. See *The Omni Switch/Router Hardware Routing Engine (HRE-X)* on page 1-22 for more information on the HRE-X.

Omni Switch/Router switching modules perform software filtering, translations between dissimilar network interfaces, and hardware-based switching. Omni Switch/Router switching modules have an additional on-board interface connector for the HRE-X.

Currently, Omni Switch/Router switching modules consist of Gigabit Ethernet modules, auto-sensing Ethernet modules, Fast 10/100 Ethernet modules, 10 Mbps Ethernet modules, WAN modules, and Voice Over IP (VOIP) modules. See Chapter 3, “Omni Switch/Router Switching Modules,” for documentation.

◆ **Important Note** ◆

Omni Switch/Router modules require the use of an Omni Switch/Router chassis (see *Omni Switch/Router Chassis and Power Supplies* on page 1-7). Do *not* install an Omni Switch/Router module in an OmniSwitch chassis and do *not* install an OmniSwitch module in an Omni Switch/Router chassis.

Omni Switch/Router User Interface (UI) Software

Omni Switch/Router hardware uses the same User Interface (UI) commands and Network Management Software (NMS) as OmniSwitch hardware. Omni Switch/Router modules support broadcast management, multicast management, any-to-any switching, virtual LANs (VLANs), firewalls, user authentication, WAN access, and policy-based configuration.

◆ **Important Note** ◆

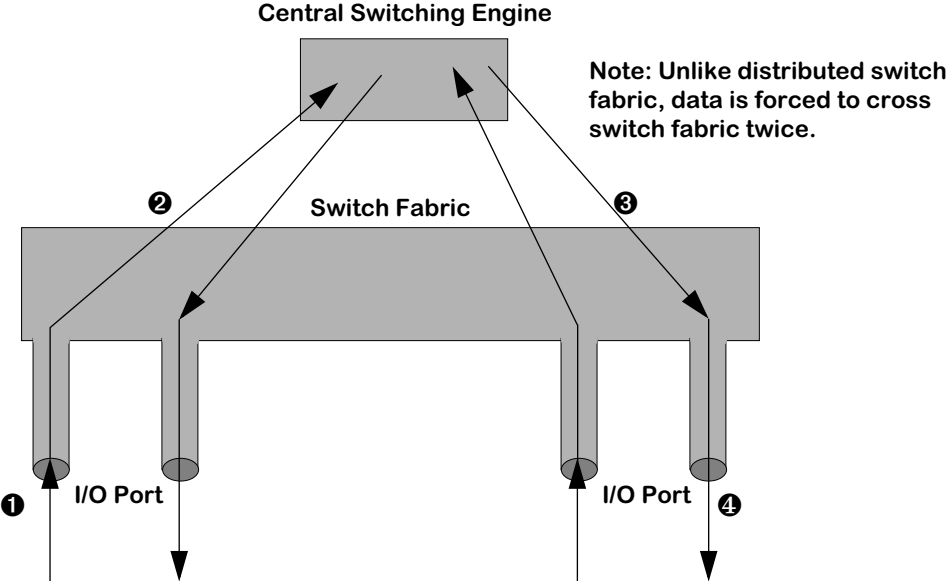
In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. Chapter 4, “The User Interface,” includes documentation on changing from CLI mode to UI mode.

Omni Switch/Router Network Management Software (NMS)

You need Release 3.4, or higher, of Alcatel’s X-Vision Network Management Software (NMS) to operate with Omni Switch/Router hardware.


Omni Switch/Router Distributed Switching Fabric

Many switches in the market employ a shared memory architecture, which uses a central switching engine to send data to the appropriate port. As shown in the figure below, data enters the input port (1 below), crosses the switching fabric on its way to the central switching engine(2 below), and *again* crosses the switching fabric (3 below) before exiting the appropriate output port (4 below).



Traditional Shared Memory Architecture


In contrast, Omni Switch/Router switches use a distributed switching fabric. As shown in the figure below, data enters the input port and crosses the switching fabric *only once* before exiting the appropriate output port. Compared to the shared memory architecture, only half as much bandwidth is required since data just crosses the switching fabric once.



Omni Switch/Router Distributed Switching Fabric

Omni Switch/Router Fabric Capacity

In a chassis with Omni Switch/Router modules only, each Omni Switch/Router module provides 2.4 Gbps of switching capacity in full-duplex mode. In a chassis with all Omni Switch/Router modules, the Omni Switch/Router architecture provides up to a 22 Gbps distributed switching fabric. As shown in the figure below, an OmniS/R-9 with an MPX and eight (8) Omni Switch/Router switching modules provides 22 Gbps of switching capacity. An OmniS/R-5 with an MPX and four (4) Omni Switch/Router switching modules provides 12 Gbps of switching capacity, while an OmniS/R-3 with an MPX and two (2) Omni Switch/Router switching modules provides 7 Gbps of switching capacity.




Omni Switch/Router Fabric Capacity in OmniS/R-3, OmniS/R-5 and OmniS/R-9 Chassis

Omni Switch/Router Applications and Configurations

Omni Switch/Router hardware is ideally suited to meet the most demanding server and backbone needs. In addition, Omni Switch/Router hardware can be integrated easily with OmniSwitches and with OmniStack workgroup switches. The examples that follow show how the Omni Switch/Router can be used as a network backbone and as the central switch/router in a wiring closet.

Omni Switch/Router as the Backbone Connecting Several Networks

The figure below shows how Omni Switch/Router Gigabit Ethernet and 10/100 Ethernet modules can be used as a network backbone. In this example, two networks on two different floors need high speed access to a server farm on the first floor.




Using Omni Switch/Router in a Network Backbone

The servers each have dedicated Gigabit Ethernet connections to Omni Switch/Router modules on the first floor. The Omni Switch/Router chassis on the first floor is connected to the network on the second floor via a Gigabit Ethernet link to the OmniStack on the second floor. The Omni Switch/Router chassis on the first floor is connected via a 10/100 Ethernet link, using OmniChannel, to the OmniSwitch chassis on the third floor containing a Fast Ethernet module, such as the ESM-100C-12. See Chapter 15, “Managing Ethernet Modules,” for more information on OmniChannel.

Omni Switch/Router as the Central Backbone Switch/Router and in the Wiring Closet

The figure below shows Omni Switch/Router chassis used in the wiring closet and as a network backbone switch/router connecting the wiring closets and server farm. On the third floor, an Omni Switch/Router chassis connects a mixture of 10BaseT and 100BaseTx workstations with an auto-sensing Ethernet module. In addition, this Omni Switch/Router chassis connects the workstations to a local server with a Gigabit Ethernet module. On the second floor, an Omni Switch/Router connects legacy Token Ring workstations. On the first floor, the Omni Switch/Router connects the networks on the upper floors to the server farm using a Gigabit Ethernet module.



Using Omni Switch/Router in the Wiring Closet

Omni Switch/Router Chassis and Power Supplies

The Omni Switch/Router chassis houses the MPX, switching modules, and one or two power supplies. The modular design of the chassis provides the ability to configure your Omni Switch/Router to meet your networking needs. The Omni Switch/Router chassis also offer such failure resistant features as redundant MPXs, redundant power supplies, and hot swapping of switching modules. (See Chapter 3, “Omni Switch/Router Switching Modules,” for more information on hot swapping switching modules.)

There are three (3) different versions of the Omni Switch/Router chassis. The OmniS/R-3, a three-slot version, is documented in *OmniS/R-3* on page 1-8. The OmniS/R-5, a five-slot version, is documented in *OmniS/R-5* on page 1-10. A nine-slot version called the OmniS/R-9 is documented in *OmniS/R-9 and OmniS/R-9P* on page 1-13. The OmniS/R-3, OmniS/R-5 and OmniS/R-9 chassis, the MPX module, and several switching modules have met FCC Class B requirements.

◆ Note ◆

In the current release, a maximum of seven (7) 32-port switching modules (e.g., ESX-100C-32W) is supported in 9-slot Omni Switch/Router chassis.

Slot 1 is reserved for the MPX; you *cannot* install a switching module in Slot 1. You can install a switching module in Slot 2 (if an MPX is installed in Slot 1) or an MPX. When dual-redundant MPXs are installed, one of them must be installed in Slot 1 and the other in Slot 2. On the OmniS/R-3, Slot 3 is reserved for a switching module. On the OmniS/R-5, Slots 3 through 5 are reserved for switching modules. On the OmniS/R-9, Slots 3 through 9 are reserved for switching modules.

◆ Important Note ◆

You *must* have an MPX acting as the management module; you cannot use any version of the MPM.

Warning

If you have any empty switching module slots in either an OmniS/R-3 (3-slot) or OmniS/R-5 (5-slot) chassis, you *must* cover them with blank panels (available from Alcatel) to prevent your chassis from overheating.

Covering empty slots forces air to flow directly over the power supplies, thereby cooling them. If the power supplies are not properly cooled, they will overheat and shut down.


OmniS/R-3

The OmniS/R-3 chassis features three slots for an MPX and specific switching modules (contact your Alcatel sales representative for information on module availability). Slots are numbered from 1 to 3 starting with the topmost slot. A built-in power supply is located on the right side of the chassis, and a fan cooling system is located on the left side of the chassis. The chassis can be rack-mounted. You can view all cabling, power supplies, module interfaces, and LEDs at the front of the chassis.

The OmniS/R-3 uses a built-in AC power supply that has a capacity of 32.8 Amps at 5 volts and 3 amps at 12 volts for 200 Watts of output power. The OmniS/R-3 does not support a Backup Power Supply (BPS).

◆ Caution ◆

Do not connect the power connector on the back of the OmniS/R-3 to data communication equipment.



OmniS/R-3 Chassis

◆ Important Note ◆

Slot 1 (the top slot) on the OmniS/R-3 is reserved for an MPX module. Slot 2 can accommodate either a second (optional) MPX module or a Switching module. Slot 3 (the bottom slot) is reserved for a Switching module. Contact your Alcatel sales representative for information regarding module availability.

OmniS/R-3 Chassis Technical Specifications	
Total Module Slots	3
Total Slots for Switching Modules	2
Physical Dimensions	5.25" (13.34 cm) high, 17.13" (43.51 cm) wide, 13.00" (33.02 cm) deep
Weight	18 lb. (8.18 kg), fully populated with modules and power supplies.
Switching Backplane	Up to 7 Gbps (aggregate) switching fabric capacity
Voltage Range	85-270 VAC, 47 to 63 Hz, auto-ranging and auto-sensing
Current Draw	3.8 Amps at 100/115 VAC 1.7 Amps at 230 VAC
Watts (Output)	200
Current Provided	32.8 Amps at +5 Volts 3 Amps at +12 Volts
Heat Generation	Approximately 1020 BTUs per hour
Temperature Operating Range	0 to 45 degrees Celsius 32 to 113 degrees Fahrenheit
Humidity	5% to 90% Relative Humidity (Operating) 0% to 95% Relative Humidity (Storage)
Altitude	Sea level to 10,000 feet (3 km)
Agency Listings	UL 1950 CSA-C22.2 EN60950 FCC Part 15, Subpart B (Class A) EN55022, 1987/EN50081 FCC Class B C.I.S.P.R. 22: 1985 EN50082-1, 1992 IEC 801-2, 1991 IEC 801-3, 1984 IEC 801-4, 1988 VCCI V-3/94.04 (Class A & Class B) EN 61000-4-2: 1995 EN 61000-4-3: 1995 EN 61000-4-4: 1995 EN 61000-4-5: 1995 EN 61000-4-6: 1996 EN 61000-4-8: 1993 EN 61000-4-11: 1994 ENV 50204: 1996

OmniS/R-5


The OmniS/R-5 chassis has five slots for an MPX and switching modules (see figure below). Slots are numbered from 1 to 5 starting with the topmost slot. Slots for two power supplies are located at the bottom of the chassis.

◆ Warning ◆

If you have an OmniS/R-5 with a single power supply, do *not* remove the cover on the empty power supply slot. In addition, if you have any empty switching module slots in an OmniS/R-5, you *must* cover them with blank panels (available from Alcatel) to prevent your chassis from overheating.

Covering empty slots forces air to flow directly over the power supplies, thereby cooling them. If the power supplies are not properly cooled, they will overheat and shut down.

The entire chassis can be wall-mounted or rack-mounted. You can view all cabling, power supplies, module interfaces, and LEDs at the front of the chassis.



The OmniS/R-5

The OmniS/R-5 uses the MPX. Slot 1 is reserved for the MPX; you *cannot* install a switching module in Slot 1. You can install a switching module in Slot 2 (if an MPX is installed in Slot 1) or an MPX. When dual-redundant MPXs are installed, one of them must be installed in Slot 1 and the other in Slot 2. Slots 3 through 5 are reserved for switching modules.

The OmniS/R-5 provides bays for two power supplies. The power supplies are self-enclosed to allow safe hot-insertion and hot-removal. When two power supplies are installed, they share the electrical load. If one should fail, the remaining power supply automatically takes up the load without any disruption to the operation. See Chapter 1, “Omni Switch/Router Chassis and Power Supplies,” for more information on installing and removing power supplies. See *OmniS/R-5 Technical Specifications* on page 1-12 for more information.

The OmniS/R-5 uses one of the following power supplies:

OmniS/R-PS5-375 The standard power supply. It can provide 375 Watts of power.

OmniS/R-PS5-DC375 A -48 volt (input voltage) DC version of the OmniS/R-PS5-375 power supply. This power supply can provide 375 Watts of power. It requires the use of 12 to 14 gauge wire for connections to the DC power source. See *Connecting a DC Power Source to an OmniS/R-PS5-DC375* on page 1-24 for more information.

◆ Caution ◆

This unit may be equipped with two power connections. To reduce the risk of electrical shock, disconnect both power connections before servicing the unit.

◆ VORSICHT ◆


Das Gerät kann mit zwei Netzanschlüssen ausgestattet sein. Um einen elektrischen Schlag zu vermeiden, immer beide Anschlüsse vor der Wartung vom Netz trennen.

Omni Switch/Router Chassis and Power Supplies

OmniS/R-5 Technical Specifications	
Total Module Slots	5
Total Slots for Switching Modules	4
Physical Dimensions	12.25" (31.12 cm) high, 17.14" (43.54 cm) wide, 13" (33.02 cm) deep
Weight	approximately 55 lb. (24.09 kg), fully populated with modules and power supplies.
Switching Backplane	Up to 12 Gbps (aggregate) switching fabric capacity
Voltage Range	90-265 VAC, 47 to 63 Hz auto-ranging and auto-sensing.
Current Draw	6 Amps at 100/115 VAC; 3 Amps at 230 VAC
Watts (Output)	375
Current Provided	60 Amps at 5 Volts (V1) 5 Amps at 12 Volts (V2) 3 Amps at 3.3 Volts (V3) 5.1 Amps at 1.5 Volts (V4)
Temperature Operating Range	0 to 45 degrees Celsius 32 to 113 degrees Fahrenheit
Humidity	5% to 90% Relative Humidity (Operating) 0% to 95% Relative Humidity (Storage)
Altitude	Sea level to 10,000 feet (3 km)
Heat Generation	1280 BTUs per hour (one power supply)
Agency Listings	UL 1950 CSA-C22.2 EN60950 FCC Part 15, Subpart B (Class A) EN55022, 1987/EN50081 FCC Class B C.I.S.P.R. 22: 1985 EN50082-1, 1992 IEC 801-2, 1991 IEC 801-3, 1984 IEC 801-4, 1988 VCCI V-3/94.04 (Class A & Class B) EN 61000-4-2: 1995 EN 61000-4-3: 1995 EN 61000-4-4: 1995 EN 61000-4-5: 1995 EN 61000-4-6: 1996 EN 61000-4-8: 1993 EN 61000-4-11: 1994 ENV 50204: 1996

OmniS/R-9 and OmniS/R-9P

The OmniS/R-9 and OmniS/R-9P chassis have nine slots for an MPX and switching modules (see figure below). Slots are numbered from 1 to 9 starting with the left-most slot. Slots for two power supplies are located at the bottom of the chassis. A separate, removable fan tray containing four fans is located above the power supply module bays.



The OmniS/R-9

A fully loaded OmniS/R-9 weighs nearly 100 lbs. Therefore, it is recommended that if you are rack-mounting the chassis you use a rack mount shelf instead of just brackets. Using a shelf will ensure that the weight of the chassis can be supported. In addition, the OmniS/R-9 contains side handles to make lifting and installation easier.

The OmniS/R-9 uses the MPX. Slot 1 is reserved for the MPX; you *cannot* install a switching module in Slot 1. You can install a switching module in Slot 2 (if an MPX is installed in Slot 1) or an MPX. When dual-redundant MPXs are installed, one of them must be installed in Slot 1 and the other in Slot 2. Slots 3 through 9 are reserved for switching modules.

◆ Important Note ◆

You *must* have an MPX acting as the management module; you cannot use any version of the MPM. See Chapter 2, “The Omni Switch/Router MPX,” for more information on the MPX.

The OmniS/R-9 and OmniS/R-9P provide bays for two power supplies. The power supplies are self-enclosed to allow safe hot-insertion and hot-removal. When two power supplies are installed, they share the electrical load. If one should fail, the remaining power supply automatically takes up the load without any disruption to the operation. See Chapter 1, “Omni Switch/Router Chassis and Power Supplies,” for additional information on installing and removing power supplies.

The OmniS/R-9 uses the following power supply:

OmniS/R-PS9-650 The standard power supply. It can provide 650 Watts of power.

The OmniS/R-9P uses the following power supply:

OmniS/R-PS9-725 This power supply can provide 725 Watts of power.

The OmniS/R-9P-48V uses the following power supply:

OmniS/R-PS9-DC725 A -48 Volt (input voltage) DC version of the OmniS/R-PS9-725 power supply. This power supply can provide 725 Watts of power. It requires the use of 12 to 14 gauge wire for connections to the DC power source. See *Connecting a DC Power Source to an OmniS/R-PS9-DC725* on page 1-27 for more information.

For additional information, see *OmniS/R-9 Technical Specifications* on page 1-15, *OmniS/R-9P Technical Specifications* on page 1-16 and *OmniS/R-9P-48V Technical Specifications* on page 1-17.

◆ Caution ◆

This unit may be equipped with two power connections. To reduce the risk of electrical shock, disconnect both power connections before servicing the unit.

◆ VORSICHT ◆

Das Gerät kann mit zwei Netzanschlüssen ausgestattet sein. Um einen elektrischen Schlag zu vermeiden, immer beide Anschlüsse vor der Wartung vom Netz trennen.

OmniS/R-9 Technical Specifications	
Total Module Slots	9
Total Slots for Switching Modules	8
Physical Dimensions	24.50" (62.23 cm) high, 16.60" (42.16 cm) wide, 13.25" (36.66 cm) deep
Weight	96 lb. (43.55 kg), fully populated with modules and power supplies.
Switching Backplane	Up to 22 Gbps (aggregate) switching fabric capacity
Voltage Range	90-264 VAC, 47 to 63 Hz
Current Draw	12 Amps at 100/115 VAC; 6 Amps at 230 VAC
Watts (Output)	650
Current Provided	120 Amps at 5 Volts 4 Amps at 12 Volts 6 Amps at 3.3 Volts 8 Amps at 1.5 Volts
Temperature Operating Range	0 to 45 degrees Celsius 32 to 113 degrees Fahrenheit
Humidity	5% to 90% Relative Humidity (Operating) 0% to 95% Relative Humidity (Storage)
Altitude	Sea level to 10,000 feet (3 km)
Heat Generation	2219 BTUs per hour (one power supply)
Agency Listings	UL 1950 CSA-C22.2 EN60950 FCC Part 15, Subpart B (Class A) EN55022, 1987/EN50081 FCC Class B C.I.S.P.R. 22: 1985 EN50082-1, 1992 IEC 801-2, 1991 IEC 801-3, 1984 IEC 801-4, 1988 VCCI V-3/94.04 (Class A & Class B) EN 61000-4-2: 1995 EN 61000-4-3: 1995 EN 61000-4-4: 1995 EN 61000-4-5: 1995 EN 61000-4-6: 1996 EN 61000-4-8: 1993 EN 61000-4-11: 1994 ENV 50204: 1996

OmniS/R-9P Technical Specifications	
Total Module Slots	9
Total Slots for Switching Modules	8
Physical Dimensions	24.50" (62.23 cm) high, 16.60" (42.16 cm) wide, 13.25" (36.66 cm) deep
Weight	96 lb. (43.55 kg), fully populated with modules and power supplies.
Switching Backplane	Up to 22 Gbps (aggregate) switching fabric capacity
Voltage Range	85-270 VAC, 47 to 63 Hz
Current Draw	12 Amps at 100/115 VAC; 6 Amps at 230 VAC
Watts (Output)	725
Current Provided	120 Amps at 5 Volts 6 Amps at 12 Volts 6 Amps at 3.3 Volts 8 Amps at 1.5 Volts
Temperature Operating Range	0 to 70 degrees Celsius 32 to 158 degrees Fahrenheit
Humidity	5% to 90% Relative Humidity (Operating) 0% to 95% Relative Humidity (Storage)
Altitude	Sea level to 10,000 feet (3 km)
Heat Generation	2219 BTUs per hour (one power supply)
Agency Listings	UL 1950; CSA-C22.2 #950-M90; TUV EN60950; CB Certification IEC 950; FCC Title 47 CRF Part 15, Subpart B (Class A & Class B); IEC EN55022, 1995 (Class A & Class B) CISPR 22, 1995; IEC 1000-3-2; IEC 1000-3-3 (EN60555-2); IEC 1000-4-2 (EN61000-4-2, per EN50082-1, 1992); IEC 1000-4-3 (EN61000-4-3, per EN50082-1, 1992); IEC 1000-4-4 (EN61000-4-4) Level 4; IEC 1000-4-5 (EN61000-4-5) Level 4; IEC 1000-4-6 (EN61000-4-6); IEC 1000-4-8 (EN61000-4-8); IEC 1000-4-11 (EN61000-4-11); EN50204: 1996.

OmniS/R-9P-48V Technical Specifications	
Total Module Slots	9
Total Slots for Switching Modules	8
Physical Dimensions	24.50" (62.23 cm) high, 16.60" (42.16 cm) wide, 13.25" (36.66 cm) deep
Weight	96 lb. (43.55 kg), fully populated with modules and power supplies.
Switching Backplane	Up to 22 Gbps (aggregate) switching fabric capacity
Voltage Range	40-60 VDC
Current Draw	23 Amps
Watts (Output)	725
Current Provided	120 Amps at 5.15 VDC 6 Amps at 12 VDC 6 Amps at 3.3 VDC 8 Amps at 1.5 VDC
Temperature Operating Range	0 to 70 degrees Celsius 32 to 158 degrees Fahrenheit
Humidity	5% to 90% Relative Humidity (Operating) 0% to 95% Relative Humidity (Storage)
Altitude	Sea level to 10,000 feet (3 km)
Heat Generation	2219 BTUs per hour (one power supply)
Agency Listings	UL 1950; CSA-C22.2 #950-M90; TUV EN60950; CB Certification IEC 950; FCC Title 47 CRF Part 15, Subpart B (Class A & Class B); IEC EN55022, 1995 (Class A & Class B) CISPR 22, 1995; IEC 1000-3-2; IEC 1000-3-3 (EN60555-2); IEC 1000-4-2 (EN61000-4-2, per EN50082-1, 1992); EN55024 IEC 1000-4-3 (EN61000-4-3, per EN50082-1, 1992); IEC 1000-4-4 (EN61000-4-4) Level 4; IEC 1000-4-5 (EN61000-4-5) Level 4; IEC 1000-4-6 (EN61000-4-6); IEC 1000-4-8 (EN61000-4-8); IEC 1000-4-11 (EN61000-4-11); ENV 50204: 1996.

Omni Switch/Router Power Requirements

Always make sure that the total power requirements of the modules in your chassis do not exceed the limits of your power supply. To check the power consumption of your configuration, refer to the tables on the following pages and add up the **DC Current Draw** of all modules in your switch. The tables beginning on page 1-19 list modules *without* an HRE-X and the tables beginning on page 1-20 list modules *with* an HRE-X.

The total power consumption of all your modules should be below the current provided by your power supply, which is listed in *OmniS/R-3* on page 1-8 for the OmniS/R-3, *OmniS/R-5* on page 1-10 for the OmniS/R-5 and *OmniS/R-9 and OmniS/R-9P* on page 1-13 for the OmniS/R-9 and OmniS/R-9P. For power consumption and FCC compliance information for Omni Switch/Router VoIP modules, consult your *VoIP User Manual*.

◆ Caution ◆

It is possible, but *not recommended*, to have a configuration in which the current draw of the installed modules exceeds the power provided by a single power supply. However, such a configuration would *require two power supplies and would not allow you to have power redundancy*.

Module Power Requirements *without* an HRE-X

Module	Description	DC Current Draw (Amps)	FCC Class Approval
MPX	Management Processor Module.	3.75	B
ESX-K-100C-32W	Advanced auto-Sensing 10/100 Ethernet module with thirty-two (32) RJ-45 ports.	10.25	B
ESX-K-100FM/FS-16W	Advanced Fast Ethernet (100 Mbps) module with sixteen (16) fiber MT-RJ ports.	9.75	B
GSX-K-FM/FS-2W	Advanced Gigabit Ethernet module with two (2) fiber SC ports.	5.25	B (STP cable) A (UTP cable)
WSX-S-2W	WAN module with 2 serial ports	4.75	B
WSX-SC-4W	WAN module with 4 serial ports	6.25	B
WSX-SC-8W	WAN module with 8 serial ports	8.25	B
WSX-BRI-SC-1W	WAN ISDN module with 1 serial and 1 BRI port	5.75	B
WSX-BRI-SC-2W	WAN ISDN module with 2 serial and 2 BRI ports	7.25	B
WSX-FT1-SC-1W	WAN module with 1 serial and 1 T1 or E1 port	5.75	A
WSX-FE1-SC-1W	WAN module with 1 serial and 1 T1 or E1 port	5.75	B
WSX-FT1-SC-2W	WAN module with 2 serial and 2 T1 or E1 ports	7.25	B
WSX-FE1-SC-2W	WAN module with 2 serial and 2 T1 or E1 ports	7.25	B


Module Power Requirements *with* an HRE-X

Module	Description	DC Current Draw (Amps)	FCC Class Approval
MPX-L3	Management Processor Module.	5.25	B
ESX-K-100C-32W-L3	Advanced auto-Sensing 10/100 Ethernet module with thirty-two (32) RJ-45 ports.	11.75	B
ESX-FM-24W-L3	10 Mbps Ethernet module with twenty-four (24) fiber VF-45 ports	14.5	B
ESX-K-100FM/FS-16W-L3	Advanced Fast Ethernet (100 Mbps) module with sixteen (16) fiber MT-RJ ports.	11.25	B
GSX-K-FM/FS-2W-L3	Advanced Gigabit Ethernet module with two (2) fiber SC ports.	6.75	B (STP cable) A (UTP cable)
WSX-S-2W-L3	WAN module with 2 serial ports	6.25	B (STP cable) A (UTP cable)
WSX-SC-4W-L3	WAN module with 4 serial ports	7.75	B (STP cable) A (UTP cable)
WSX-SC-8W-L3	WAN module with 8 serial ports	9.75	B (STP cable) A (UTP cable)
WSX-BRI-SC-1W-L3	WAN ISDN module with 1 serial and 1 BRI port	7.25	B (STP cable) A (UTP cable)
WSX-BRI-SC-2W-L3	WAN ISDN module with 2 serial and 2 BRI ports	8.75	B (STP cable) A (UTP cable)
WSX-FT1-SC-1W-L3	WAN module with 1 serial and 1 T1 or E1 port	7.25	B (STP cable) A (UTP cable)
WSX-FE1-SC-1W-L3	WAN module with 1 serial and 1 T1 or E1 port	7.25	B (STP cable) A (UTP cable)
WSX-FT1-SC-2W-L3	WAN module with 2 serial and 2 T1 or E1 ports	8.75	B (STP cable) A (UTP cable)
WSX-FE1-SC-2W-L3	WAN module with 2 serial and 2 T1 or E1 ports	8.75	B (STP cable) A (UTP cable)

Grounding a Chassis

Omni Switch/Routers have two grounding screw holes on the back of the chassis. These holes use 10-32 screws and are approximately 1 inch apart. In addition, these holes do not have paint and are surrounded by a small paint-free rectangular section, which provides for a good connection contact.

The figure below shows the location of the grounding screw holes on the back of an OmniS/R-9. They are located approximately four (4) inches from the bottom of the chassis and approximately one (1) inch from the left-hand side of the rear of the chassis.




Grounding Screw Holes on an OmniS/R-9

On an OmniS/R-5, the grounding screw holes are located approximately one (1) inch from the bottom of the chassis and approximately one (1) inch from the left-hand side of the rear of the chassis.

On an OmniS/R-3, they are located approximately four (4) inches from the bottom of the chassis and approximately one (1) inch from the left-hand side of the rear of the chassis.

The Omni Switch/Router Hardware Routing Engine (HRE-X)

The Omni Switch/Router Hardware Routing Engine (HRE-X) is available for the MPX and all Omni Switch/Router switching modules. The HRE-X is a submodule, which plugs into an Omni Switch/Router module, that provides high speed Layer 3 distributed routing for IP and IPX traffic. The HRE-X intercepts frames from the switching logic and determines if a frame should be switched or routed. If a frame needs to be routed, the HRE-X will automatically add the appropriate routing information.



MPX with an HRE-X

The HRE-X has the following restrictions:

- You *must* have Release 3.4.4 software, or later, on your Omni Switch/Router.
- Do *not* install an HRE-X on an MPX unless it is Revision A10, or later.
- Do *not* install an HRE-X on a GSX-FM/FS-4W unless it is Revision B04, or later.

Each HRE-X routes up to 1.5 million packets per second. In an OmniS/R-9 with an HRE-X on every switching module, for example, you could have up to 12 Mpps routed throughput. On a per switch basis, the HRE-X also supports over 256,000 route entries and 64,000 Next Hop destinations.

Valid HRE-X Configurations

You can configure an Omni Switch/Router chassis in one of two ways: with an HRE-X on every single Omni Switch/Router switching module (distributed routing) or a single HRE-X on the MPX (centralized routing).

Distributed Routing. In this configuration, you *must* install an HRE-X on every single switching module in the chassis. In addition, you *cannot* install an HRE-X on the MPX. For example, in an OmniS/R-9 with a single MPX, you would need eight (8) HRE-Xs for all the switching modules. As a general rule, this configuration is recommended in networks of more than four subnets from any one switch.

Centralized Routing. In this configuration, you *must* install the HRE-X on the MPX but not on any Omni Switch/Router switching modules. The HRE-X will perform routing for all Omni Switch/Router switching modules in the chassis. As a general rule, this configuration is recommended for networks of two to four subnets from any one switch.

HRE-X Router Registers versus Feature Limitations

The HRE-X has three (3) registers that can be programmed with a MAC address and mask that allows it to recognize which destination MAC addresses it should act as a router for. IP Routing, Virtual Router Redundancy Protocol (VRRP), ATM Classical IP (CIP), and Channelized DS3 (i.e., M013) utilize at least one of these registers for their operation. This leads to a restriction of the combination of these features that can be supported on an Omni Switch/Router at any given time.

◆ Important Note ◆

ATM and M013 are not supported in Release 4.5.

The HRE-X registers are programmed on a first come, first served basis. Any attempt to program more than three registers fails. In current release, the order which these features program the HRE-X is as follows:

1. ATM CIP
2. IP Routing (**Note:** If there is a second base MAC configured on the MPX, then it will also take a second register.)
3. M013
4. VRRP

For example, if a switch has two base MACs and a CIP group, then no other features can be configured. Any combination of the above features will work given the available HRE-X registers. IP routing always takes one register (two in the dual base MAC case), leaving the other features to compete for the remaining two (one in the dual base MAC case). The other features attempt to program a register only if they are enabled.

◆ Note ◆

ATM CIP is limited to 128 end node route cache entries.

Connecting a DC Power Source to an OmniS/R-PS5-DC375

The OmniS/R-5 can use a DC power supply called the OmniS/R-5-DC375. This power supply contains a female power connector as shown in the figure below. This supply requires the use of 12 gauge wire. A clamp inside each connector keeps the power wire tightly in place during operation. This connector has side screws that can be used to remove the connector.

OmniS/R-PS5-DC375




OmniS/R-5 DC Power Supply Connector Style

Installing DC Power Source Wire Leads

These instructions describe how to connect your 3-wire DC power source to the power connector on your DC power supply. A small flat-tip screwdriver and a wire stripper are required for this procedure.

1. Prepare the three (3) wires—12 gauge—that will plug into the power supply. First, **make sure they are not plugged into the 48-volt power source.**
2. Next, use a wire stripper to carefully strip about a half-inch off the end of each wire, removing the outer insulation to expose the copper core.
3. Twist the loose strands of copper wire together so that they form a tight braid. If possible, solder the entire braid of wire together for better conductivity.
4. Open the wire bay door for one of the three (3) power connector holes. The front of this connector contains a row of square holes. It also contains three (3) circular holes on top that contain screws; you loosen the screws in these holes to open the wire bay doors (square holes) on the connector front so that you can insert the wire lead.
 - a. Insert a small flat-tip screwdriver into one of the top three (3) screw holes.
 - b. Loosen the screw so that the door for the wire bay on the connector front opens.



Opening Wire Bay on Screw-Style Connector


5. Insert the appropriate wire lead into the open circular hole. The silkscreen above each hole indicates which power lead—negative (-), positive (+), or ground (GND)—to plug into which hole. The lead you insert *must* match the lead attached to the 48-volt power source (i.e., negative to negative, positive to positive, ground to ground).

◆ Warning ◆

You must plug DC wire leads into the correct holes in the DC power connector. Use the labels above the DC power connector as a guide to positive, negative, and ground connections.

If you plug wire leads into wrong holes the power supply will not work and could result in damage.

Push the wire in far enough such that it reaches the back wall of the connector, about a half inch inside.



This end would plug into the negative (-) power source. The middle lead would plug into the positive (+) power source and the rightmost lead would plug into the ground (GND).

Inserting the Wire Lead Into the Circular Hole

6. Close the wire bay. Use the small screwdriver (from Step 4a) to tighten the screw above the wire bay into which you inserted the wire lead. The wire lead should be securely attached inside the connector. You should be able to pull on the wire and not dislodge it.


7. Repeat Steps 4 through 6 for the remaining two wire leads. Be sure that the end of each lead attaches to the same power source that you connected to on the power supply (i.e., negative to negative, positive to positive, ground to ground).


Connecting a DC Power Source to an OmniS/R-PS9-DC725

The OmniS/R-9P can use a DC power supply called the OmniS/R-PS9-DC725. This power supply contains a female power connector as shown in the figure below. This supply requires the use of 10 gauge wire. A clamp inside each connector keeps the power wire tightly in place during operation.

OmniS/R-PS9-DC725

GND/(+)/(-)



GND = 

OmniS/R-9P DC Power Supply Connector Style

Installation Requirements


Caution: To reduce the risk of electric shock or energy hazards:

- The branch circuit overcurrent protection must be rated at a minimum of 30 A (amperes) for the OmniS/R-9P PS9-DC725.
- Use 10 gauge (AWG - American Wire Gauge) solid copper conductors only for the OmniS/R-9P PS9-DC725.
- A readily-accessible disconnect device that is suitably approved and rated shall be incorporated in the field wiring.
- This device is to be installed in a restricted access area in accordance with the NEC (National Electrical Code) or the authority having jurisdiction.
- Connect this device to a reliably grounded SELV (Safety Extra Low Voltage) or a centralized DC source.

Installing DC Power Source Wire Leads

These instructions describe how to connect your 3-wire DC power source to the power connector on your DC power supply. A small flat-tip screwdriver and a wire stripper are required for this procedure.

1. Prepare the three (3) wires—10 gauge—that will plug into the power supply. First, **make sure they are not plugged into the 48-volt power source.**
2. Next, use a wire stripper to carefully strip about a half-inch off the end of each wire, removing the outer insulation to expose the copper core.
3. Twist the loose strands of copper wire together so that they form a tight braid. If possible, solder the entire braid of wire together for better conductivity.
4. Open the wire bay door for one of the three (3) power connector holes. The front of the power connector contains a row of square holes. It also contains three (3) circular holes (located directly above the square holes) that contain screws; you loosen the screws in these holes to open the wire bay doors (square holes) on the connector front so that you can insert the wire leads into the power connector.
 - a. Insert a small flat-tip screwdriver into one of the three (3) screw holes.
 - b. Loosen the screw so that the door for the wire bay on the connector front opens.



Opening Wire Bay on DC Power Supply Connector

5. Insert the appropriate wire lead into the open circular hole. The silkscreen above each hole indicates which power lead—ground (GND), positive (+), or negative (—)—to plug into which hole. The lead you insert *must* match the lead attached to the 48-volt power source (i.e., ground to ground, positive to positive, negative to negative).


◆ Warning ◆

You *must* plug DC wire leads into the correct holes in the DC power connector. Use the labels above the DC power connector as a guide to ground, positive and negative connections.

If you plug wire leads into the wrong holes, the power supply will not work and could result in damage.

Push the wire in far enough so that it reaches the back wall of the connector, about a half inch inside.

GND/(+)/(-)



This end would plug into the ground (GND). The middle lead would plug into the positive (+) power source and the rightmost lead would plug into the negative (-) power source.

Inserting the Wire Lead Into the Circular Hole

6. Close the wire bay door. Use the small screwdriver (from Step 4a) to tighten the screw above the wire bay into which you inserted the wire lead. The wire lead should be securely attached inside the connector. You should be able to pull on the wire and not dislodge it.
7. Repeat Steps 4 through 6 for the remaining two wire leads. Be sure that the end of each lead attaches to the same power source that you connected to on the power supply (i.e., ground to ground, positive to positive, negative to negative).

Replacing Power Supplies (9-Slot Chassis)

If a power supply ever needs to be replaced in an Omni Switch/Router 9-slot Chassis (e.g., OmniS/R-9 or OmniS/R-9p), it is strongly recommended that power supplies not be mixed, except under the conditions and exceptions shown in the following table.

◆ **Note** ◆

In all cases, swapping operations must be made with the power switch of the replacement power supply turned OFF. Failure to turn the power switch off during the swapping operation may cause the data switch to reset and restart.

Replacing Power Supplies (9-Slot Chassis)

If One of Two Power Supplies Fails	Revision	Replace	With
650-watt	Pre-M1	Both Power Supplies	Two 650-watt (Revision M1+) or two 725-watt Power Supplies
650-watt	M1 or later	Failed Power Supply	One 650-watt (Revision M1+) or one 725-watt Power Supply
725-watt	Any	Failed Power Supply	One 725-watt Power Supply

2 The Omni Switch/Router MPX

Omni Switch/Router Management Processor Module (MPX) Features

The MPX provides such system services as maintenance of user configuration information, downloading of switching module software, basic bridge management functions, basic routing functions, the SNMP management agent, access to the User Interface software, and Advanced Routing. In addition, the MPX can operate in a redundant configuration with another MPX.

◆ Important Note ◆

If you have a single MPX in your chassis, it *must* be installed in Slot 1.

With the optional HRE-X, which is described in Chapter 1, “Omni Switch/Router Chassis and Power Supplies,” you can increase routing performance to 1.5 million packets per second.

MPX Technical Specifications	
Flash Memory	8 MB (32 MB maximum); 16 MB required for Release 4.4 and later
SIMM (DRAM) Memory	32 MB (128 MB maximum); 64 MB required for Release 4.4 and later
SDRAM Memory	16 MB
MAC Addresses Supported	4096
Switching Backplane	Up to 22 Gbps (aggregate) switching fabric capacity
Serial Ports	2 (1 male DB9 modem connector and 1 female DB9 console connector)
Ethernet (10 Mbps) Switch Management Ports	1 copper RJ-45 or fiber (ST) port for switch management functions.
Current Draw	3.75 amps without an HRE-X 5.25 amps with an HRE-X

◆ Warning ◆

Do *not* install any version of the MPM (i.e., MPM-C, MPM-1G, MPM-II, MPM-III, or original MPM) in a chassis with an MPX or any OmniSwitch switching module. Installing an MPM in a chassis with an MPX can cause physical damage.

Omni Switch/Router Management Processor Module (MPX) Features

Warning Label. This label indicates that the module contains an optical transceiver (on the MPXs with fiber ST Ethernet ports only).

OK1 (Hardware Status). This dual-state LED is on Green when the MPX has passed power-on hardware diagnostics successfully. On Amber when the hardware has failed diagnostic tests. If the **OK1** LED is alternating Green and Amber, then file system compaction is in progress.


Caution

Do not power down the Omni Switch/Router or insert any modules while the **OK1** LED is alternating Green and Amber. If you do, file corruption may result and you will not be able to restart the switch.

OK2 (Software Status). Blinking Green when the MPX has successfully loaded software to the switching modules. Blinking Amber when the MPX is in a transitional state, such as when it first boots up. If the **OK2** LED blinks Amber for an extended period of time (i.e., more than a minute), then you should reboot the switch.

Caution

Do not insert or remove any modules while the MPX **OK2** LED is blinking Amber. If you do, file corruption may result and you will not be able to restart the switch.



Label. This label will indicate the Ethernet management port type. It will read either **MPX 10 mm** (multimode fiber Ethernet port) or **MPX 10** (copper RJ-45 Ethernet port).

Module Status LEDs

Module Status LEDs

PS1 (Power Supply 1 Status). This dual-state LED is on Green when the switch is receiving the proper voltage from Power Supply 1. It is on Amber when Power Supply 1 is on, but not supplying the correct amount of voltage to power the switch, or is installed and turned off. The **PS1** LED is Off when the Power Supply 1 is not present.


PS2 (Power Supply 2 Status). This dual-state LED is on Green when the Omni Switch/Router is receiving the proper voltage from Power Supply 2. It is on Amber when Power Supply 2 is on, but not supplying the correct amount of voltage to power the switch, or is installed and turned off. The **PS2** LED is Off when Power Supply 2 is not present.

TEMP (Temperature). On Yellow to warn that the internal switch temperature is approaching maximum operating limits. Note that this LED comes on *before* the temperature limit is reached.

PRI (Primary MPX). On Green when this MPX is the active, or controlling, MPX. It is also on Green when this is the only MPX installed in the switch.

SEC (Secondary MPX). On Green when this MPX is the secondary MPX in a redundant MPX configuration. As the secondary MPX, this module is in hot standby mode.

Omni Switch/Router Management Processor Module (MPX) Status LEDs



MPX Management Connectors

MPX Serial and Ethernet Management Ports


You can gain access to switch management software through one of the two serial (RS-232) ports on the MPX or the Ethernet management port. The two serial ports are configured with 9-pin “D” connectors (DB-9) per the IBM AT serial port specification. One port, called the “modem” port, is male and the other, called the “console” port, is female. See *MPX Management Connectors* on page 2-3 for illustrations of these ports.

The modem port is a Data Terminal Equipment (DTE) connector, which is typically connected to a modem. You can also connect directly from this port to a PC or terminal with a standard null-modem cable available in most computer equipment stores.

◆ **Note** ◆


The modem port is hard-wired for DTE communication; you do not need to set any jumpers.

The console port is a Data Communication Equipment (DCE) connector, which can be directly connected to a PC, terminal, or printer.



MPX Console Port Specifications		
Pin Number	Standard Signal Name	Direction
1	Not Used	
2	RD	From MPX
3	TD	To MPX
4	Not Used	
5	GND	
6	Not Used	
7	Not Used	
8	Not Used	
9	Not Used	
Shell	Shield GND	

MPX Console Port



MPX Modem Port Specifications		
Pin Number	Standard Signal Name	Direction
1	Not Used	
2	RD	To MPX
3	TD	From MPX
4,	DTR	From MPX
5	GND	
6	DSR	To MPX
7	RTS	From MPX
8	CTS	To MPX
9	Not used	
Shell	Shield GND	

MPX Modem Port

Ethernet Management Port

The MPX also supports an out-of-band Ethernet port for high-speed uploads and switch management functions. With this port, you can access the Omni Switch/Router over a network via Telnet or FTP.

You can use the Boot prompt to configure an IP address for the Ethernet management port or you can use the **ethernetc** command, which is described in Chapter 6, “Configuring Management Processor Modules.” After you have assigned an IP address to the Ethernet management port, you can use it to Telnet into the UI.

See Appendix A, “The Boot Line Prompt,” for documentation on configuring the Ethernet management port with the boot prompt.

◆ Important Note ◆

On some revisions of the MPX, you *must* configure the Ethernet management port with the boot prompt before you can use the **ethernetc** command.

See the table on the following page for available Ethernet management port types.

MPX Model	Ethernet Management Port Type (Cable Type)	Max. Cable Distance
MPX-T	RJ-45 (UTP)	100 meters
MPX-FL	ST (Multimode fiber)	2 kilometers

Configuring MPX Serial Ports

The serial communications parameters for the two MPX serial ports are set by default to the following:

- 9600 bits per second (bps)
- 8 data bits
- 1 stop bit
- no parity
- no hardware flow control (Windows 95)

Each serial port supports serial data rates of 1200, 9600, 19200, and 38400 bps. However, you must remove the default baud rate shunt (E1), which fixes the baud rate at 9600 bps, before you can change the baud rate. This shunt is located near the front end of the MPX's circuit board, just to the right of the Ethernet management port.

To change the serial port configuration parameters, use the **ser** command, which is described in detail in Chapter 6, "Configuring Management Processor Modules."

Flash Memory and Omni Switch/Router Software

Flash memory on the MPX holds the Omni Switch/Router's executable images and configuration data. When a switching module comes online, the MPX downloads the appropriate image file for that module to that module's memory. Image files (those with the **img** extension) contain executable code for different switching modules and software features.

The following table lists Omni Switch/Router image files that may be present in MPX flash memory along with the module(s) or feature with which the file is used.

File Name	Modules/Function Used With
mpx.img mpx.cmd mpm.cfg mpm.cnf	MPX
desx.img	Ethernet port stress test software
diagx.img	Diagnostics software
esx.img	All GSX and ESX modules
fwdx.img	IP Fastpath and Firewall software
gated.img	Advanced Routing software
ipcntrl.img	IP control software
ipms.img	IPMS software
isdn.img	WSX-BRI-SC
mrd.img	Advanced Routing software
ntp.img	Network Time Protocol (NTP) software
policy.conf	PolicyManager file comprised of a MAC address and time that uniquely identifies the switch(es) to which the policy applies
policy.img	PolicyView software
qos.img	Quality of Service (QOS) software
rav.img	RADIUS authentication software
t1e1drv.img	WSX-FT1/E1-SC
text_cfg.img	Text-based configuration software
vrrp.img	VRRP software
vsmboot.asc	Boot file for Voice Over IP (VOIP) modules
vsx.img	Voice Over IP (VOIP) modules
web.img	HTTP browser client software
wsx.img	WSX-S-2W, WSX-SC-4W, WSX-SC-8W (Frame Relay and PPP software)

Flash Memory Guidelines

The switch alters flash memory contents when a software command requests a configuration change, when a remote administrator downloads a new executable image, or when the switch fails and a record of the failure is written to flash memory. These operations require available space in flash memory.

In general the flash memory on the switch should always have at least 75000 bytes available at all times. In a switch with 8 MB of flash memory, for example, the images in flash should never exceed 7.45 MB. (You can view how much flash memory is available through the **ls** command.) This will allow enough room in flash for booting and configuration file expansions. If your flash memory exceeds this amount, then you need to delete some images from flash.

In addition, the flash file system has a limit of 256 files, including configuration, logging, and other files. When this 256-file limit is reached, configuration file expansions will cease and new files will not be able to be loaded. This file limit applies even if there is enough memory available in flash.

Not all image files in flash memory are required—only those that must be used with the switching modules in your Omni Switch/Router. You can remove any files that are not required for your Omni Switch/Router configuration by using the **rm** command. For example, if you do not have T1/E1 ports, you could remove the **t1e1drv.img** file.

MPX Redundancy

In order to provide greater reliability, Omni Switch/Router supports two MPXs in a primary/secondary redundant configuration. If the primary MPX fails, the secondary MPX takes over without any operator intervention.

◆ Warning ◆

Do *not* install any version of the MPM (i.e, MPM-C, MPM 1G, MPM II, or original MPM) in a chassis with an MPX. Installing an MPM in a chassis with an MPX can cause physical damage. If you want to configure an Omni Switch/Router chassis in a redundant configuration, you *must* use two MPXs.

When you have two MPXs in one chassis, they must be installed in Slots 1 and 2, and only one can be active. MPXs will assume one of the following roles.

- Primary - The MPX that is currently active and processing commands. It is also the MPX that is communicating via Telnet, FTP, etc.
- Secondary - An MPX that is currently not the primary. It has sufficient software to communicate with the primary MPX. (For full redundancy, the secondary MPX should also have the same software version as the primary and its configuration should be in sync with the primary.) In this state, it is capable at any time of assuming the primary role.

The LEDs on each MPX reflect the same status with the exception that the primary's **PRI** LED is on whereas the secondary's **SEC** LED is on. Also, the secondary MPX's **OK2** LED will not flash amber during board transitions. See *Omni Switch/Router Management Processor Module (MPX) Status LEDs* on page 2-2 for locations of the LEDs.

◆ Important Note ◆

To support redundancy, your MPX *must* be Revision A14 or higher.

Change-Over Procedure

The secondary MPX continuously monitors the primary MPX. This monitoring serves two purposes: 1) to notify the secondary MPX that the primary is alive and processing, and 2) to update the configuration and thus keep the two MPXs in sync. If the secondary MPX detects that the primary is no longer operational, it will begin to take over as primary. When a secondary MPX becomes primary it resets all the other modules in the chassis and performs a primary MPX initialization.

There are four states for an MPX configuration. You can view the current MPX state through the **slot** command. These states are described in the table below. Note that for a primary/secondary configuration to be in a “redundant” state, the relationship between the two MPXs must meet the conditions shown in the table.

MPX State	Requirement for State
Redundant	Both MPXs are running the same version of software and the configurations are in sync.
Configuration Fallback	Both MPXs are running the same version of software but the configurations are different.
Software Fallback	The MPXs are running different versions of software, and their configurations may be the same or different.
None	There is only one MPX installed in the chassis.

The primary MPX has the ability to transfer files to and from the secondary MPX. In the condition where the secondary MPX has an older version of software (Software Fallback), it is not desirable to update the configuration file of the secondary. It is therefore the default not to update the configuration file on the secondary if the secondary is running an earlier version of software. You can force the update using appropriate commands in the **mpm** menu. (See Chapter 6, “Configuring Management Processor Modules,” for more information on commands in the **mpm** menu.)

◆ **Note** ◆

Do *not* remove a primary MPX without performing a **renounce** command (described in Chapter 6, “Configuring Management Processor Modules”) first.

MPX Redundancy Commands

A set of commands exists to monitor the primary and secondary MPXs. These commands are covered in detail in Chapter 6, “Configuring Management Processor Modules.” Note that you can attach a terminal to both MPXs in a chassis; however, you will see a different responses depending on which is primary and which is secondary. You should execute all UI commands from the primary MPX except for those commands specifically addressing the secondary MPX. For example, commands are available to control and monitor the secondary MPX from the primary MPX (e.g., the **sls** command lists files on the secondary MPX from the primary MPX).

3 Omni Switch/Router Switching Modules

Omni Switch/Router switching modules perform software filtering, translations between dissimilar network interfaces, and hardware-based switching. Omni Switch/Router switching modules have an additional on-board interface connector for the HRE-X.

Currently, Omni Switch/Router switching modules consist of Gigabit Ethernet modules, auto-sensing 10/100 Ethernet modules, Fast (100 Mbps) Ethernet modules, 10 Mbps Ethernet modules, Voice Over IP (VOIP) modules, and WAN modules.

◆ Important Note ◆

Omni Switch/Router modules require the use of an Omni Switch/Router chassis (see Chapter 1, “Omni Switch/Router Chassis and Power Supplies”). Do *not* install an Omni Switch/Router module in an OmniSwitch chassis and do *not* install an OmniSwitch module in an Omni Switch/Router chassis.

Gigabit Ethernet Modules

- GSX-K-FM/FS/FH-2W Advanced 2-port Gigabit Ethernet switching module

10/100 Ethernet Modules

- ESX-K-100C-32W Advanced 32-port auto-sensing 10/100 Ethernet switching module

Fast (100 Mbps) Ethernet Modules

- ESX-K-100FM/FS-16W Advanced 16-port Fast Ethernet (100 Mbps) switching module

WAN Modules

- WSX-S-2W 2 serial ports that support the frame relay or PPP protocol.
- WSX-SC-4W/8W 4 or 8 serial ports that support the frame relay or PPP protocol.
- WSX-FT1/E1-SC-1W/2W 1 or 2 T1/E1 ports and one or two serial ports that support the frame relay or PPP protocol
- WSX-BRI-SC-1W/2W 1 or 2 UPS (Universal Serial Port) and 1 or 2 ISDN-BRI ports that support Frame Relay or PPP

Voice Over IP Modules

Voice Over IP (VOIP) modules for the Omni Switch/Router are listed below and are documented in the *VoIP User Manual*.

- VSX-A 4, 6, 8, 14, or 16 analog RJ-11 ports supporting FXS and FXO interfaces, including T.38 FAX
- VSX-VSD 2 or 4 digital T1 or E1 (Euro PRI and Qsig) ports, including T.38 FAX

Omni Switch/Router Hardware Routing Engine

The HRE-X offers high-speed Layer 3 switching from 1.5 to 12.0 million packets per second (Mpps) in a fully loaded chassis. See Chapter 1, “Omni Switch/Router Chassis and Power Supplies,” for more information on the HRE-X.

◆ Important Note ◆

Omni Switch/Router switching modules require an MPX. You cannot install any version of the MPM (i.e., MPM-III, MPM-C, MPM-1G, MPM-II, or original MPM) in a chassis with an MPX. See Chapter 2, “The Omni Switch/Router MPX,” for more information on the MPX.

Required Image Files

See the table below for the required images files for the MPX and switching modules. You *must* load the image file (or files) listed for the corresponding module or it will not run.

Required Image Files

Module	Image File(s)
MPX	mpx.img, fpx.img
ESX-K-100C-32W	esx.img
ESX-K-100FM/FS-16W	esx.img
GSX-K-FM/FS/FH-2W	esx.img
VSX-VSA	vsx.img, text_cfg.img, vsmboot.asc
VSX-VSD	vsx.img, text_cfg.img, vsmboot.asc
WSX-S-2W	wsx.img
WSX-SC-4W	wsx.img
WSX-SC-8W	wsx.img
WSX-BRI-SC-1W/2W	wsx.img, isdn.img
WSX-FT1-SC-1W/2W	wsx.img, t1e1drv.img
WSX-FE1-SC-1W/2W	wsx.img, t1e1drv.img

Installing a Switching Module

All switching modules can be inserted and removed from the switch chassis while power is on or off without disrupting the other modules. *A standard screwdriver is required for installing and removing switching modules.* You can also hot swap modules of the same type while the switch is active.

Switching modules may be installed in any slot other than Slot 1. (Slot 1 is reserved for an MPX.) In a setup with redundant MPX modules, Slots 1 and 2 are reserved for the MPXs. Additional modules can be installed in any available slot. (OmniS/R-3 slots are numbered 1 to 3 starting from the topmost slot. OmniS/R-5 slots are numbered 1 to 5 starting from the topmost slot. OmniS/R-9 slots are numbered 1 to 9 starting from the left.)


◆ Anti-Static Warning ◆

Before handling a switching module, free your hands of static by wearing a grounding strip, or by grounding yourself properly. Static discharge can damage the components on the switching module.


To insert a switching module follow these instructions:

1. Holding the module firmly in both hands, carefully slide it into the card guide. The front panel connectors and LEDs should face outward. In a 9-slot Omni Switch/Router, the component side of the board should face right (toward the power supply). In a 3- or 5-slot Omni Switch/Router, the component side should face up.


The module should slide in easily. A large amount of force is not necessary and should not be used. If any resistance is encountered, check to be sure that the module is aligned properly in the card guide.



2. Once the module is in the slot, close the two card ejectors (one on each end of the module) by pressing them in toward the module until they snap into place.



3. Use a standard screwdriver to tighten the two screw fasteners to secure the module inside the chassis. The screws should be tight enough such that a screwdriver would be necessary to loosen the screws.



Removing a Switching Module

To remove a switching module, follow the instructions below. If you are “hot swapping” the modules (i.e., removing and inserting while power is on), see *Hot Swapping a Switching Module* on page 3-7.

◆ Anti-Static Warning ◆

Before handling a switching module, free your hands of static by wearing a grounding strip, or by grounding yourself properly. Static discharge can damage the components on your switching module.

1. Loosen the screw fasteners at the top and bottom of the switching module using a standard screwdriver.
2. Gently unlock the two card ejectors by pulling them out away from the module.
3. With both hands, carefully pull the module free of the chassis enclosure.

Hot Swapping a Switching Module

You may remove and insert switching modules while the switch is running. This technique is referred to as “hot swapping.” When you hot swap, you must replace the module with the same module type as the one you removed. For example, if you remove an ESX switching module you must replace it with another ESX switching module.

◆ Note ◆

You *cannot* hot swap a module into a previously empty slot. To use an empty slot, you *must* power down your chassis.

Perform the following steps to safely hot swap a switching module. (You cannot hot swap a primary MPX module.) Since this procedure could possibly disrupt the network, it is best to hot swap during network down times.

1. At the system prompt, enter

```
swap on <minutes>
```

where **minutes** is the number of minutes you want the switch to be in swap mode (the default is 5 minutes). A message similar to the following will be displayed.

```
Swap is ON for 5 minutes
```

The swap mode *must* be enabled (**ON**) to insert a switching module. If not, the system may halt or restart. (See Chapter 6, “Configuring Management Processor Modules,” for more information on the **swap** command.)

◆ Caution ◆

Modules can only be reset and hot-swapped when the MPX’s **OK2** light is in its normal flashing green state.

2. Enter **reset**, followed by the slot number of the switching module you want to hot swap, then followed by the word **disable**. (See Chapter 36, “Running Hardware Diagnostics,” for more information on the **reset** command.) For example, if you want to hot swap the switching module in slot 4, you would enter

```
reset 4 disable
```

at the system prompt. Next, the switch will prompt you to confirm the reset. The following is an example of the display for an ESX module. The display for other types of switching modules will be similar.

```
Resetting slot of type F-Ether/M may crash system
Attempt reset anyway {Y/N}? (N) :
```

Press **y** and then press **<Enter>**. If the switching module is in slot 4, a message similar to the following will be displayed.

```
resetting slot 4 to disable
```

3. The MPX’s **OK2** LED will flash amber 1 or 2 times, then return to normal flashing green. The switching module’s **OK1** LED will turn amber and the **OK2** LED will *not* be illuminated. Remove all cables attached to ports on the switching module that you are going to swap out.

- Carefully remove the switching module from the chassis and put it in a safe place. (See *Removing a Switching Module* on page 3-6 for instructions on removing a switching module.) The MPX's **OK2** LED will flash amber 1 or 2 times, then return to normal flashing green. In addition, the swap time will reset to its original value. (For example, if you set the swap time to 15 minutes in step 1, you will have 15 minutes again, regardless of how much time has elapsed.)

◆ **Warning** ◆

Removing or inserting the switching module while the MPX's **OK2** LED is flashing amber can cause the system to reset.

- Carefully insert the new switching module into the chassis. (See *Installing a Switching Module* on page 3-4 for instructions on inserting a switching module.)

◆ **Caution** ◆

When re-installing a module during a hot swap, it must make a proper connection to the switch backplane. The connection is made when you close the card ejectors. Always close the card ejectors firmly and briskly, without hesitation. Closing them too slowly can cause the switch to halt or restart.

The MPX's **OK2** LED will flash amber 1 or 2 times, then return to normal flashing green. If, after hot-swapping modules, the MPX's **OK2** LED continues to flash amber for more than about 8 seconds, it means that the switch needs to be reset.

The swap time will again reset to its original value.

- Re-insert the cables that were removed in step 3 into the new switching module.
- Enter **reset** followed by the slot number for the new switching module. For example, if the new switching module is in slot 4, you would enter

reset 4

at the system prompt. Next, the switch will prompt you to confirm the reset. The following is an example of the display for an ESX module. The display for other types of switching modules will be similar.

```
Resetting slot of type F-Ether/M may crash system  
Attempt reset anyway {Y/N}? (N) :
```

Press **y** and then press **<Enter>**. If the switching module is in slot 4, a message similar to the following will be displayed.

```
resetting slot 4 to enable
```

- The MPX's **OK2** LED will flash amber 1 or 2 times, then return to normal flashing green. The switching module's **OK1** LED will turn from amber to solid green and the **OK2** LED will be blinking green. If the **OK1** LED on the switching module is amber, then the hardware has failed diagnostics or the corresponding image file for the module is not in flash memory. If the **OK2** LED on the switching module is solid amber, then the module failed to download software from the MPX.

9. If the hot swapping mode has not timed out, enter

```
swap off
```

at the system prompt. Something like the following will then be displayed.

```
Swap is OFF, timeout is 5 minutes  
usage swap { ON [ minutes ] | OFF [ minutes ] }
```

Diagnostic Tests

All switching modules are subjected to extensive power-on diagnostics during the Power-On Self-Test cycle (POST). These diagnostics are designed to be as extensive as possible without causing disruption to external networks or requiring special test connections. While the diagnostics are running, the MPX **OK2** LED will be flashing green. LEDs on the switching module can provide information on the success or failure of these tests. Also refer to Chapter 35, "Troubleshooting," for information on error conditions reflected in the LED displays.

More extensive diagnostic tests are available for off-line testing of switching modules. See Chapter 36, "Running Hardware Diagnostics," for further information.

Handling Fiber and Fiber Optic Connectors

Using fiber is extremely simple, but a few important rules should always be followed:

Step 1. Use Premium Grade Jumper Cables with Duplex SC Connectors

There are many brands of fiber optic jumper cables, with a wide range of quality between each manufacturer. Premium cables do three things well:

- They provide a good polish on the fiber optic connector endface (where the light exits the cable). Endface geometries must be exceptionally precise and aligned to extremely tight tolerances. The better the endface geometry, the lower the loss and more consistent the connection. Poor connector interfaces will reflect light back into the laser, causing an increase in laser noise.
- They mate well with other connector interfaces. Chances are the manufacturer of the jumper cable will not be the same as the manufacturer of the transceiver connector interface. Premium jumper cables mechanically align themselves well into most transceiver interfaces. This provides both better performance as well as better repeatability. You will always see a variance in transceiver power due to connector alignment, often as much as 0.3 to 0.7 dB. Good jumper cables help reduce this variance.
- They continue to mate well after many insertions and removals. Premium grade jumper use premium grade connectors that maintain their mechanical integrity up to and beyond 2000 insertion cycles.

For better repeatability, always use duplex (two connectors fused together and terminated to two cables) SC connectors on your jumper cables when connecting to a fiber-optic transceiver. Two simplex connectors inserted into a transceiver interface will often have up to 3 dB greater variation in repeatability compared to duplex connectors.

Never bend the fiber optic cable beyond its recommended minimum bend radius (1.2 inches minimum). This introduces bend losses and reflections that will degrade the performance of your system. It can also damage the fiber, although fiber is much tougher than most would assume. Still, it is highly recommended to buy only jumper cables with 3mm Kevlar jacketing, which offer superior protection and longer life.

Step 2. Keep Your Fiber Optic Connectors Clean

Unlike electrical connectors, fiber-optic connectors need to be extremely clean to ensure good system performance. Microscopic particles on the connector endface (where the light exits the connector) can degrade the performance of your system, often to the point of failure. If you have low-power output from a fiber-optic transceiver or a fault signal from your equipment, cleaning your fiber-optic connectors should always be done before trouble shooting.

Follow the steps below to clean your fiber optic connector:

1. Hold the connector cleaner tool in the palm of your left hand and, with the silver shutter upwards, rotate the cloth-forwarding lever (located on the right side of the tool) with your thumb away from your body. As the lever winds the cleaning cloth inside the case, it simultaneously opens the silver shutter located at the top of the unit.

2. Keeping your thumb pressed on the cloth-forwarding lever, press the optical plug ferrule endface against the cleaning cloth and drag the plug down toward your body (there should be arrows on the top of the tool that indicate the proper wiping direction). The connector is now clean.
3. Release the cloth-forwarding lever, allowing it to return to its initial position.

A cleaning cloth reel can enable over 400 cleanings and is replaceable. When cables are not being used, always put the plastic or rubber endcaps back on the connector to ensure cleanliness.

Step 3. Keep the Transceiver Interface Clean

If you have cleaned your connectors, but still experience low-power output from a fiber-optic transceiver or a fault signal from your equipment, you should clean the transceiver interface by blowing inert dusting gas inside the transceiver interface. This removes dust and other small particles that may block the optical path between the optics of the transceiver and the connector's endface.

Step 4. Attenuate Properly

Often equipment using laser-based transceivers need to have the optical path attenuated when performing loop-back testing or testing between two pieces of equipment. Too much optical power launched into the receiver will cause saturation and result in system failure. If you are using single mode fiber and you do not know the power output of the laser, it is always best to use a 10 dB attenuator when testing. Using the wrong type of attenuator will introduce problems, most notably reflection of light back into the laser, often resulting in excess noise and causing system failure.

Inline attenuators eliminate the need for additional jumper cables and thus reduce the number of connection interfaces. This increases the integrity of the optical path resulting in a more accurate test.

Gigabit Ethernet Modules

Gigabit Ethernet connections can be used as network backbones or in a wiring closet. The following Omni Switch/Router Gigabit Ethernet modules are available:

- **GSX-K-FM/FS/FH-2W** Advanced switching module with two (2) Gigabit Ethernet backbone connections using fiber (SC) connectors.

This module is described and illustrated in the following sections.

◆ Note ◆

Wait at least five (5) seconds after a cable is pulled from a GSX module before reinserting it. This will prevent packets from being dropped.

GSX-K-FM/FS/FH-2W

The GSX-K-FM/FS/FH-2W Gigabit Ethernet backbone switching module contains two fiber SC connectors that support two fully switched 1000Base-LX (long-distance fiber transmissions) or 1000Base-SX (short-distance fiber transmission ports). The GSX-K-FM/FS/FH-2W can be used as a backbone connection in networks where Gigabit Ethernet is used as the backbone media.

The GSX-K-FM/FS/FH-2W can be factory configured with intermediate-reach single mode or multimode fiber ports (see *GSX-K-FM/FS/FH-2W Technical Specifications* on page 3-13 for more information). The intermediate-reach single mode version is referred to as the GSX-K-FS-2W; the long-reach single mode version is referred to as the GSX-K-FH-2W; and the multimode version is referred to as the GSX-K-FM-2W.

The ports are color coded to differentiate the mode: multimode connectors are black, long-haul single mode connectors are yellow, and intermediate-reach single mode connectors are blue. (See *Handling Fiber and Fiber Optic Connectors* on page 3-10 for proper handling of SC connectors and fiber-optic cable.)

The GSX-K-FM/FS/FH-2W takes advantage of new Gigabit Ethernet/Fast Ethernet ASIC technology known as “Kodiak.” This module provides 4 priority levels and 256 queues per Kodiak ASIC.

◆ Note ◆

Kodiak-based modules support up to 4 levels of priority (0-1, 2-3, 4-5, 6-7). This is *not* compatible with the implementation of VLAN priority of Mammoth-based modules. Kodiak based priority VLANs can only be used with other Kodiak based priority VLANs.

With the optional HRE-X you can increase routing performance to 1.5 million packets per second per module and up to 12 Mpps in a fully-loaded 9-slot chassis.

GSX-K-FM/FS/FH-2W Technical Specifications	
Number of ports	2
Connector Type	SC
Standards Supported	802-3z, 1000Base-LX, and 1000Base-SX
Data Rate	1 Gigabit per second (full duplex)
Maximum Frame Size	1,518 bytes
MAC Addresses Supported	8,192
Connections Supported	1000Base-LX or 1000Base-SX connection to backbone or server
Cable Supported	Multimode and single mode
Output Optical Power	-9.5 to -4 dBm (Multimode) -9.5 to -3 dBm (Intermediate-reach single mode) 0 to +5 dBm (Long-reach single mode)
Input Optical Power	-17 to 0 dBm (Multimode) -20 to -3 dBm (Intermediate-reach single mode) -24 to -3 dBm (Long-reach single mode)
Cable Distance	Multimode fiber: \approx 220 m Intermediate-reach single mode fiber: \approx 10 km Long-reach single mode fiber: \approx 70 km
Current Draw	5.25 amps without an HRE-X 6.75 amps with an HRE-X

◆ Special Note ◆

The single mode version of this module has been deemed:

CLASS 1 LASER PRODUCT
LASER KLASSE 1
LUOKAN 1 LASERLAITE
APPAREIL A LASER DE CLASSE 1

to IEC 825:1984/CENELEC HD 482 S1.

Warning Label. This label indicates that the module contains an optical transceiver.

This Gigabit Ethernet module includes one row of LEDs for each port. The LEDs for a given port display in the row labeled with the port number. Definitions for the LEDs are given below.

RX (Receive). On Green when the corresponding port is receiving data.

TX (Transmit). On Green when the corresponding port is transmitting data.

LINK (Link Status/Disabled). On Green when the corresponding port has a valid physical link and a signal is present. Under normal conditions, this LED should always be on when a cable is connected.

Port LEDs


Module LEDs

Module Label. This label will indicate the GSX-K-FM/FS/FH-2W type. It will read either **GSX-K mm** (multimode cable), **GSX-K sm** (intermediate-reach single mode cable), or **GSX sm K long reach** (long-reach single-mode cable).

OK1 (Hardware Status). On Green when the module has passed diagnostic tests successfully. On Red when the hardware has failed diagnostics.

OK2 (Software Status). Blinking Green when the module software was downloaded successfully and the module is communicating with the MPX. Blinking Red when the module is in a transitional state. On solid Red if the module failed to download software from the MPX.

SC connectors will be color coded to indicate multimode (Black) or intermediate-reach single mode (Blue).



2-Port Advanced Gigabit Ethernet Switching Module

Auto-Sensing 10/100 Ethernet Modules

Alcatel's Omni Switch/Router 10/100 Ethernet modules can be used to connect networks with a mix of 10 Mbps and 100 Mbps workstations or as a network backbone.


The following Omni Switch/Router 10/100 and Fast Ethernet modules are available:

- ESX-K-100C-32W Advanced switching module with thirty-two (32) auto-sensing 10/100 Mbps desktop connections using RJ-45 ports.

This module is described and illustrated in the following sections.

Ethernet RJ-45 Pinouts

The figure and table below illustrate the pinouts used on RJ-45 ports in Omni Switch/Router 10/100 Ethernet modules.



Ethernet RJ-45 Specifications	
Pin Number	Standard Signal Name
1	RD +
2	RD -
3	TD +
4,	Not Used
5	Not Used
6	TD -
7	Not Used
8	Not Used

ESX-K-100C-32W

The ESX-K-100C-32W Omni Switch/Router 10/100 Ethernet switching module contains 32 ports that each support a fully switched 10 or 100 Mbps connection in full- or half-duplex mode. This module offers high density 10/100 connectivity for desktop connections. Each port can auto-sense the connection speed and automatically switch at that speed. You configure whether you want to use the auto-sensing functionality through the **10/100cfg** command.

By default, each port is configured to operate in half-duplex, auto-sensing mode. You can configure full-duplex mode on each port through **10/100cfg**. Auto-sensing may be disabled to allow you to manually configure ports through the **10/100cfg** command. An additional software command, **10/100vc**, allows you to view the current line speed and link mode of each port connection. The **10/100cfg** and **10/100vc** commands are described in Chapter 15, "Managing Ethernet Modules."

The 32 RJ-45 ports may connect to unshielded or shielded twisted pair (UTP) cable (see *ESX-K-100C-32W Technical Specifications* on page 3-17 for more information). Each port may connect to a single high-speed device or a hub serving multiple devices. The ESX-K-100C-32W can be used in the wiring closet with a mix of 100 Mbps Ethernet devices and 10 Mbps Ethernet devices that are transitioning to higher speed connections.

Module ports are divided into four (4) banks of eight (8) ports. Ports are numbered from 1 to 8 within each of the four banks. The four banks are labelled **A**, **B**, **C**, and **D**. This grouping simplifies the display of LEDs, which are organized as a matrix (see *32-Port Advanced Auto-Sensing 10/100 Ethernet Switching Module* on page 3-18). Software commands will number these ports 1 through 32, with Port **A1** as 1, Port **B1** as 9, **C1** as 17, **D1** as 25, etc.

The ESX-K-100C-32W takes advantage of new Gigabit Ethernet/Fast Ethernet ASIC technology known as “Kodiak.” This module provides 4 priority levels and 256 queues per Kodiak ASIC.

◆ **Note** ◆

Kodiak-based modules support up to 4 levels of priority (0-1, 2-3, 4-5, 6-7). This is *not* compatible with the implementation of VLAN priority of Mammoth-based modules. Kodiak based priority VLANs can only be used with other Kodiak based priority VLANs.

With the optional HRE-X you can increase routing performance to 1.5 million packets per second per module and up to 12 Mpps in a fully-loaded 9-slot chassis.


ESX-K-100C-32W Technical Specifications	
Number of ports	32
Connector Type	RJ-45
Standards Supported	IEEE 802.3; IAB RFCs 826, 894
Data Rate	10 or 100 Mbps (full or half duplex)
Maximum Frame Size	1,518 bytes
MAC Addresses Supported	ESX-K-100C-32W: 1,024 ESX-K-100C-32W4: 4,096
Connections Supported	10BaseT hub or device 100BaseTx hub or device
Cable Supported	10BaseT Unshielded twisted-pair (UTP) 100BaseTx Unshielded twisted-pair: Category 5, EIA/TIA 568 Shielded twisted-pair Category 5, 100 ohm
Maximum Cable Distance	100 m
Current Draw	10.25 amps without an HRE-X 11.75 amps with an HRE-X

OK1 (Hardware Status). On Green when the module has passed diagnostic tests successfully. On Amber when the hardware has failed diagnostics or if the corresponding image file for the module is not in flash memory.

OK2 (Software Status). Blinking Green when the module software was downloaded successfully and the module is communicating with the MPX. Blinking Amber when the module is in a transitional state. On solid Amber if the module failed to download software from the MPX.

Module LEDs

Port LEDs



Each LED corresponds to a port on the module. When an LED is on Green continuously, a good cable connection exists. The LED will blink Green when traffic is transmitted or received on the port.

32-Port Advanced Auto-Sensing 10/100 Ethernet Switching Module

Fast (100 Mbps) Ethernet Modules

Alcatel's Omni Switch/Router Fast Ethernet modules can be used to connect networks with 100 Mbps workstations or as a network backbone.

The following Omni Switch/Router Fast Ethernet modules are available:

- ESX-K-100FM/FS-16W Advanced switching module with sixteen (16) Fast Ethernet (100 Mbps) backbone connections using MT-RJ ports.

This module is described and illustrated in the following sections.

ESX-K-100FM/FS-16W

The ESX-K-100FM/FS-16W Omni Switch/Router Fast Ethernet switching module has sixteen (16) fiber MT-RJ ports that each support a fully-switched 100 Mbps connection in full-duplex mode. This module provides high-speed backbone connectivity. It also supports backbone features such as 802.1q and OmniChannel. Each port uses the full 100 Mbps of bandwidth in each direction (see *ESX-K-100FM/FS-16W Technical Specifications* on page 3-20). The single mode version is referred to as the ESX-K-100FS-16W; the multimode version is referred to as the ESX-K-100FM-16W. Multimode and single mode connectors are differentiated by color: multimode connectors are black and single mode connectors are blue.

◆ Note ◆

If your network currently uses SC connectors, you can order MT-RJ-to-SC cables from Alcatel.

The MT-RJ fiber port supports full-duplex operation. You can configure half-duplex mode on each port through **10/100cfg**. An additional software command, **10/100vc**, allows you to view the current line speed and link mode of each port connection. The **10/100cfg** and **10/100vc** commands are described in Chapter 15, "Managing Ethernet Modules."

The ESX-K-100FM/FS-16W is best used as a backbone connection in networks where Fast Ethernet is used as the backbone media. Each 100Base-Fx port may also connect to a single high-traffic device, such as a mail or file server.

The ESX-K-100FM/FS-16W takes advantage of new Gigabit Ethernet/Fast Ethernet ASIC technology known as "Kodiak." This module has provides 4 priority levels and 256 queues per Kodiak ASIC.

◆ Note ◆

Kodiak-based modules support up to 4 levels of priority (0-1, 2-3, 4-5, 6-7). This is *not* compatible with the implementation of VLAN priority of Mammoth-based modules. Kodiak based priority VLANs can only be used with other Kodiak based priority VLANs.


With the optional HRE-X you can increase routing performance to 1.5 million packets per second per module and up to 12 Mpps in a fully-loaded 9-slot chassis.

ESX-K-100FM/FS-16W Technical Specifications	
Number of ports	16
Connector Type	MT-RJ
Standards Supported	IEEE 802.3; IAB RFCs 826, 894
Data Rate	100 Mbps (full duplex)
Maximum Frame Size	1,518 bytes
MAC Addresses Supported	8,192
Connections Supported	100Base-Fx connection to backbone or server
Cable Supported	Multimode: 62.5/125 micron multimode fiber Single mode: single mode fiber
Optical output power	Multimode: -19 to -14 dBm Single-mode: -20 to -14 dBm
Optical receiver sensitivity	Multimode: -31 dBm Max. Single-mode: -31 dBm Max.
Cable Distance	Multimode: approximately 2 km Single-mode: approximately 15 km
Current Draw	9.75 amps without an HRE-X 11.25 amps with an HRE-X

Warning Label. This label indicates that the module contains an optical transceiver).

Each LED corresponds to a port on the module. When an LED is on Green continuously, a good cable connection exists. The LED will blink Green when traffic is transmitted or received on the port.

MT-RJ connectors will be color coded to indicate multimode (Black) or single mode (Blue).



Module Label. This label will indicate the ESX-100FM/FS-16W type. It will read either **ESX-K 100 mm** (multimode cable) or **ESX-K 100 sm** (single mode cable).

Module LEDs
OK1 (Hardware Status). On Green when the module has passed diagnostic tests successfully. On Red when the hardware has failed diagnostics.

OK2 (Software Status). Blinking Green when the module software was downloaded successfully and the module is communicating with the MPX. Blinking Red when the module is in a transitional state. On solid Red if the module failed to download software from the MPX.

16-Port Advanced Fast Ethernet Switching Module

WAN Modules

The Omni Switch/Router currently supports the following Wide Area Network (WAN) modules:

- WSX-S-2W Provides two serial ports that support Frame Relay or PPP.
- WSX-SC Provides four or eight serial ports that support Frame Relay or PPP with data compression.
- WSX-FT1/E1-SC Provides one or two T1/E1 ports and one or two serial ports that support Frame Relay or PPP with data compression.
- WSX-BRI-SC Provides one or two Universal Serial Ports (USPs) ports and one or two ISDN-BRI ports that support Frame Relay or PPP with data compression.


All of these modules are described and illustrated in the sections beginning on page 3-27.

A WSX switching module is actually a submodule, or daughtercard, that attaches to an Omni Switch/Router High-Speed Module (HSX). The HSX contains RISC processors, RAM for holding software image files, ASICs for performing switching, and Content Addressable Memory (CAM) for storing MAC addresses. You plug your cable into the WSX submodule, but it is the HSX module that connects to the switch's backplane.


WAN Pinouts

The figures and tables on the following pages illustrate the pinouts used on Omni Switch/Router WAN modules. Please note that the signal commonly known as "remote loop-back" (LL) is not supported on the WAN serial port (see *WAN Serial Port Specifications* on page 3-25). In addition, CTP2, CTP1, and CTP0 are assigned to CS(B), DR(B), and CD(B), respectively, on the serial port. The latter are not used in the cable configurations that require the former.


See Appendix B, "Custom Cables," for information on cables used to connect the serial connector to different interface types.




WAN BRI Port Specifications (S/T Interface)	
Pin Number	Standard Signal Name
1	Not Used
2	Not Used
3	Rcv + from TE
4,	Rcv - from TE
5	Xmt + from TE
6	Xmt - from TE
7	Not Used
8	Not Used



WAN BRI Port Specifications (U Interface)	
Pin Number	Standard Signal Name
1	Not Used
2	Not Used
3	Xmt to /Rcv from Network
4,	Xmt to /Rcv from Network
5	Not Used
6	Not Used
7	Not Used
8	Not Used



WAN T1/E1 Port Specifications	
Pin Number	Standard Signal Name
1	Rx_Ring
2	Rx_Tip
3	Chassis GND
4,	Tx_Ring
5	Tx_Tip
6	Chassis GND
7	Chassis GND (A jumper is provided for connecting Pins 7 and 8 to the chassis ground, if required.)
8	Chassis GND (A jumper is provided for connecting Pins 7 and 8 to the chassis ground, if required.)



WAN Serial Port Numbering

WAN Serial Port Specifications							
Generic Signal Name	Source	Alcatel SPI		EIA-530		RS-449	
		Mnemonic	Pin	Mnemonic	Pin	Mnemonic	Pin
Shield	--	Shield	1	--	1	--	1
Signal Ground	--	AB	7	AB	7	SG	19
Transmitted Data	DTE	TD(A)	2	BA(A)	2	SD(A)	4
		TD(B)	14	BA(B)	14	SD(B)	22
Received Data	DCE	RD(A)	3	BB(A)	3	RD(A)	6
		RD(B)	16	BB(B)	16	RD(B)	24
Transmit Clock	DCE	TC(A)	15	DB(A)	15	ST(A)	5
		TC(B)	12	DB(B)	12	ST(B)	23
Receive Clock	DCE	TC(A)	17	DD(A)	17	RT(A)	8
		TC(B)	9	DD(B)	9	RT(B)	26
Ext. Transmit Clock	DTE	XC(A)	24	DA(A)	24	TT(A)	17
		XC(B)	11	DA(B)	11	TT(B)	35
Request To Send	DTE	RS(A)	4	CA(A)	4	RS(A)	7
		RS(B)	19	CA(B)	19	RS(B)	25
Clear To Send	DCE	CS(A)	5	CB(A)	5	CS(A)	9
		CS(B)	13	CB(B)	13	CS(B)	27
Data Set Ready	DCE	DR(A)	6	CC(A)	6	DM(A)	11
		DR(B)	22	CC(B)	22	DM(B)	29
Data Terminal Ready	DTE	TR(A)	20	CD(A)	20	TR(A)	12
		TR(B)	23	CD(B)	23	TR(B)	30
Data Carrier Detect	DCE	CD(A)	8	CF(A)	8	RR(A)	13
		CD(B)	10	CF(B)	10	RR(B)	31
Local Loopback	DTE	LL	18	LL	18	LL	10
Remote Loopback	DTE	RL	21	RL	21	RL	14
Ring Indicator	DCE	RI/TM	25	--	--	--	--
Test Mode	DCE	RI/TM	25	TM	25	TM	18
Cable Type 4	--	CTP4	18		n/c		n/c
Cable Type 3	--	CTP3	26		n/c		n/c
Cable Type 2	--	CTP2	13				
Cable Type 1	--	CTP1	22				
Cable Type 0	--	CTP0	10				

continued on next page...

WAN Serial Port Specifications (cont.)							
Generic Signal Name	Source	X.21/X.26		V.35		RS232	
		Mnemonic	Pin	Mnemonic	Pin	Mnemonic	Pin
Shield	--	--	1	--	A	--	1
Signal Ground	--	G	8	102	B	AB	7
Transmitted Data	DTE	T(A)	2	103(A)	P	BA	2
		T(B)	9	103(B)	S		
Received Data	DCE	R(A)	4	104(A)	R	BB	3
		R(B)	11	104(B)	T		
Transmit Clock	DCE	--	--	114(A)	Y	DB	15
				114(B)	AA		
Receive Clock	DCE	S(A)	6	115(A)	V	DD	17
		S(B)	13	115(B)	X		
Ext. Transmit Clock	DTE	B(A)	7	113(A)	U	DA	24
		B(B)	14	113	W		
Request To Send	DTE	C(A)	3	105	C	CA	4
		C(B)	10				
Clear To Send	DCE	--	--	106	D	CB	5
Data Set Ready	DCE	--	--	107	E	CC	6
Data Terminal Ready	DTE	--	--	108	H	CD	20
Data Carrier Detect	DCE	I(A)	5	109	F	CF	8
		I(B)	12				
Local Loopback	DTE	--	--	141	L	LL	18
Remote Loopback	DTE	--	--	140	N	RL	21
Ring Indicator	DCE	--	--	125	J	CE	22
Test Mode	DCE	--	--	142	NN	TM	25
Cable Type 4	--		n/c		n/c		
Cable Type 3	--		n/c		n/c		
Cable Type 2	--						
Cable Type 1	--						
Cable Type 0	--						

WSX-S-2W

The WSX-S-2W supports two (2) serial ports, which can provide access rates from 9.6 Kbps to 2 Mbps. The WSX-S-2W also supports three types of clocking (internal, external, and split). See *WSX-S-2W Technical Specifications* on page 3-27 for more information.

◆ **Note** ◆


The WSX-S-2W does not support hardware compression.

The WSX-S-2W can sense and auto-configure for any of five serial cable types (RS-232, V.35, X.21, RS-530, and RS-449). A WSX-S-2W port is normally considered a physical DTE device. It can be turned into a physical DCE device—for speed or clocking purposes—by plugging in a DCE cable. The WSX-S-2W senses whether a DCE or DTE cable is connected.

Software in the switch allows you to configure parameters for the Frame Relay or Point-to-Point Protocol (PPP). Software commands allow you to view the status of the WAN connection at the WSX-S-2W board, port, or virtual circuit level. Extensive statistics are provided at each level. Software commands for Frame Relay are described in Chapter 29, “Managing Frame Relay”; commands for PPP are described in Chapter 30, “Point to Point Protocol.”

With the optional HRE-X you can increase routing performance to 1.5 million packets per second per module and up to 12 Mpps in a fully-loaded 9-slot chassis.

WSX-S-2W Technical Specifications	
Number of ports	2
Connector Type	High-density 26-pin shielded serial
Protocols Supported	Frame Relay and Point-to-Point (PPP)
Data Rates Supported	9.6, 19.2, 56, 64, 128, 256, 512, 768, 1024, 1536, 2048 Kbps
Clocking	Internal, External, or Split
Virtual Circuits Supported	Permanent Virtual Circuits (PVCs)
MAC Addresses Supported	4,096
Connections Supported	Physical Data Terminal Equipment (DTE) or Data Communication Equipment (DCE)
Cable Supported	DTE or DCE in the following types: R2-232, V.35, X.21, RS-530, RS-449
Power Consumption	5.25 amps (without an HRE-X) 6.75 amps (with an HRE-X)



2-Port WAN Frame Relay Switching Module

WSX-SC

The WSX-SC supports 4 or 8 serial ports, each of which can provide access rates from 9.6 Kbps to 2 Mbps. The 4-port version is referred to as the WSX-SC-4W, and the 8-port version is referred to as the WSX-SC-8W. The WSX-SC supports STAC hardware compression and three types of clocking (internal, external, and split). See *WSX-SC Technical Specifications* on page 3-30 for more information.

The WSX-SC can sense and auto-configure for any of five serial cable types (RS-232, V.35, X.21, RS-530, and RS-449). A WSX-SC port is normally considered a physical DTE device. It can be turned into a physical DCE device—for speed or clocking purposes— by plugging in a DCE cable. The WSX-SC board senses whether a DCE or DTE cable is connected.

Software in the switch allows you to configure parameters for the Frame Relay or Point-to-Point Protocol (PPP). Software commands allow you to view the status of the WAN connection at the WSX-SC board, port, or virtual circuit level. Extensive statistics are provided at each level. Software commands for Frame Relay are described in Chapter 29, “Managing Frame Relay”; commands for PPP are described in Chapter 30, “Point to Point Protocol.”

With the optional HRE-X you can increase routing performance to 1.5 million packets per second per module and up to 12 Mpps in a fully-loaded 9-slot chassis.

WSX-SC Technical Specifications	
Number of ports	4 or 8
Connector Type	High-density 26-pin shielded serial
Protocols Supported	Frame Relay and Point-to-Point (PPP)
Data Rates Supported	9.6, 19.2, 56, 64, 128, 256, 512, 768, 1024, 1536, 2048 Kbps
Compression	Hardware-based using STAC 9705
Clocking	Internal, External, or Split
Virtual Circuits Supported	Permanent Virtual Circuits (PVCs)
MAC Addresses Supported	4,096
Connections Supported	Physical Data Terminal Equipment (DTE) or Data Communication Equipment (DCE)
Cable Supported	DTE or DCE in the following types: R2-232, V.35, X.21, RS-530, RS-449
Power Consumption	WSX-SC-4W without an HRE-X: 6.25 amps WSX-SC-4W with an HRE-X: 7.75 amps WSX-SC-8W without an HRE-X: 8.25 amps WSX-SC-8W with an HRE-X: 9.75 amps

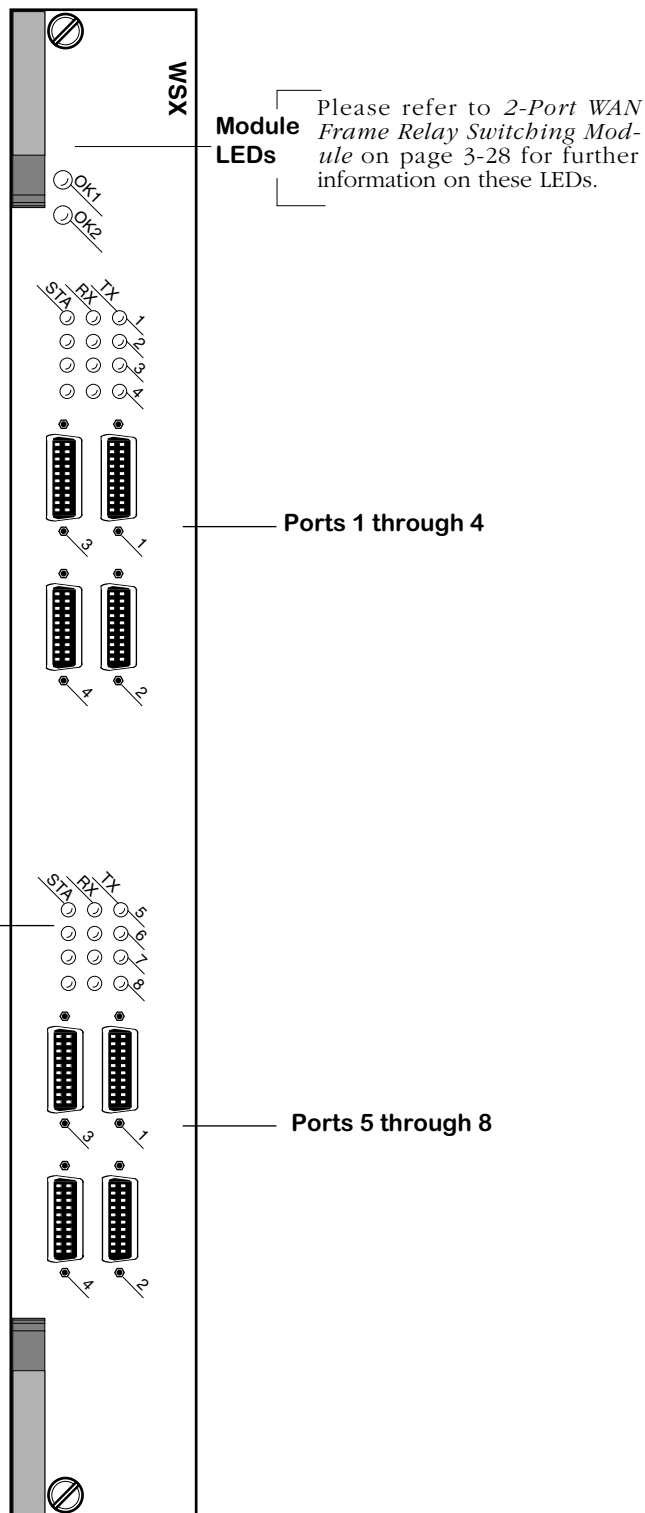
The module includes one row of LEDs for each port. The LEDs for a given port are located in the row labeled with the port number. If the WSX module includes a total of eight ports, then the module contains two sets of four rows of LEDs. The second set of LEDs are located above the second set of ports.

STA (Status). On Green continuously when the port connection is operational. Off when the port is disabled or the cable is detached. Blinking On/Off if cable is attached but receive control data is detected as down.

This LED also blinks during initialization, diagnostics, or when invalid data is being exchanged on the port.

TX (Transmit). On “half-bright” Green when idle and Green with occasional flickers when the port is transmitting data.

RX (Receive). On “half-bright” Green when idle and Green with occasional flickers when the corresponding port is receiving data.



8-Port WAN Frame Relay Switching Module

WSX-FT1/E1-SC

The WSX-FT1/E1-SC module contains one or two T1 or E1 ports and one or two serial ports. T1 and E1 ports use RJ-48C connectors. The T1 version of this module is referred to as the WSX-FT1-SC; the E1 version is referred to as the WSX-FE1-SC. You can configure these ports to run either Frame Relay or the Point-to-Point Protocol (PPP). See *WSX-FT1/E1-SC Technical Specifications* on page 3-33 for more information.

This module includes an integrated CSU/DSU to enable direct connection to a T1/E1 device, such as a PBX, or a T1/E1 line to a service provider.

You can configure physical port parameters through software commands. Configuration options include frame format, facility datalink, and line coding. In addition, the switch can store up to 24 hours of local and remote statistics. See Chapter 33, “Managing T1 and E1 Ports,” for more information on software-configurable parameters.

The WSX-FT1/E1-SC also supports STAC hardware compression.

With the optional HRE-X you can increase routing performance to 1.5 million packets per second per module and up to 12 Mpps in a fully-loaded 9-slot chassis.

WSX-FT1/E1-SC Technical Specifications	
Number of ports	1 or 2 T1 or E1 ports 1 or 2 Universal Serial ports
Connector Types	T1/E1: RJ-48C Serial: High-density, 26-pin shielded
Standards Supported	RFCs 1406, 1213, 1659
Frame Formats	T1: Superframe, Extended Superframe, Unframed E1: E1, E1-CRC, E1-MF, E1-CRC-MF, Unframed
Line Coding	T1: B8ZS or AMI E1: HDB3 or AMI
Data Rates Supported	T1: 1.544 Mbps E1: 2.048 Mbps Serial: 56, 64, 128, 256, 384, 512, 768, 1024, 1536, 1544, 2048 Kbps
Compression	Hardware-based using STAC 9705
Facility Datalink Protocol	ANSI T1.403 and AT&T 54016
MAC Addresses Supported	4,096
Connections Supported	Physical Data Terminal Equipment (DTE) or Data Communication Equipment (DCE)
Cable Supported	Serial Ports DTE or DCE of the following types: R2-232, V.35, X.21, RS-530, RS-449
Cable Distance	T1/E1 (short haul): 200 meters T1/E1 (long haul): 1829 meters
Power Consumption	WSX-FT1/E1-SC-1W without an HRE-X: 5.75 amps WSX-FT1/E1-SC-1W with an HRE-X: 7.25 amps WSX-FT1/E1-SC-2W without an HRE-X: 7.25 amps WSX-FT1/E1-SC-2W with an HRE-X: 8.75 amps

This module includes one set of LEDs for each port. The LEDs for a given port are located above the port. If the WSX module includes four ports, then the module contains two sets of LEDs. The second set of LEDs are located above the third and fourth ports.

STA (Status). On Green continuously when the port connection is operational. Off when the port is disabled or the cable is detached. Blinking On/Off if cable is attached but receive control data is detected as down.

This LED also blinks during initialization, diagnostics, or when invalid data is being exchanged on the port.

TX (Transmit). On “half-bright” Green when idle and Green with occasional flickers when the port is transmitting data.

RX (Receive). On “half-bright” Green when idle and Green with occasional flickers when the corresponding port is receiving data.

Serial Port LEDs

Module LEDs

Please refer to *2-Port WAN Frame Relay Switching Module* on page 3-28 for further information on these LEDs.

Port 1: T1 or E1

Port 2: Serial

T1/E1 Port LEDs

ALM (Alarm). On Green when the port is enabled and a signal is present. On Yellow when an error has occurred on the port.

ACT (Activity). On Green when the T1 or E1 port is transmitting or receiving data.

STA (Status). On Green continuously when the port connection is operational. Off when the port is disabled or the cable is detached.

Port 3: T1 or E1

Port 4: Serial




WAN 2-Port Serial and 2-Port Fractional T1/E1 Switching Module

WSX-FE1-SC Cabling/Jumper Settings

The WSX-FE1-SC supports both twisted pair (120 Ohm) and coaxial (75 Ohm) cable types. The default is 120 Ohm. You must set a pair of jumpers (JP2 and JP4) on the back of the board to correspond to the type of cable you are using. For more detailed information on the types of cables to use with this module, see Appendix B, "Custom Cables." The illustration below shows the correct jumper positions.

◆ **Note** ◆

JP3 is reserved. Do not set a jumper across JP3.



Cable Termination Jumpers for WSX-FE1-SC

WSX-BRI-SC

The ISDN Basic Rate Interface WAN Switching Module (WSX-BRI-SC) supports either one (1) serial port and one (1) BRI port or two (2) serial ports and two (2) BRI ports. The version with 1 serial port and 1 BRI port is referred to as the WSX-BRI-SC-1W; the version with 2 serial ports and 2 BRI ports is referred to as the WSX-BRI-SC-2W. See *WSX-BRI-SC Technical Specifications* on page 3-37 for more information.

The serial port on a WSX-BRI-SC module is essentially the same as the serial ports found on the WSX-SC module. A WSX-BRI-SC serial port can detect, and configure itself, for any of five serial cable types (RS-232, V.35, X.21, RS-530, and RS-449). A WSX-BRI-SC serial port is normally considered a physical DTE device, but it can be turned into a physical DCE device—for speed or clocking purposes—by simply plugging in a DCE cable. The WSX-BRI-SC internally senses whether a DCE or DTE cable is connected and configures itself appropriately.

The BRI port on the WSX-BRI-SC board can be configured as either a “U” or an “S/T” type of interface (the board is shipped set to “U”). Either type of interface supports two “B” channels operating at 56/64 Kbps and one “D” channel operating at 16 Kbps.

Software running in the switch allows you to configure the operation of the Point-to-Point Protocol (PPP) over the serial port or the BRI port. The serial port can also support the Frame Relay protocol. The software commands used to configure PPP are described in Chapter 30, “Point-to-Point Protocol.” The software commands used to configure Frame Relay are described in Chapter 29, “Managing Frame Relay.” The software commands used to configure the WAN “links” that support PPP connections are described in Chapter 31, “WAN Links.” Finally, the software commands used to manage the ISDN ports are described in Chapter 32, “Managing ISDN Ports.”

With the optional HRE-X you can increase routing performance to 1.5 million packets per second per module and up to 12 Mpps in a fully-loaded 9-slot chassis.

WSX-BRI-SC Technical Specifications	
Number of ports	1 or 2 pairs of a serial port and an ISDN Basic Rate Interface (BRI) port
Serial Connector Type	High-density 26-pin shielded serial
BRI Connector Type	RJ-45
Protocols Supported	Point-to-Point Protocol (PPP); Frame Relay (supported on the serial port only)
Data Rates Supported	2 "B" Channels at 56/64 Kbps 1 "D" Channel at 16 Kbps
Compression	Hardware-based using STAC 9705
MAC Addresses Supported	4,096
Serial Port Connections Supported	Physical Data Terminal Equipment (DTE) or Data Communication Equipment (DCE)
Serial Cables Supported	DTE or DCE in the following types: R2-232, V.35, X.21, RS-530, RS-449
BRI Port Connections Supported	"U" interface or "S/T" interface (jumper-selectable; "U" is shipping default)
Maximum Cable Distance	BRI: 100 m
Switch Types Supported	National ISDN-1, AT&T 5ESS, Northern Telecom DMS100, ETSI Euro-ISDN Net3
ISDN Standards Supported	Q.921, Q.931, I.430, T1.601
Power Consumption	WSX-BRI-SC-1W without an HRE-X: 4.75 amps WSX-BRI-SC-1W with an HRE-X: 6.25 amps WSX-BRI-SC-2W without an HRE-X: 5.25 amps WSX-BRI-SC-2W with an HRE-X: 6.75 amps

The WSX-BRI module includes one set of LEDs for each port. The LEDs for a given port are located in the set labeled with the port number. If the HSX module contains two WSX-BRI daughter cards, the second set of ports (one Serial and one BRI) are numbered as Ports 3 and 4 respectively, and include their own separate set of LEDs that function exactly like those related to Ports 1 and 2.

STA (Status). On Green continuously when the port connection is operational. Off when the port is disabled or the cable is detached. Blinking On/Off if cable is attached but receive control data is detected as down.

This LED also blinks during initialization, diagnostics, or when invalid data is being exchanged on the port.


TX (Transmit). On “half-bright” Green when idle and Green with occasional flickers when the port is transmitting data.

RX (Receive). On “half-bright” Green when idle and Green with occasional flickers when the corresponding port is receiving data.

ACT (Activity). On Green when the ISDN-BRI port is sending or receiving data.

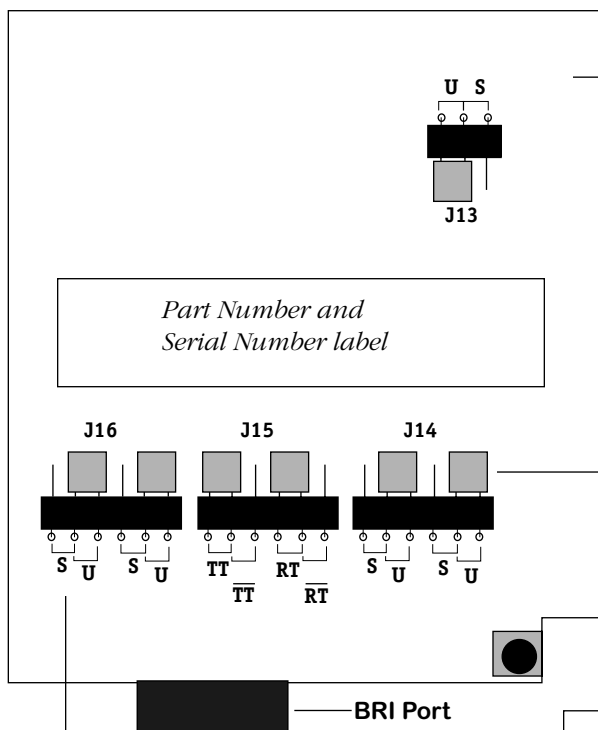
UIF (“U” Interface). On Green when the ISDN-BRI port is configured as a “U” type of interface. Off when the port is configured as an “S/T” type of interface.

STA (Port 2/4 Status). On Green continuously when the port connection is operational. Off when the BRI port is disabled or the cable is detached. This LED blinks during initialization.



WAN 2-Port Serial and 2-Port BRI-ISDN Switching Module

**Jumper Configuration for the "U" Interface
(this is how the board is shipped)**




This is a simplified view of the bottom lower-right quadrant of the WSX-BRI submodule. Immediately above the BRI port are three jumper blocks labelled J14, J15, and J16. About two inches above and to the right is another jumper labeled J13. J13, J14, and J16 are used to switch between the "U" and "S/T" interfaces. J15 is used to set transmit and receive termination for the "S/T" interface.

The gray boxes are the jumper blocks

The small labels next to the jumper pins at J13, J14, and J16 indicate which pins must be bridged to set the BRI port to either the "U" or the "S/T" interface.

Small labels under the pins at J15 indicate which pins must be bridged to set Transmit Termination (tt) and Receive Termination (rt) to the "on" or "off" position (the two sets of letters with a line over them indicate the "off" settings).



**Jumper Configuration for the "S/T" Interface
(transmit/receive termination are set to "on")**

4 The User Interface

In order to configure parameters and statistics on the switch, you may connect it to a terminal, such as a PC or UNIX workstation, using terminal emulation software. The command interfaces used on the switch are part of the MPX executable image. When a switch boots up, the boot monitor handles the loading of this executable image and system startup. Once the image is loaded and initialized, the CLI starts.

You access the command interfaces through a connection with the switch. This connection can be made directly to the serial port, through a modem, or over a network via Telnet. You can have up to four simultaneous connections to an Omni Switch/Router. (Please see Multiple User Sessions on page 4-33 for further details.) For Telnet access, you must first set up an IP address for the switch. See the *Getting Started Guide* that came with your switch for information on setting up an IP address and logging in.

Overview of Command Interfaces

The Alcatel Omni Switch/Router has two different command interfaces available for configuring parameters and viewing statistics. They are the User Interface (UI) and the Command Line Interface (CLI). Prior to software Release 4.4, the switch automatically booted up in the UI mode. In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in the CLI mode.

◆ Terminology Notes◆

Command interface generically refers to any mechanism resident in the software that allows a user to change switch configurations or to display statistics.

The *UI* is the original command interface used exclusively on all Alcatel Omni Switch/Router and OmniAccess products. The UI has its commands grouped into functional menus. Prior to software Release 4.1, the UI was the only command interface supported on the Omni Switch/Router products.

The *CLI* is Alcatel's text-based configuration interface that allows you to configure Omni Switch/Router and OmniAccess products using single-line text commands. The CLI was implemented in software Release 4.1 and higher. In release 4.4 and later it is the default interface.

Changing Between the CLI and UI Modes

Once you log on to the switch, the following screen displays. You must press the **<Enter>** key to start the command interface.

```
*****  
  
Alcatel Omni Switch/Router  
Copyright (c), 1994-2002 Alcatel Internetworking, Inc. All rights reserved.  
Omni Switch/Router is a trademark of Alcatel Internetworking, Incorporated,  
registered in the United States Patent and Trademark Office.  
Press ENTER to start  
->
```

After you press **<Enter>**, the CLI starts automatically and the following text displays.

```
Entering command line interface.  
->
```

At this point, you are in the CLI mode and may configure the switch or display statistics using the commands described in the *Text-Based Configuration CLI Reference Guide*. If you want to use the UI command interface, type **ui** and press **<Enter>**. This causes the switch to leave the CLI mode and enter the UI mode, provided you are using a login with Read/Write privileges. You can verify that you are in the UI mode by typing **?** to display the top-level menu for the UI as shown below.

<i>l%?</i> Command	Main Menu
File	Manage system files
Summary	Display summary info for VLANs, bridge, interfaces, etc.
VLAN	VLAN management
Networking	Configure/view network parameters such as routing, etc.
Interface	View or configure the physical interface parameters
Security	Configure system security parameters
System	View/set system-specific parameters1
Services	View/set service parameters
Switch	Enter Any to Any Switching menu
Help	Help on specific commands
Diag	Display diagnostic level commands
Quit/Logout	Log out of this session
?	Display the current menu contents

To change from the UI mode back to the CLI mode, type **cli** and press **<Enter>**.

◆ Note ◆

Note the default command prompt for the UI is **l%**. The default command prompt for the CLI is **->**. You can change the UI system prompt by using the **uic** command.

Exit the Command Interface

To exit your current session with the switch from the CLI or the UI mode, type either **quit** or **logout** at the prompt, then press **<Enter>**. Your session is immediately terminated.

◆ Note ◆

If you forget which command interface mode you are in, type the **?** character. If you are in the UI mode, the Main Menu will display as shown above. If you are in the CLI mode, the switch will show the following display.

```
^NO, SHOW, VOICE, SYSTEM, ACCOUNTING, . . .  
->
```

UI to CLI Command Cross Reference

The chapters in this Users Guide are organized around the UI commands as they are grouped into menus and sub-menus. Even though the Omni Switch/Router software has been changed to boot up in the CLI mode, the Users Guide conforms to its original design. The CLI commands are fully documented in the *Text-Based Configuration CLI Reference Guide*.

This section presents the key UI commands that are explained in this User's Manual along with their CLI equivalents. Where the CLI commands support partition management, these tables also list the partition management family to which the commands belong.

Hardware Commands

The hardware section of this manual set consists of Chapters 1 through 3. There are relatively few UI commands in this section because these chapters cover the hardware elements of the switch. The commands defined in these chapters are listed in the Hardware Table beginning on page 4-4.

Hardware Table

Chapter	UI Command	Equivalent CLI Commands	PM Family
1, "OSR Chassis/Power Supplies"	No UI commands are defined in this chapter.	N/A	N/A
2, "MPX"	ethernetc	ethernet management port view ethernet manage port	GF-interface
3, "OSR Switching Modules"	10/100cfg 10/100vc	ethernet view interface fastethernet	GF-interface

Basic Switch Management Commands

The table beginning on page 4-5 summarizes the features supported in the UI and the CLI for Chapters 4 through 11.

Basic Switch Management Table

Chapter	UI Command	Equivalent CLI Commands	PM Family
4, "The User Interface"	alert, echo, history, kill, ping, pwd, timeout, who lookup, save, summary, uic, write	alert, echo, history, kill, ping, password, timeout, who Unsupported	No PM Support
5, "Installing Switch Software"	ftp load primary, secondary	ftp load primary, secondary	GF-Ftp GF-File
6, "Configuring Management Processor Modules"	configsync ethernetc imgsync mpm mpmget mpmload mpmreplace mpmrm mpmstore renounce secreset slipc sls swap syncctl takeover	configuration copy ethernet management port image copy view mpm command load primary mpm file load secondary mpm file replace secondary mpm file remove secondary mpm file store secondary mpm file takeover reload secondary mpm slip view secondary mpm file swap configuration auto-copy takeover	GF-File
7, "Managing Files"	cd cp load newfs ftp ls pwd rm imgcl	cd copy load newfs ftp ls password rm imgcl	GF-CD GF-System GF-System GF-System GF-FTP GF-LS 18-User GF-RM GF-System
8, "Switch Security"	pw reboot useradd userdel usermod userview asacfg secdefine secapply layer2auth, privs, secapply, secdefine, seclog, security	password reboot now user no user user view user ldap server secure access filter secure access no filter view secure access filter security security custom security no custom Unsupported	18-User GF-Reboot 18-User 18-User 18-User 18-User 1-Configuration GF-System GF-System GF-System GF-System GF-System No PM Support

continued on next page...

Basic Switch Management Table (continued)

Chapter	UI Command	Equivalent CLI Commands	PM Family
9, "Switch-Wide Parameters"	cacheconfig camstat dt hrexassign hrexdisplay hrexhashopt hrexutil info memstat modvp newfs saveconfig slot syscfg systat camcfg, fsck, sc, si, ss, taskstat	configuration cache camstat dt hrexassign hrexdisplay hrexhashopt hrexutil info memstat modvp newfs configuration cache save slot syscfg systat Unsupported	No PM Support
10, "Switch Logging"	secdefine secapply caplog, cmdlog, syslog, conlog, debuglog, swlogc	secure access filter secure access no filter view secure access filter security security custom security no custom Unsupported	GF-System
11, "Health Statistics"	hdcfg health hmstat hpstat hreset	health threshold view health statistics view health statistics view health statistics health statistics reset	GF-System

Network Management Commands

The table on page 4-6 summarizes the commands supported in the UI and the CLI for Chapters 12 through 14.

Network Management Table

Chapter	UI Command	Equivalent CLI Commands	PM Family
12, "Network Time Protocol"	ntconfig, ntstats, ntadmin, ntaccess	Unsupported	No PM Support
13, "Configuring SNMP"	snmpc snmps	view snmp set snmp	6-SNMP
14, "RMON and DNS Resolver"	res probes events names chngmac	res view rmon probes view rmon events view dns Unsupported	GF-System

Layer II Switching Commands

The table on page 4-7 summarizes the features supported in the UI and the CLI for Chapters 15 through 18.

Layer II Switching Table

Chapter	UI Commands	Equivalent CLI Commands	PM Family
15, “Managing Ethernet Modules”	addprtchl chnlinfo crechnl delchnl delprtchl eth10/100vc eth10/100cfg	static agg view statis linkagg number static linkagg number type no static linkgg number static agg no view interface fastethernet interface ethernet	GF-Interface
16, “Managing 802.1Q Groups”	cas, das, mas, vas	All commands used to create, delete, modify and view a service, plus the message command are supported.	GF-System
17, “Configuring Bridging Parameters”	fddi, fsmt, fsid, fsmtc, fsstatus, fmac, fmaddr, fmstats, fmctrs, fport, fportstatus, fportctrs, fportc, macstat, slipc maccirstat, selgp, srsf, srtbcfg, srtbclrrif, srtbrif	Supported Unsupported	5-Bridge
18, “Configuring Frame Translations”	actfstps, bps, dbrmap, fc, flc, fls, fs, fstps, fwt, macinfo, modvp, rts, srtbrif, stc, sts, stpc, stps, swchmac autoencaps, ethdef, facdef, propipx, swchmac, trdef	Supported Unsupported	5-Bridge

Groups, VLANs, Policies Commands

The table beginning on page 4-8 summarizes the features supported in the UI and the CLI for Chapters 19 through 24.

Groups, VLANs, Policies Table

Chapter	UI Command	Equivalent CLI Commands	PM Family
19, “Managing Groups and Ports”	swch vi autoencaps, ethdef, facdef, propipx, swchmac, trdef	port encapsulation view group rules Unsupported	2-Group
20, “Group and VLAN Policies”	addqgp addvp cas cats crgp dats delqgp gmcfg gmstat gp modvl pmapcr pmapdel pmapmod pmapv pmcfg pmon pmstat pmp prty_mod prty_disp rmgp rmvp vats ve vi viqgp vs via vpl at, br, pmd, prty_mod, vlan, vigl, viqgp	group num 802.1q group num interface fddi svc, group 802.1q atm service group elan group group num no 802.1q group no elan group mobility group mobility view group group router, vlan router port mapping ingress no port mapping port mapping view port mapping port monitor configuration port monitor view port monitor resume port monitor group priority num view group priority no group group no interface view group auto view group virtual errors view group rules view ethernet view group virtual statistics view group virtual (ports) view group mobility Unsupported	2-Group
21, “InterSwitch Protocols”	atvl fwtl modatvl vap viatrl vivl vlap	view vlan rules view group mac group mac, vlanmac, vlan user, vlan port, vlan chcp port, vlan dhcp mac, vlan protocol, vlan binding ip, vlan binding vap port vlan ip, vlan ipx view vlan rules view vlan rules vlap	6-Group GF-System GF-System

continued on next page...

Group, VLANs, Policies Table (continued)

Chapter	UI Commands	Equivalent CLI Commands	PM Family
22, "Managing AutoTracker VLANs"	gmap, gmapst gmappaptime gmapholdtime gmapuptime xmapst xmapls xmapcmntime xmapdisctime	gmap gmap gap time gmap hold time gmap up time xmap, view xmap status view xmap, view xmap xmap common time xmap discovery time	6-Group
23, "Multicast VLANs"	cats cratvl crmcvl defvl fwtvl gmcfg gmstat mag mcvl modatvl rmatvl vag vats viatrl vimcvl vivil vpl atvl, vigl, xip	group elan vlan, vlan router ip, vlan router ipx, vlan mac, vlan user, vlan dhcp port vlan dhcp nac, vlan protocol, vlan binding ip, vlan binding mac, vlan binding port vlan ip, vlan ipx multicast vlan, multicast vlan port multicast vlan mac, vlan protocol vlan binding ip, vlan binding mac vlan binding port, multicast vlan descr vlan default view group mac view group authenticated group mobility group authentication, group authentication protocol view multicast vlan group mac, group mac range, group user, group port, group dhcp port, group dhcp mac, group dhcp range group protocol, group binding ip, group protocol mac, group binding port, group ip, group ipx, vlan mac, vlan user, vlan port, vlan dhcp port, vlan dhcp mac, vlan protocol, vlan binding protocol, vlan binding mac, vlan binding port, vlan ip, vlan ipx no vlan view group authenticated view group auto view vlan rules view multicast vlan ports view group ports, view group vports view group mobility	6-Group
24, "AutoTracker VLAN Examples"	crmcvl, modmcvl rmmcvl vimcrl vimcvl	multicast vlan no multicast vlan view multicast vlan rules view multicast vlan	GF-System

Routing Commands

The table beginning on page 4-10 summarizes the features supported in the UI and the CLI for Chapters 25 through 27.

Routing Table

Chapter	UI Command	Equivalent CLI Commands	PM Family
25, "IP Routing"	All IP Routing commands are supported in the CLI.	All IP Routing commands are supported in the CLI.	3-IP Routing GF-System
26, "UDP Forwarding"	aisr events icmps ipfilter ipmac ipr ips names ping probes ripflush rips risr snmpc snmps telnet tcpc tcps traceroute udpl udps xlat chnghmac, flush, flconfig, ipclass, ipdirbrcast, names, probes	iproute view rmon events view icmp rip filter view mac view ip route view ip traffic ip [no] domain-lookup ping view rmon probes ripflush rips no ip route snmp config, snmp communi- ty, snmp trap, broadcast, snmp trap unicast snmp station view snmp telnet ip-address view tcp users view tcp trace view udp users view ucp arp, clear arp-cache, view arp Unsupported	3-IP Routing
27, "IPX Routing"	relayc relays avlbootmode, edit	ip helper view ip helper stats Unsupported	No PM Support

Troubleshooting Diagnostics Commands

The table beginning on page 4-13 summarizes the features supported in the UI and the CLI for Chapters 35 and 36 and Appendices A and B.

Troubleshooting/Diagnostics Table

Chapter/ Appendices	UI Command	Equivalent CLI Commands	PM Family
35, "Troubleshoot- ing"	uic	Unsupported	No PM Support
36, "Running Hardware Diagnostics"	diag	Unsupported	No PM Support
A, "Boot Line Prompt"	ethernetc	ethernet manager port	No PM Support
B, "Custom Cables"	No UI commands in this Appendix.	No CLI commands in this Appendix	No PM Support

User Interface Menu

This menu provides a top-level view of all UI menus. The commands are grouped together in the form of sub-menus. Within each sub-menu there is a set of commands and/or another sub-menu.

Command	Main Menu
File	Manage system files
Summary	Display summary info for VLANs, bridge, interfaces, etc.
VLAN	VLAN management
Networking	Configure/view network parameters such as routing, etc.
Interface	View or configure the physical interface parameters
Security	Configure system security parameters
System	View/set system-specific parameters ¹
Services	View/set service parameters
Switch	Enter Any to Any Switching menu
Help	Help on specific commands
Diag	Display diagnostic level commands
Quit/Logout	Log out of this session
?	Display the current menu contents

◆ **Note** ◆

Although the commands are grouped in a sub-menu structure, any command may be entered from any sub-menu. You are not restricted to the commands listed in the current sub-menu.

Main Menu Summary

These menus, their sub-menus, and sub-options are described in this manual. The following provides a brief overview of each item on this main menu.

File. Contains options for downloading system software, listing software files, copying files, editing files, and deleting files. This menu is fully described in Chapter 7, “Managing Files.”

Summary. Provides very basic information on the physical switch, such as its name, MAC address, and resets. It also provides options for viewing the virtual interface and information on the MIB. This menu is described in Chapter 9, “Switch-Wide Parameters.”

VLAN. The main menu for configuring Groups, virtual ports, and AutoTracker VLANs. This menu also contains a sub-menu for configuring bridging parameters, such as Spanning Tree. Groups and ports are described in Chapter 19, “Managing Groups and Ports.” VLANs are described in Chapter 22, “Managing AutoTracker VLANs” and Chapter 23, “Multicast VLANs.” Bridging parameters are described in Chapter 17, “Configuring Bridging Parameters.”

Networking. Contains menu options for managing internetworking protocols, such as SNMP and RMON (described in Chapters 13 and 14, respectively), IP (described in Chapter 25, “IP Routing,”) and IPX (described in Chapter 27, “IPX Routing”).

Interface. The main menu for configuring parameters and viewing statistics for switching modules. This menu has sub-menus for managing Frame Relay and Fast Ethernet switching modules. In addition it includes a sub-option for configuring SLIP. These sub-menus are described in Chapters 15 through 16 and Chapter 29.

Security. This menu contains options for changing a password and rebooting the system. It is described in Chapter 8, “Switch Security.”

System. Contains a wide array of options for configuring and viewing information on a variety of switch functions. Options include displays of switch slot contents, configuring serial ports, and viewing CAM information. Commands used to configure User Interface display options are described in User Interface Display Options on page 4-30. Other System menu commands are described in Chapter 9, “Configuring Switch-Wide Parameters.” The System menu also includes a sub-menu option that provides additional commands for configuring the MPX module. This sub-menu is described in Chapter 6, “Configuring Management Processor Modules.”

Services. Provides options for creating, modifying, viewing, and deleting Frame Relay services. Frame Relay services include bridging, routing, and trunking. Frame Relay services are described in Chapter 29, “Managing Frame Relay.”

Switch. Provides options to precisely define frame translations. A MAC-layer type (Ethernet, Token Ring, etc.) may have more than one type of frame format, such as Ethernet or 802.3. But, by default, each MAC-layer type defaults to certain frame format upon translation. This menu allows you to define translations for each frame format. This menu is described in Chapter 18, “Configuring Frame Translations.”

Help. Provides textual help on how to use the UI and on each menu or sub-menu. For the item of interest, enter

help <sub-menu name>

Diag. This menu, fully available to the **diag** login account, contains commands to run diagnostic tests. It is described in Chapter 36, “Running Hardware Diagnostics.”

Quit. Logs you out of the UI. You can also enter **logout** to exit.

? Displays the options for current menu.

General User Interface Guidelines

You can monitor and configure your switch in the following various ways:

- The User Interface (UI): The UI is the original method of switch configuration. It is a text-based and menu-driven interface to which you can connect through the serial port, through a modem, or over a network via Telnet. You can have up to four simultaneous UI connections to an Omni Switch/Router. For Release 4.4 and later, the default for switch monitoring and configuration is the CLI mode. If you are using a login account with permission to use the UI command, you can enter the UI mode by entering the **ui** command at the CLI system prompt.
- X-Vision: This purchasable network management software program consists of several powerful sub-applications that help you manage and monitor your network. X-Vision allows you to connect and configure multiple switches simultaneously. For more information, refer to X-Vision’s on-line help.
- The Command Line Interface (CLI): The CLI is a new feature included with Release 4.1 that allows you to configure Omni Switch/Routers using single-line text-based commands that are entered through the local console. Improved readability, easy text editing of the configuration files, and simple cloning of switch configurations are among some of the advantages of the CLI. For more information, refer to the *Text-Based Configuration CLI Reference Guide*.

Entering Command Names

The UI is not case sensitive for commands, meaning that you may enter upper or lower case as you desire. However, command line assignments, configuration input, and logins *are* case sensitive.

Except for the **logout** and **quit** commands, you only need to enter as much of the command that is unique. For example, if you want to execute the **switch** command you need only enter **swi**. If you enter only **sw**, the system will respond with a choice of the following:

switch swch swchmac swap

If you set the switch to the verbose mode you will see additional information on the screen (see Setting Verbose/Terse Mode for the User Interface on page 4-22).

Non-unique command match, possible commands:

switch	Enter Any to Any Switching Menu
swch	Configure Any To Any Switching Port Translations
swchmac	View Per Mac Translation Options
swap	Change swap status of chassis
swlogc	Configure Switch Logging source/destination mapping and priority levels

◆ Note ◆

If you cannot see a UI command confirmation prompt or if you do not get the command prompt after the completion of a command, press the **<Enter>** key to regain the prompt.

Quitting a Command

Many of the commands give you a list of parameters to change. With most commands you can enter in **quit** if you want to exit the command without making changes. If the **quit** parameter is not available, press **Ctrl-d** to abort the command without making changes.

Scrolling

If the screen scrolls up too far to read you can stop the incoming data by pressing **Ctrl-s**. The screen will stop and allow you to read the data. Press **Ctrl-q** to continue the data transmission.

The UI Configuration Menu

The User Interface (UI) Configuration menu consolidates the following UI commands into a single, easy-to-use menu:

- **chpr**
- **more**
- **ver**
- **ter**
- **timeout**

◆ Note ◆

The switch's *prompt*, *more*, *verbose/terse*, and *timeout* functions remain fully supported. However, if you enter any of the commands listed above, you will be redirected to the UI Configuration menu.

To access the UI Configuration menu, type

uic

at the system prompt and press **<Enter>**. The following screen will be displayed:

UI Configuration

```
1) Prompt : '$Menu-Path%'
2) More   : on
  21) Lines : 22 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Refer to the following sections for information on using the UI Configuration menu.

Configuring the System Prompt

The **uic** submenu is listed under the system menu. The **uic** submenu allows you to change the system prompt. The prompt can be made up of literal information, system variable information, or a combination of the two.

Literal information means that the prompt will reflect exactly what you type at the **uic** submenu. For example, **Marketing 1** or **Enter command:**.

System variable information means that the prompt will reflect the switch's variable information, such as the current menu-path or the system name. Use **\$Menu-Path** (case sensitive) to have the system prompt display the current menu-path name. Use **\$SysName** to have the system prompt display the system name.

You can also mix variables and literals such as **\$Menu-Path ->** or **\$SysName Enter command:**.

◆ Note ◆

The default system prompt is **->**.

To change the system prompt, type **uic** at the user prompt and press **<Enter>**.

A screen similar to the following will be displayed.

UI Configuration

```
1) Prompt : '$Menu-Path%'
2) More   : on
  21) Lines : 22 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Next, type **1=**, followed by the desired prompt information, and press **<Enter>**. For example:

```
1=$SysName ->
```

After you press **<Enter>**, the screen will be redrawn. Note that the prompt information at line 1 of the **uic** submenu has been changed.

UI Configuration

```
1) Prompt : '$SysName ->'
2) More   : on
  21) Lines : 22 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Type **save** at the submenu prompt and press **<Enter>**. The system prompt has been successfully changed.

Configuring More Mode for the User Interface

Enabling More Mode

The more mode allows you to specify the maximum number of lines that will be scrolled to your workstation's display. However, before you can specify the maximum number of lines that can be displayed, you must first verify that the more mode is enabled. To enable the more mode, type **uic** at the user prompt and press **<Enter>**. A screen similar to the following will be displayed.

UI Configuration

```
1) Prompt : '$Menu-Path%'
2) More   : off
  21) Lines : 22 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Next, type **2=on** at the submenu prompt and press **<Enter>**. The screen will be redrawn. Note that more mode is now set to **on**.

UI Configuration

```
1) Prompt : '$Menu-Path%'
2) More   : on
  21) Lines : 22 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

The switch's default output display is 22 lines. If you want to change this value, type **21=**, followed by the maximum number of lines to be displayed, and press **<Enter>**. For example:

21=50.

After you press **<Enter>**, the screen will be redrawn. Note that the output display value at line 21 of the **uic** submenu has been changed.

UI Configuration

```
1) Prompt : '$Menu-Path%'
2) More   : on
  21) Lines : 50 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Be sure to type **save** at the submenu prompt and press **<Enter>**. More mode is now enabled.

Changing the More Mode Line Value

If the switch's more mode has already been enabled and you want to change the maximum number of lines to be displayed on your workstation, type **uic** at the user prompt and press **<Enter>**.

A screen similar to the following will be displayed.

UI Configuration

- 1) Prompt : '\$Menu-Path% '
- 2) More : on
- 21) Lines : 22 lines
- 3) Verbose : off
- 4) Timeout : 5 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Type **21=**, followed by the maximum number of lines to be displayed, and press **<Enter>**. (The value may range from 0 to 2147483647.) For example:

21=2000.

After you press **<Enter>**, the screen will be redrawn. Note that the output display value at line 21 of the **uic** submenu has been changed.

UI Configuration

- 1) Prompt : '\$Menu-Path% '
- 2) More : on
- 21) Lines : 2000 lines
- 3) Verbose : off
- 4) Timeout : 5 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Type **save** at the submenu prompt and press **<Enter>**. The more mode line value has been successfully changed.

Disabling More Mode

To disable more mode, type **uic** at the user prompt and press **<Enter>**.

A screen similar to the following will be displayed.

UI Configuration

```
1) Prompt : '$Menu-Path% '
2) More   : on
  21) Lines : 22 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Next, type **2=off** at the submenu prompt and press **<Enter>**. The screen will be redrawn. Note that more mode is now set to **off**.

UI Configuration

```
1) Prompt : '$Menu-Path% '
2) More   : off
  21) Lines : 22 lines
3) Verbose : off
4) Timeout : 5 minutes
```

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Type **save** at the submenu prompt and press **<Enter>**. More mode is now disabled.

◆ Reminder ◆

The switch's table filtering feature *cannot* be used when the **more** mode is disabled. For more information on UI table filtering, refer to UI Table Filtering (Using Search and Filter Commands) on page 4-38.

Setting Verbose/Terse Mode for the User Interface

Enabling Verbose Mode

When verbose mode is enabled, you are not required to enter a question mark in order to view the switch's configuration menus. Instead, menus are displayed automatically. For example, if verbose mode is enabled and you enter

summary

at the user prompt, the Summary menu will be displayed automatically, as shown below:

<u>Command</u>	<u>Summary Menu</u>
ss	Display MIB-II System group variables
sc	OmniSwitch chassis summary
si	Current interface status

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The switch's default verbose mode setting is **off**, or disabled. To enable verbose mode, type **uic** at the user prompt and press **<Enter>**.

A screen similar to the following will be displayed.

UI Configuration

- 1) Prompt : '\$Menu-Path% '
- 2) More : on
 - 21) Lines : 22 lines
- 3) Verbose : off
- 4) Timeout : 5 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Next, type **3=on** at the submenu prompt and press **<Enter>**. The screen will be redrawn. Note that verbose mode is now set to **on**.

UI Configuration

- 1) Prompt : '\$Menu-Path% '
- 2) More : on
 - 21) Lines : 22 lines
- 3) Verbose : on
- 4) Timeout : 5 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Type **save** at the submenu prompt and press **<Enter>**. You will be returned to the user prompt. Verbose mode is now enabled.

Disabling Verbose Mode

Although the **terse** command is no longer supported as of Release 4.1, disabling verbose mode via the **uic** submenu is the command equivalent. When verbose mode is disabled, configuration menus *will not* be displayed automatically. To display a current menu when verbose mode is disabled, you must type a question mark (?) and then press **<Enter>**.

To disable verbose mode, type **uic** at the user prompt and press **<Enter>**.

A screen similar to the following will be displayed.

UI Configuration

- 1) Prompt : '\$Menu-Path% '
- 2) More : on
 - 21) Lines : 22 lines
- 3) Verbose : on
- 4) Timeout : 5 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Next, type **3=off** at the submenu prompt and press **<Enter>**. The screen will be redrawn. Note that verbose mode is now set to **off**.

UI Configuration

- 1) Prompt : '\$Menu-Path% '
- 2) More : on
 - 21) Lines : 22 lines
- 3) Verbose : off
- 4) Timeout : 5 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :

Type **save** at the submenu prompt and press **<Enter>**. Verbose mode is now disabled.

Configuring the Auto Logout Time

When the switch detects no user activity on the UI for a certain period of time, it automatically logs the user out of the system. By default, this automatic logout occurs after 4 minutes of console inactivity. You can configure the automatic logout to range from 1 minute to 35,791,394 minutes.

To set a new automatic logout time, type **uic** at the user prompt and press **<Enter>**.

A screen similar to the following will be displayed.

```

                                UI Configuration
1) Prompt   : '$Menu-Path%'
2) More     : off
  21) Lines  : 22 lines
3) Verbose  : off
4) Timeout  : 5 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :
```

Next, type **4=on**, followed by the desired automatic logout time, and press **<Enter>**. For example:

```
4=15.
```

After you press **<Enter>**, the screen will be redrawn. Note that the automatic logout time at line 4 of the **uic** submenu has been changed.

```

                                UI Configuration
1) Prompt   : '$Menu-Path%'
2) More     : on
  21) Lines  : 22 lines
3) Verbose  : off
4) Timeout  : 15 minutes

Command {Item=Value/?/Help?Quit?Redraw?Save} (Redraw) :
```

Be sure to type **save** at the submenu prompt and press **<Enter>**. The automatic logout time has been successfully changed.

◆ Note ◆

The automatic logout value you enter takes effect immediately; you do not have to reboot the switch. In addition, the timeout parameter you enter is saved. Later sessions using this account will have the same automatic logout parameter until you change it.

Viewing Commands

If at any time you are not sure of the commands available, enter **?** and you will be given a list of the commands in the current sub-menu. Following each list of commands is a list of sub-menus. You can go directly to any sub-menu in the list.

You can specify whether the full menu will be displayed when you enter a command for a menu or sub-menu and the amount of information you receive when you run the help command. (Refer to Setting Verbose/Terse Mode for the User Interface on page 4-22 for more information.) Additionally, there is a lookup facility to assist with administrative tasks. You can look up any command name or prefix as follows:

lookup vlans

or to see all commands starting with **v** use:

lookup v*

To see all commands available, enter:

lookup *

Changing Passwords

The **pw** command is used to change passwords and is described in Chapter 8, “Switch Security.”

Command History and Re-Executing Commands

The **history** command displays up to 50 commands numbered in order with the most recently executed command listed last. The following is a typical example of the **history** command.

```
1: view mpx.cmd
2: vlan
3: at
4: atvl
5: vimcvi
6: mcvi
7: vivi
8: fwtvl
9: xlat
10: history
```

In the example above, the **history** command is listed last because it is the one that was executed most recently. If you want to re-execute the last command, enter two exclamation points (!!). In the example above, you could re-execute the **history** command by entering

```
!!
```

at the system prompt.

You can also display a specific number of commands by entering **history** followed by a number less than or equal to the number of commands in the history buffer. For example, if you entered

```
history 5
```

in the example above you would see the following:

```
7: vivi
8: fwtvl
9: xlat
10: history
11: history 5
```

The UI also provides several other ways to re-execute earlier commands. For example, you can re-execute a specific command shown in the **history** list by entering an exclamation point (!) followed by the number to the left of that command shown in the **history** list. In the example at the beginning of this section, entering

```
!2
```

would re-execute the **vlan** command.

You can also re-execute a command a set number of commands back by entering an exclamation point and a minus sign (!-) followed by that set number of commands back. In the example at the beginning of this section, entering

```
!-3
```

would re-execute the **fwtvl** command.

In addition, you can re-execute a command by entering an exclamation point (!) followed by the first character(s) of the most recently executed command. In the example at the beginning of this section, entering

!vim

would re-execute the **vimcvl** command. Entering

!vi

however, would re-execute the **vi** command because it is the most recently executed command beginning with **vi**.

You can also re-execute the most recently executed command containing a string of characters by entering an exclamation point and a question mark (!?), followed by the string of characters, and an optional question mark (?) which acts as a “wild card.” In the example at the beginning of this section, entering

!?!an?

at the system prompt would re-execute the **vlan** command. Entering

!?!a?

however, would re-execute the **xlat** command because it is the most recently executed command containing **la**.

Commands in the history buffer can be modified by adding a parameter, when it is applicable. For example, if you entered

!7 3/1

in the example at the beginning of this section you would execute the command **vim 3/1**.

Abbreviating IP Addresses

The Omni Switch/Router software provides the user with a more concise way to enter the dotted decimal format of a 32-bit IP address. The new syntax conforms to the traditional Internet interpretation. Several examples of abbreviated IP addresses are shown in the table below. The first column of the table lists examples of abbreviated IP addresses, and the second column shows how the system interprets the abbreviated address.

Abbreviated IP Address Formats

Sample User Entry	IP Address
198	0.0.0.198
198.	198.0.0.0
198..	198.0.0.0
198...	198.0.0.0
198.206	198.0.0.206
198..206	198.0.0.206
198..206.	198.0.206.0
198...206	198.0.0.206
198.206.	198.206.0.0
198.206..	198.206.0.0
198.206.182	198.206.0.182
198..206.182	198.0.206.182
198.206..182	198.206.0.182
198.206.182.	198.206.182.0
198.206.182.158	198.206.182.158

As shown in the table above, the system performs two important steps to ensure that the IP address is valid. First, it puts zeroes when you do not specify the number. Second, the system will insert as many zeroes as needed to the right of a period.

This abbreviated IP address format can be used with the **ftp**, **telnet**, **crpg**, **modvl**, **ping**, **snmpc**, and **xlat** commands. For example, to ping the IP address 198.0.0.2, you can abbreviate this IP address by entering

```
ping 198.2
```

at the system prompt. After you answer a few prompts (see Chapter 25, “IP Routing” for more information on the **ping** command), something similar to the following will be displayed.

```
Ping starting, hit <Enter> to stop  
PING 198.0.0.2: 64 data bytes
```

```
[0 ] T
```

```
----198.0.0.2 PING Statistics----
```

```
1 packets transmitted, 0 packets received, 100% packet loss
```

In addition, the IP subnet mask 255.255.0.0 can be abbreviated in the following ways:

- 255.255.
- 255.255..

User Interface Display Options

The System menu several commands to configure help information, character display, and the system prompt for the UI. Enter

system

at the system prompt to enter the System menu. Press the question mark (?) to see the System menu commands, as shown below.

<u>Command</u>	<u>System Menu</u>
info	Basic info on this system
dt	Set system date and time
ser	View or configure the DTE or DCE port
mpm	Configure a Management Processor Module
slot	View Slot Table information
systat	View system stats related to system, power and environment
taskstat	View task utilization stats
memstat	View memory use statistics
fsck	Perform a file system check on the flash file system
newfs	Erase all files from /flash & create a new file system
syscfg	View/Configure info related to this system
uic	UI configuration; change - prompt, timeout, more, verbose.
camstat	View CAM info and usage
camcfg	Configure CAM info and usage
hrex	Enter HRE-X management command sub-menu
ver/ter	Enables/disables automatic display of menus on entry (obsolete, use 'uic' command)
echo/noecho	Enable/disable character echo
chpr	Change the prompt for the system (obsolete, use 'uic' command)
logging	View system logs.
health	Set health parameters or view health statistics
cli/exit	Enter command line interface
saveconfig	Dump the cache configuration content to the mpm.cfg file.
cacheconfig	Set the flag to use cache configuration only.

Main File Summary VLAN Networking
Interface Security System Services Help

For information on the **info**, **dt**, **ser**, **slot**, **systat**, **taskstat**, **memstat**, **fsck**, **newfs**, **syscfg**, **camstat**, **camcfg**, and **hrex** commands, refer to Chapter 9, “Switch-Wide Parameters.” The **mpm** command is described in Chapter 6, “Configuring Management Processor Modules.” The **ver/ter** and **chpr** commands are described earlier in Setting Verbose/Terse Mode for the User Interface on page 4-22. The **echo/noecho** command is described in the following section. The **cli** command is described earlier in Changing Between the CLI and UI Modes on page 4-2. The **logging** command is described in Chapter 10, “Switch Logging.”

◆ **Note** ◆

The **ver/ter**, and **chpr** commands now appear as items in the UI Configuration menu (displayed through the **uic** command). If you enter the **ver/ter** and **chpr** commands, a message will advise you to use the **uic** command, and the UI Configuration menu will automatically display. For more information on the UI Configuration menu, refer to The UI Configuration Menu on page 4-17.

Setting Echo/NoEcho for User Entry

You can determine whether your entries will appear by enabling the echo for user entries. The default is to echo all characters.

To enable the echo, enter

```
echo
```

at the system prompt. Everything you enter will be displayed. For example, if you enter

```
history
```

at the system prompt, it will be displayed on your terminal, as shown in the example below.

```
/%history
```

If your terminal echoes characters locally it is a good idea to set the UI to **noecho** to avoid repeated characters. To disable the echo, enter

```
noecho
```

at the system prompt. For example, if your terminal echoes characters locally, you would see something like the following if you entered **history**.

```
/%history
```

If your terminal does not echo characters locally, nothing you enter will be displayed. For example, if you enter

```
history
```

at the system prompt, it will *not* be displayed on your terminal, as shown in the example below.

```
/%
```

Setting the Login Banner

The login banner feature allows you to change the banner that displays whenever someone logs into the UI. This feature can be used to display messages about user authorization and security. You can display the same message for all login sessions or you can display different messages for login sessions initiated by the console, ftp or Telnet access. The default login message looks like this:

```
This product includes software developed by the University of California  
Berkeley and its contributors.
```

```
Welcome to the Alcatel Omni Switch/Router ! Version 4.4
```

```
login:
```

Here is an example of a banner that has been changed:

```
This product includes software developed by the University of California  
Berkeley and its contributors.
```

```
*** LOGIN ALERT ***
```

```
This is a secure device. Unauthorized use of this switch will result  
in criminal prosecution.
```

```
login:
```

Creating a new Banner

Three steps are required to change the login banner. They are listed here.

- Create a text file containing the new banner in the switch's flash directory.
- Add the **UI_add_do_alert()** command syntax to the switch's `mpx.cmd` file.
- Enable the feature by executing the **alert {console | telnet | ftp}** command.

To create the text file containing your banner you may use the **create file** command in the UI's edit buffer sub-menu. This method allows you to create the file in the flash directory without leaving the UI console session. You can also create the text file in an external editor (such as MS Wordpad) and ftp the file to the switch's flash directory. In either case, be sure to remember the name of your file.

To add the **ui_add_do_alert()** command syntax to the switch's `mpx.cmd` file, use the edit command of the UI's **file** sub-menu. (For information on using the file sub-menu, refer to Chapter 7, "Managing Files").

To enable the new login banner, add the **alert {console |telnet | ftp}** syntax to the `mpx.cmd` file, using the **edit** command of the UI's **file** sub-menu. This command will cause the banner message to display at each login until the switch is rebooted. After a reboot, the switch will not display the banner unless the **alert** command is executed again.

Permanent Banner

If you want the banner message to display after the system has been rebooted, you must add additional lines to the `mpx.cmd` file. The following example lists the commands you must add to the `mpx.cmd` file. This example uses a banner text file with the name **"banner.txt"**.

```
cmDoDump=1
cmlnit
ui_add_do_alert()
change_prompt_file("console", "banner.txt")
change_prompt_file("telnet", "banner.txt")
```

◆ Note ◆

Any commands added to the `mpx.cmd` file must be added after the lines **cmDoDump=1** and **cmlnit**. If the commands in the `mpx.cmd` file are not in the proper order the switch may not boot properly.

Banners for Different Access Methods

You may use different banners for sessions accessed by console, Telnet or ftp methods. To do this, create different text files for each banner with unique filenames. When you add the commands to the `mpx.cmd` file, use the filenames to associate the banner with the session access methods. Here is an example:

```
cmDoDump=1
cmlnit
ui_add_do_alert()
change_prompt_file("console", "console_banner.txt")
change_prompt_file("telnet", "telnet_banner.txt")
change_prompt_file("ftp", "telnet_banner.txt")
```

Login Accounts

The UI provides three default login accounts—Administrator, User and Diagnostics. The Administrator login provides full access to all functions. The initial login name for an Administrator account is **admin**. The Diagnostics login also has full access to all switch functions plus a special sub-menu with a set of switching module tests. The initial login name for Diagnostics is **diag**. The User login has read-only privileges to the switch. The initial login name for a User account is **user**. The password for each of these default login accounts is **switch**.

◆ **Note** ◆

In software release 4.3, the **user** login account with read-only privileges is not included automatically.

◆ **Note** ◆

You can configure new and delete existing login accounts with the **useradd** UI command, that is described in Chapter 8, “Switch Security.”

Multiple User Sessions

You can have up to four simultaneous connections to an Omni Switch/Router. One connection can be made to the console port, two can be made through Telnet, and one connection can be made to the modem port if you are connecting to an Omni Switch/Router.

◆ **Note** ◆

For software Releases 4.4 and later, more than one login account with write privileges *can be* active at the same time.

For software Release 4.3 and earlier, only one login account with write privileges was allowed on the switch at the same time. In this case, the first switch user who logged on as either **admin** or **diag** would be the only user with the write privilege. Subsequent users who logged on as either **admin** or **diag** would not have the write privilege and would be unable to perform any functions that change switch parameters. These users would also see a message that informs them they do not have the write privilege when they log on. For example, a user who logs on as **admin** when another user already has the write privilege will see the following message:

You are logged in as 'admin' without the WRITE privilege.

The WRITE privilege is currently in use by another user.

However, users who log on as either **admin** or **diag** without the write privilege can “kill” the session of the user with the write privilege and gain that privilege for themselves. This is described in Deleting Other Sessions on page 4-35.

If you try to log on when the limit of user has been reached (e.g., you attempt a Telnet connection when there are two users currently connected through Telnet), you will see the following message:

Sorry, reached maximum number of sessions.

Listing Other Users

To display all the users currently logged on to the switch, type

```
who
```

at the system prompt. The following is an example of the display shown where two Telnet sessions are logged in, one as **admin** and the other as **user**.

SESSION	USER	READ	PRIVILEGES WRITE	GLOBAL	TTY
3	admin (123.456.78.910)	00000008007fff	00000008007fff	0000000007ffff	/pty/telnetA
4	rrtest1 (123.456.78.910)	00000008007fff	00000008007fff	0000000000000000	/pty/telnetB

You can also display information about just your session by typing

```
who am i
```

at the system prompt. The following is a typical example of the output.

SESSION	USER	READ	PRIVILEGES WRITE	GLOBAL	TTY
3	admin (123.456.78.910)	00000008007fff	00000008007fff	0000000007ffff	/pty/telnetA

The following sections describe the parameters shown by the **who** command.

SESSION. The session number of the user. A **0** indicates that the user is connected through the console port, a **1** indicates that the user is connected through the modem port, and a **2** or **3** indicates that the user is connected through Telnet. The session number is used with the **write** and **kill** commands described in Communicating with Other Users on page 4-35 and Deleting Other Sessions on page 4-35, respectively.

USER. The administrative level of the user. This will be **admin**, **user** or **diag**.

PRIVILEGES. The privilege level of the user. The **READ**, **WRITE** and **GLOBAL** privileges are indicated in hexadecimal numbers.

TTY. Type of connection. This shows whether the user is connected by Telnet, the modem port, or the console port. If the connection is via Telnet, the IP address of the connecting workstation is also shown.

Communicating with Other Users

If you want to send a message to another user, enter **write** followed by the user's session number. If you wanted to send a message to a user connected on the console port (session 0), you would enter

```
write 0
```

at the system prompt. The switch would then display

```
Enter message. (End with CTRL-D or 'exit')
```

Everything you type now will be sent to the user connected on the console port until you press **CTRL-D** or enter **exit** on a line by itself. Here is an example of the **write** command:

```
write 0
I need the write privilege
exit
```

The user receiving the message would see the following:

```
Message from user 'admin' on session 3.
I need the write privilege
End of message.
```

If you enter an invalid session number, the switch will display an error message. For example, if you entered

```
write 1
```

at the system prompt and no user was connected through the modem port (session 1), the switch would display

```
ERROR: Session 1 is an invalid session number.
```

Note

After you have received a message or after you have written a message you must press the **<Enter>** key to regain the system prompt.

Deleting Other Sessions

If you are logged on as **admin** or **diag**, you can kill the session of another user. For example, if you want the write privilege and you are logged on as **diag** or **admin**, you must end the session of the user who currently has the write privilege with the **kill** command. The syntax for the **kill** command is as follows:

```
kill [-t <timeout>] -f <session_number>
```

The **session_number** is assigned by the switch and can be displayed with the **who** command, which is described in Listing Other Users on page 4-34. If you do not use the **-f** option, then the system will wait until the other user presses **<Enter>** or finishes his current command. If you do use this option, then the other user's session will be terminated immediately.

The **-t** option can be used with the **-f** option to set the amount of time before the other user's session is terminated. See Advanced Kill Command Options on page 4-37 for descriptions of the **-f** and **-t** options.

Multiple User Sessions

For example, to end the session of the user connected to the console port (session 0) and let him finish his current command, you would enter

kill 0

at the system prompt. The system would then display something similar to the following:

Press <Enter> to cancel.

Trying.....

The user losing the write privilege would see something similar to the following:

**Your session will be killed by user 'admin' on session 3
as soon as you finish this command or press return.**

After the user with the session being killed has finished his work, he will be logged off. If the user who was logged off had the write privilege, you will gain the write privilege and a message similar to the following will be displayed.

Done.

You have gained the WRITE privilege

You can use the **who** command to confirm that you now have the write privilege.

In addition, the session number used in the **kill** command must be valid. If, for example, you entered

kill 1

and no user was connected to the modem port (session 1), the system would display the following:

ERROR: Session 1 is an invalid session number.

Also, you cannot use the **kill** command to end your own session. For example, if your session number is **3** and you entered

kill 3

the system would display the following:

ERROR: You cannot kill your own session.

Instead, use the **quit** or **logout** command if you want to log out.

Advanced Kill Command Options

You can also kill the session of a user immediately by adding the parameter **-f** followed by the session number of the user. This option will kill the user's session before he can finish his current command. In addition, this option will end the user's sessions without waiting for him to press **<Enter>**. This option can be used to log off a user with the write privilege who forgot to log out and then gain the write privilege for yourself.

If you wanted to kill the session of the user with a session number of 2 immediately, you would enter

```
kill -f 2
```

at the system prompt.

The default timeout for the **kill** command is 2 seconds. You can modify the duration of the timeout by using **-t** option in conjunction with the **-f** option. To use the timeout option, enter **kill**, followed by **-t**, the number of seconds for the timeout, **-f**, and the session number of the user. For example, if you wanted to kill the session of the user with a session number of 2 in 15 seconds, you would enter

```
kill -t15 -f 2
```

at the system prompt. The valid range for the timeout is 1 to 240 seconds.

◆ Note ◆

You *cannot* use the timeout option (**-t**) unless you also use the **-f** option.

UI Table Filtering (Using Search and Filter Commands)

The amount of information displayed in UI tables can be extensive, especially with larger networks. Common UI commands, such as **ipr**, **vipl**, **macinfo**, and **fwl**, often return multi-page tables. The user can locate specific information in these large tables through the **More?** UI prompt.

The **More?** prompt appears whenever the maximum number of table entries designated by the **more** command has been reached (the **more** command's default is 22 lines). Note that if a table exceeds 22 lines, and the **more** mode has been configured to display *more than* 100 lines, the following message appears:

Screen Size larger than 100 Lines, Displaying with 22 Lines (Press Any Key)

After pressing any key, only the page of the table is displayed, followed by the **More?** prompt.

◆ Important Note ◆

The switch's **more** mode is active by default. If the **more** mode is turned off, the Search and Filter commands cannot be used. For more information on the **more** command, see The UI Configuration Menu on page 4-17.

A typical **More?** UI prompt will look like this:

```
1 4/6 Brg/ 1/ na 0020da:030995 Tns DFLT Enabl Inactv Disabl AutoSw
1 4/7 Brg/ 1/ na 0020da:030996 Tns DFLT Enabl Inactv Disabl AutoSw
1 4/8 Brg/ 1/ na 0020da:030997 Tns DFLT Enabl Inactv Disabl AutoSw
1 5/1 Brg/ 1/ na 0020da:954050 Tns DFLT Enabl Inactv Disabl AutoSw
More? [<SP>,<CR>,/,F,N,Q,?]
```

At the **More?** prompt, the user is given a list of options, which includes the Search (**I**) and Filter (**F**) commands:

- <SP>** Press **<SP>** (space bar) to display the next page of information.
- <CR>** Press **<CR>** (character return) to display the next line of information.
- /** Press **/** to enter the Search mode.
- F** Press **F** to enter the Filter mode.
- N** Press **N** to renew the search, starting from the next line in the UI table.
- Q** Press **Q** to exit the **More?** prompt.
- ?** Press **?** to enter the **More?** command Help Menu.

These commands are available for **admin** and **diag** login sessions. Please refer to the following sections for more information on the Search and Filter commands, as well as renewing a search, combining Search and Filter commands, and using wildcards.

The Search Command

Starting from the page being displayed, the Search command (*/*) searches all lines of a UI table for a specified text pattern (up to 80 characters). The first line containing the pattern is brought to the top of the page, followed by any remaining lines in the table.

Searches *cannot* be limited to a specific column or heading.

To use the Search command, type */* at the **More?** prompt, followed by the text pattern you are looking for, then press **<Enter>**.

◆ Important Note ◆

The Search command is case sensitive. When using this command, be sure to type the text pattern exactly as it would appear in the UI table.

Real World Example

The following example uses the Search command to locate a specific MAC address in the **macinfo** table. (Before using this example, be sure that the **more** mode is enabled and the default is set at 22 lines. For more information, refer to page 4-38.)

1. Type **macinfo** and press **<Enter>**. The following screen will be displayed:

Enter MAC address ([XXYYZZ:AABBCC] or return for none) :

Press **<Enter>** again. A screen similar to the following will be displayed:

Enter Slot Number (1-5) :

Type the slot number for the module containing the relevant MAC address information (e.g. **3**), then press **<Enter>**. A table similar to the following will be displayed:

Total number of MAC addresses learned for this slot: 58

Sl/ If/ Service/ In	MAC Address	Non-Canonical MAC Address	T	Group ID	CAM Indx	S	Last Seen	Exp Timer
3/ 1/ Brg/ 1	00A0C9:064D04	000593:60B220	E	1	7024	T	134	300
3/ 1/ Brg/ 1	006008:C1D7C2	000610:83EB43	E	1	7030	T	115	300
3/ 1/ Brg/ 1	0020DA:88F110	00045B:118F08	E	1	70E6	T	46	300
3/ 1/ Brg/ 1	0020DA:B6FF12	00045B:6DFF48	E	1	7094	T	66	300
3/ 1/ Brg/ 1	0020DA:8A7DC0	00045B:51BE03	E	1	705A	T	83	300
3/ 1/ Brg/ 1	0020DA:A67FA2	00045B:65FE45	E	1	7120	T	27	300
3/ 1/ Brg/ 1	0020DA:024F75	00045B:40F2AE	E	1	710C	T	34	300
3/ 1/ Brg/ 1	0020DA:9B88E4	00045B:D91127	E	1	70EE	T	45	300
3/ 1/ Brg/ 1	0020DA:9C062B	00045B:3960D4	E	1	7074	T	76	300
3/ 1/ Brg/ 1	0020DA:79F062	00045B:9E0F46	E	1	70D2	T	52	300
3/ 1/ Brg/ 1	006008:991CA7	000610:9938E5	E	1	701C	T	117	300
3/ 1/ Brg/ 1	0020DA:936A8F	00045B:C956F1	E	1	712A	T	23	300
3/ 1/ Brg/ 1	0020DA:9CEAC5	00045B:3957A3	E	1	70CC	T	53	300
3/ 1/ Brg/ 1	0020DA:9B9B54	00045B:D9D92A	E	1	70D6	T	50	300
3/ 1/ Brg/ 1	0020DA:7AAE24	00045B:5E7524	E	1	70B8	T	58	300
3/ 1/ Brg/ 1	0020DA:A9EEB3	00045B:9577CD	E	1	710A	T	34	300
3/ 1/ Brg/ 1	0020DA:8DB20B	00045B:B14DD0	E	1	7080	T	72	300
3/ 1/ Brg/ 1	0020DA:9F6B82	00045B:F9D641	E	1	70F4	T	42	300
3/ 1/ Brg/ 1	0020DA:8762A3	00045B:E146C5	E	1	7126	T	24	300
3/ 1/ Brg/ 1	006008:C1D7C2	000610:83EB43	E	1	7030	T	115	300

More? [**<SP>**,**<CR>**,/,F,N,Q,?]

Note that, because the information in the table exceeds the **more** command's default page size of 22 lines, the **More?** prompt appears at the bottom of the screen.

UI Table Filtering (Using Search and Filter Commands)

2. Type `/` at the **More?** prompt. The Search prompt (`/`) will appear automatically. At the Search prompt, enter the text pattern for the desired MAC address. For example:

```
/0020DA:9E479D
```

Press **<Enter>**. A screen similar to the following will be displayed:

```
Searching .....
```

```
3/ 1/ Brg/ 1 0020DA:9E479D 00045B:79E2B9 E 1 702C T 138 300
3/ 1/ Brg/ 1 0020DA:9D0D1B 00045B:B9B0D8 E 1 7030 T 67 300
3/ 1/ Brg/ 1 0020DA:97CDE0 00045B:E9B307 E 1 70E6 T 122 300
3/ 1/ Brg/ 1 00A0C9:8DED5B 000593:B1B7DA E 1 7094 T 114 300
3/ 1/ Brg/ 1 0020DA:92A152 00045B:49854A E 1 705A T 97 300
3/ 1/ Brg/ 1 0020DA:8528D5 00045B:A114AB E 1 7120 T 102 300
3/ 1/ Brg/ 1 0020DA:93BF73 00045B:C9FDCE E 1 710C T 130 300
3/ 1/ Brg/ 1 0020DA:B956B5 00045B:9D6AAD E 1 70EE T 56 300
3/ 1/ Brg/ 1 0020DA:730F03 00045B:CEF0C0 E 1 7074 T 68 300
3/ 1/ Brg/ 1 0020DA:8BA710 00045B:D1E508 E 1 70D2 T 99 300
```

Note that the line containing information for the specified MAC address (**0020DA:9E479D**) now appears at the top of the screen, followed by any remaining lines in the UI table. (In this case, the last line of the **macinfo** UI table contains MAC address **0020DA:8BA710**, as shown).

Renewing a Search

If you execute the Search command and the resulting page still exceeds the maximum number of table entries designated by the **more** command, you can renew the Search. Do this by typing **n** at the **More?** prompt. The Search command will scan the remainder of the table and display the next line containing the desired text pattern at the top of the screen.

The Filter Command

The Filter command filters unwanted information from a UI table by displaying only those lines containing a specified text pattern (up to 80 characters). Once the Filter command has been executed, the Filter mode remains active until the end of the UI table has been reached, or until the user exits the current UI table.

Like the Search command, the Filter command *cannot* be limited to a specific column or heading.

To use the Filter command, type **f** at the **More?** prompt, followed by the text pattern you want displayed in the UI table, then press **<Enter>**.

◆ Important Note ◆

The Filter command is case sensitive. When using this command, be sure to type the text pattern exactly as it would appear in the UI table.

Real World Example

The following example uses the Filter command to display only those lines containing Lane services in the **vi** table. (Before using this example, be sure that the **more** mode is enabled and the default is set at 22 lines. For more information, refer to page 4-38.)

1. Type **vi** and press **<Enter>**. A table similar to the following will be displayed:

Virtual Interface VLAN Membership					
Slot / Intf / Service / Instance	Group	Member of VLAN#			
1 /1 /Rtr /1	1	1			
1 /1 /Rtr /2	33	1			
1 /1 /Rtr /3	111	1			
1 /1 /Rtr /4	33	2			
1 /1 /Rtr /5	1	3			
1 /1 /Rtr /6	1	4			
1 /1 /Rtr /7	33	7			
1 /1 /Rtr /8	33	3			
1 /1 /Rtr /9	1	5			
1 /1 /Rtr /10	1	6			
1 /1 /Rtr /11	33	5			
1 /1 /Rtr /12	33	6			
1 /1 /Rtr /13	999	1			
2 /1 /Lne /1	1	1			
2 /1 /Lne /2	111	1			
3 /1 /Brg /1	33	1 4			
3 /2 /Brg /1	1	1			
3 /3 /Brg /1	1	1			
3 /4 /Brg /1	1	1			

More? [**<SP>**,**<CR>**,/,F,N,Q,?]

Note that, because the information in the table exceeds the **more** command's default of 22 lines, the **More?** prompt appears at the bottom of the screen.

UI Table Filtering (Using Search and Filter Commands)

2. Type **f** at the **More?** prompt. The Filter prompt (**f/**) will appear automatically. At the Filter prompt, enter the desired text pattern (remember to type the text pattern exactly as it would appear in the UI table):

f/Lne

Press **<Enter>**. A screen similar to the following will be displayed:

Filtering

```
2 /1 /Lne /1 1 1
2 /1 /Lne /2 111 1
/ %
```

Note that only those lines containing Lane services are now displayed on the screen. All other table entries have been filtered from the UI.

Combining Search and Filter Commands

If you receive a **More?** prompt after using the Filter command, the filtered information still exceeds the maximum number of table entries designated by the **more** command. To further refine your results, you can combine the Search and Filter commands.

To combine the Search and Filter commands, type **/** at the Filter mode's **More?** prompt, followed by a revised text pattern of up to 80 characters. Note that you can combine the Search and Filter commands only after you have executed a Filter command *and* received a **More?** prompt at the bottom of the resulting page.

◆ Reminder ◆

Both the Search and Filter commands are case sensitive. When using these commands, be sure to type the text pattern exactly as it would appear in the text UI table.

Real World Example

The following example combines the Search and Filter commands to find specific IP address information in the **ipr** table. (Before using this example, be sure that the **more** mode is enabled and the default is set at 22 lines. For more information, refer to page 4-38.)

1. Type `ipr` and press `<Enter>`. A table similar to the following will be displayed:

IP ROUTING TABLE

128 routes in routing table

Network	Mask	Gateway	Metric	Group:VLAN	
				Id	Protocol
155.5.0.0	255.255.0.0	155.5.4.33	1	1:5	DIRECT
155.6.0.0	255.255.0.0	155.6.4.33	1	1:6	DIRECT
155.155.0.0	255.255.0.0	155.155.4.33	1	1:1	DIRECT
172.17.0.0	255.255.0.0	172.17.6.122	1	999:1	DIRECT
172.31.0.0	255.255.0.0	172.31.4.33	1	33:3	DIRECT
172.32.0.0	255.255.0.0	172.32.4.33	1	33:2	DIRECT
172.33.0.0	255.255.0.0	172.33.4.33	1	33:1	DIRECT
172.35.0.0	255.255.0.0	172.35.4.33	1	33:5	DIRECT
172.36.0.0	255.255.0.0	172.36.4.33	1	33:6	DIRECT
172.37.0.0	255.255.0.0	172.37.4.33	1	33:7	DIRECT
172.111.0.0	255.255.0.0	172.111.4.33	1	111:1	DIRECT
198.168.12.0	255.255.0.0	192.168.12.1	1	1:1	DIRECT
198.168.13.0	255.255.0.0	192.168.13.1	1	1:1	DIRECT

More? [`<SP>`;`<CR>`;`!``F``N``Q``?`]

Note that, because the information in the table exceeds the `more` command's default of 22 lines, the `More?` prompt appears at the bottom of the screen.

2. Use the Filter command to display all IP network addresses within the **IP Routing** table that contain **198**. To do this, type `f` at the `More?` prompt, followed by the specified text pattern:

`f/198`

Press `<Enter>`. A screen similar to the following is displayed:

Filtering

198.168.12.0	255.255.0.0	198.168.12.1	1	1:1	DIRECT
198.168.13.0	255.255.0.0	198.168.13.1	1	1:1	DIRECT
198.168.236.0	255.255.0.0	172.16.255.254	4	1:1	DIRECT
198.168.237.0	255.255.0.0	172.16.255.254	4	1:1	DIRECT
198.168.238.0	255.255.0.0	172.16.255.254	4	1:1	DIRECT
198.168.239.0	255.255.0.0	172.16.255.254	4	1:1	DIRECT
198.168.240.0	255.255.0.0	172.16.255.254	4	1:1	DIRECT
198.168.241.0	255.255.0.0	172.16.255.254	4	1:1	DIRECT
198.168.242.0	255.255.0.0	172.16.255.254	4	1:1	DIRECT
198.206.181.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.183.0	255.255.255.0	172.16.255.254	3	1:1	DIRECT
198.206.184.0	255.255.255.0	172.16.255.254	3	1:1	DIRECT
198.206.185.0	255.255.255.0	172.16.255.254	3	1:1	DIRECT
198.206.186.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.187.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.188.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.189.0	255.255.255.0	172.16.255.254	3	1:1	DIRECT
198.206.190.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.191.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.192.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.193.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT
198.206.194.0	255.255.255.0	172.16.255.254	2	1:1	DIRECT

More? [`<SP>`;`<CR>`;`!``F``N``Q``?`]

Because the filtered information in the table still exceeds the `more` command's default of 22 lines, the `More?` prompt appears at the bottom of the screen.

UI Table Filtering (Using Search and Filter Commands)

3. In order to further refine your results, you can now combine the Search and Filter commands. In this example, you will search for IP addresses beginning **198.206.2**. To do this, enter */* at the Filter mode's **More?** prompt, followed by the specified text pattern:

```
/198.206.2
```

Press **<Enter>**. A screen similar to the following is displayed:

Filtering and Searching ...

```
198.206.200.0 255.255.255.0 172.16.255.254 2 1:1 DIRECT
198.206.201.0 255.255.255.0 172.16.255.254 2 1:1 DIRECT
198.206.202.0 255.255.255.0 172.16.255.254 2 1:1 DIRECT
198.206.203.0 255.255.255.0 172.16.255.254 2 1:1 DIRECT
/Networking/IP %
```

Note that the IP address, **198.206.200.0**, now appears at the top of the screen, followed by any remaining lines in the table. (In this case, the last line of the **ipr** table contains information for IP address **198.206.203.0**, as shown).

Using Wildcards with Search and Filter Commands

Wildcards allow users to substitute symbols (***** or **?**) for text patterns while using the Search and Filter commands.

Any number of wildcards can be used within a single search string. In addition, multiple character (*****) and single character (**?**) wildcards can be combined within a single search string.

Wildcard Command Options

Multiple Characters

An asterisk (*****) is used as a wildcard for multiple characters in a text pattern. For example, the Filter pattern

```
/*.img
```

will filter out all lines from the UI table except those containing any text followed by **.img**.

This wildcard can also be used *within* a specific text pattern. For example, the Filter pattern

```
/1*6
```

will filter out all lines from the UI table except those containing **1**, followed by any number of characters, then **6**. For example:

```
1:3/6
```

or

```
33:3/1 Virtual port (#66)
```

or

```
16.
```

Single Characters

A question mark (?) is used as a wildcard for a single character in a text pattern. For example, the Search pattern

f/127.?0.1

will locate the first line in a UI table containing **127.** followed by *any single character*, and then the remaining text pattern **.0.1**. For example:

127.0.0.1.

◆ Note ◆

If you use a wildcard at the Search command and the resulting page still exceeds the maximum number of table entries designated by the **more** command, you can renew the search, starting from the next line containing the text pattern. Do this by typing **n** at the **More?** prompt. Note that you can renew a search only while in Search and Search/Filter modes.

5 Installing Switch Software

User Interface software comes pre-loaded on your MPX. You do not have to reload unless you are upgrading, backing up, or reloading due to file corruption.

There are different methods for loading software into your switch. The method you use depends on your hardware configuration and the condition of the switch. These methods are:

- FTP Server - The Omni Switch/Router has a built-in FTP server. If you have FTP client software, you can FTP to the switch and load new software.
- FTP Client - The Omni Switch/Router can also be an FTP client. You can use this by connecting a terminal to the switch and using the set of FTP commands in the User Interface. You can also do this through a telnet session.
- ZMODEM - You can load software directly through the serial port with any terminal emulator that supports the ZMODEM protocol. You can do this using the file commands in the User Interface or through the boot line prompt. Note that a ZMODEM transfer of larger files can take several minutes to complete.

Do Not Mix Software Versions

When loading software, ensure that the versions of software for all the modules are from the same release. Mixing earlier versions of software with current versions can cause the switch to reset or hang.

File Transfer/Corruption Problems

If at anytime, a file transfer fails, a fragment of the file may be left on your system. This remaining file is corrupted. You should delete the file fragment and reload the file before continuing. If the MPX image file (**mpx.img**) is corrupted, you will receive a message during the boot sequence requesting you to delete the file. You should delete the file and reload it using ZMODEM through the boot line prompt. See *Using ZMODEM With the Boot Line Prompt* on page 5-5 for information on loading through the boot prompt.

Using FTP Server

The Omni Switch/Router is an FTP server. Using any compatible FTP client software you can load software to and from the switch. Consult the manual that came with your FTP client software package. The following are general instructions on how to FTP to the switch.

1. You will need to configure the IP address in the switch. If you have not done this, refer to the *Getting Started Guide* that came with your switch.
2. Use your FTP client software just as you would with any FTP server. When you connect to the switch you will be able to see the files contained in the flash directory. It is the only directory in the switch.
3. Note that because of the organization of files in the switch, any time a file is deleted, the flash memory is compacted. Depending on the number of files in the switch and where they are located in memory, this compaction can take anywhere from a few seconds to a couple of minutes.
4. When you transfer a file to the switch and one of the same name exists, the old file must first be deleted. You first delete the old file, then the compaction takes place, and then you can transfer the new file. When you begin your transfer, you may not see anything happening for approximately 2 minutes due the file compaction procedure. After compaction, the file will be transferred.

Using FTP Client

The User Interface contains several FTP commands. Using these commands is similar to using FTP on a UNIX system. Follow the steps below to start the FTP Client.

1. Log on to the switch and type **ftp**. For instructions on logging into the switch see the *Getting Started Guide* that came with your switch.
2. The system will prompt for a host. It saves the last host name or IP address used. If it's the one you want, press **<Enter>** or enter the new address.
3. The system will prompt for a user name. It saves the last user name. If it's the one you want, press **<Enter>** or enter the new user name.
4. The system will prompt for a password. Enter your password.
5. After logging onto the system you will receive the **ftp>** prompt. Type a question mark (?) to review the ftp commands. These commands are described in Chapter 7, "Managing Files." The following screen displays:

Supported commands:

ascii	binary	bye	cd	delete
dir	get	help	hash	ls
put	pwd	quit	remotehelp	user
lpwd				

ascii	Set transfer type to ASCII (7-bit).
binary	Set transfer type to binary (8-bit).
bye	Close gracefully.
cd	Change to a new directory on the remote machine.
delete	Delete a file on the remote machine.
dir	Obtain a long listing on the remote machine.
get	Retrieve a file from the remote machine.
hash	Print the hash symbol (#) for every block of data transferred. This command toggles hash enabling and disabling.
ls	Summary listing of the current directory on the remote host.
put	Send a file to the remote machine.
pwd	Display the current (present) working directory on the remote host.
quit	Close gracefully.
remotehelp	List the commands that the remote FTP server supports.
user	Send new user information.
lpwd	Display the current (present) working directory on the local host.
?	Summarize this list.

If you lose communications while running ftp, you may receive the following message:

Waiting for reply (Hit ^C to abort).....

6. You may press **<ctrl-c>** to abort the ftp or wait until the communication failure is resolved and the ftp transfer will continue. Note that Sun OS systems lose echo when you use the **ctrl-c** key combination.

Using ZMODEM

Normally you use FTP to transfer files to and from the switch. It is faster than using the serial port. A ZMODEM transfer can take several minutes. There are generally two situations which would require you to use the serial port to load software:

- You do not have access to an FTP client or server program. If the switch is up and running, you can use the File commands to load software.
- You have deleted the image software files in the switch. If you are in this situation, the only way to load software is using ZMODEM with the boot line prompt.

To use ZMODEM, you must have a terminal emulator that supports the ZMODEM protocol. There are many packages on the market and they operate differently; therefore instructions on how to use them are beyond the scope of this document. Consult the user manual which came with your terminal emulation software.

Before doing a serial port transfer, you should set the baud rate to the highest possible (however, it is not recommended that you run it at 38.4 Kbps). Running at 19200 is twice as fast as 9600. To set the baud rate, use the **ser** command. For more information on the **ser** command, see Chapter 6, “Configuring Management Processor Modules.”

Note

If a file you are transferring already exists in the switch's flash memory, you must remove the file before transferring the new file via ZMODEM.

Using ZMODEM with the load Command

If your switch is up and running, log on to the switch. Type **ls** to list the files in flash memory. If the file you are going to transfer exists, you must delete it first with the **rm** command.

From the File menu, type **?** to list the file commands. The command you use to start the ZMODEM process is **load**. The **load** command does not support speeds greater than 19,200 bauds.

```
/File % load
```

```
The Console (DCE) port is currently running at 19200 baud  
Type 'y' to start ZMODEM download, 'q' to quit (y) : y
```

```
Upload directory: /flash  
ZMODEM ready to receive file, please start upload (or send 5 CTRL-X's to abort).
```

```
**B0100000023be50
```

Activate the ZMODEM transfer according to the instructions that came with your terminal emulation software. When the transfer is completed use **ls** again to list the file or files you have loaded.

Using ZMODEM With the Boot Line Prompt

If you encounter the situation where you have deleted some or all of the files in your switch, you may need to load files through the boot line prompt. This load procedure is done before the switch has booted. If there is no software available in the switch, then it cannot boot until you reload the software.

Using ZMODEM with the boot prompt is similar to using it with the load command. This section covers only specific step-by-step instructions to load a file using ZMODEM at the [boot]: prompt. Before doing this you may want to familiarize yourself with the boot line commands. See Appendix A, “Boot Line Prompt,” for more information.

◆ Important Note ◆

Loading software through the boot prompt should only be done when the switch is off line and not being used for normal network traffic.

Set Up the Correct Baud Rate

1. Connect a terminal to the console port. The terminal must be set to the last values set in the switch before it was powered down. For example if you were running at 19200,8,n,1, you must set your terminal to these values.

Note

If you have deleted or lost your configuration file (**mpm.cfg**), the console port values will revert back to the factory settings which are 9600,8,n,1.

If you are not sure what baud rate your switch is running, try the last known value. If your terminal displays garbage, keep changing the baud rate on your terminal emulator until you see normal ASCII characters.

2. If the switch is on, switch it off for a few seconds, then back on. You should see the boot start up on your screen. You will see the following:

```
System Boot
Press any key to stop auto-boot...
2
```

The number 2 shown above counts down to 0. To stop the boot, you must press a key before the number counts down to 0. If you miss this, simply turn the switch off for a few seconds, then back on to restart the process. Note that if there is no software in the switch it will not be able to boot and will eventually end up at the [boot] prompt anyway.

The [boot] Prompt

The [boot] prompt has its own set of commands that are built into the switch. You do not need to have files or software loaded to use this set of commands. You can perform many of the functions that the MPX software does; however, the purpose of these commands are to reload software in order to get the switch up and running.

To see a list of the boot commands, type **?** at the [boot]: prompt. The following screen displays:

```
[Boot]: ?
?          - print this list
Q          - boot (load and go)
p          - print boot params
c          - change boot params
l          - load boot file
g adrs    - go to adrs
d adrs [,n] - display memory
m adrs    - modify memory
f adrs, nbytes, value - fill memory
t adrs, adrs, nbytes - copy memory
e          - print fatal exception
n netif   - print network interface device address
L          - list ffs files
P          - Purge system: remove ALL ffs files
R file [files] - remove ffs file(s)
S          - save boot configuration
V          - display bootstrap version
$dev(0,procnum)host:/file h=# e=# b=# g=# u=usr [pwr=passwd] f=#
           tn=targetname s=script o=other

Boot flags:
0x02 - load local system symbols
0x04 - don't autoboot
0x08 - quick autoboot (no countdown)
0x20 - disable login security
0x40 - use bootp to get boot parameters
0x80 - use tftp to get boot image
0x100 - use proxy arp
0x1000 - factory reset

available boot devices: sl ffs zm
[Boot:]
```

Note that these commands are all case sensitive.

Type **L** to lists the files in flash memory. This will help you determine what files may be missing. If the file you are going to transfer exists, you must delete it first with the **R** command.

You may want to purge memory and reload all the files. To purge the flash memory, type in the **P** command.

Warning

After using the **P** command, there will be no files in flash and you will have to reload them all with ZMODEM.

Starting a ZMODEM Transfer at the [boot] Prompt

1. Type **c** to change boot parameters. You will be changing the boot device to **zm**. This will tell the system to load files from a ZMODEM connection instead of flash memory.

```
[Boot]: c
'.' = clear field; '-' = go to previous field; ^D = quit
```

```
Boot device : zm
```

2. Type **zm** at this prompt. You will be prompted for more parameters. Just hit **<Enter>** to accept the defaults.

```
Boot file : /flash/mpx.img
Local SLIP adr :
Startup script: /flash/mpx.cmd
Console params : 9600,n8lc
Modem params : 9600,n8l
Boot flags :0xb
Other: dvip:no_name, 198.206.183.253, 255.255.255.0, 198.206.183.255;
```

```
[Boot]:
```

3. When you complete the command, the system will return to the **[Boot]:** prompt. Type in the “at” command (**@**) to load the boot parameters.

```
[Boot]: @
```

```
Boot device : zm
Boot file : /flash/mpx.img
Startup script: /flash/mpx.cmd
Console params : 9600,n8lc
Modem params : 9600,n8l
Boot flags :0xb
Other: dvip:no_name, 198.206.183.253, 255.255.255.0, 198.206.183.255;
```

```
Attaching network interface lo0... done.
Disk load or Boot load (D/B/Q)? -> d
```

4. At the **Disk load or Boot load {D/B/Q}? ->** prompt, type in **d** to tell the system to load from a disk. The system is prepared to accept a ZMODEM transfer, and displays the following:

```
Upload directory: /flash
ZMODEM ready to receive file, please start upload (or send 5 CTRL-X's to abort).
```

```
**B0100000023be50
```

5. Activate the ZMODEM transfer according to the instructions that came with your terminal emulation software.
6. When the transfer is completed use **L** (case sensitive) to list the files you have loaded.
7. Repeat this procedure for every file that you want to load.

6 Configuring Management Processor Modules

The management processor module (MPX on the Omni Switch/Router) coordinates control of the Omni Switch/Router by providing access to the User Interface (UI) software, maintaining user configuration information, downloading switching module software, managing basic bridge functions, maintaining basic routing functions, and managing the SNMP management agent. Switching modules are dependent on the MPX for downloading software and for receiving initialization and configuration information. In addition, the Network Management System (NMS) depends on the MPX to send and receive SNMP messages for managing the switch.

◆ Important Note ◆

All of the UI commands described in this chapter also work with the Omni Switch/Router MPX.

The Omni Switch/Router also support two MPXs with one acting as the primary and with one acting as the secondary. If the primary MPX fails, the secondary MPX can take over automatically. Operating with redundant MPXs can also help avoid network downtime.

◆ Note ◆

When you have two MPXs in one chassis, they must be installed in slots 1 and 2, and only one will be active.

The primary MPX executes all the commands and, when needed, sends requests to the secondary MPX. The secondary MPX continuously monitors the primary MPX. For more information on MPXs, see Chapter 2, “The Omni Switch/Router MPX.”

The UI provides commands to configure the serial port, to configure the Ethernet management port, and a set of commands to monitor and configure primary and secondary MPXs. These commands are described in the pages that follow.

Changing Serial Port Communication Parameters

The serial communications parameters for the two MPX ports are set by default to the following:

- 9600 bits per second (bps)
- 8 data bits
- 1 stop bit
- no parity

To change the serial port configuration parameters, follow the steps below:

1. Log into the switch. For instructions on logging in, see your *Getting Started Guide*.
2. At the system prompt, type **ser**.
3. You will see the following message:

Port to configure? {(C)onsole,(M)odem} (Console) :

Press **C** if you want to configure the console port (female, DCE) parameters, or type **M** to configure the modem port (male, DTE) parameters. The default is the Console Port (**C**).

4. The current port values are shown, followed by a prompt to change the speed value.

Current Console (DCE) configuration:

**9600 bps, 8 data bits, None parity, 1 stop bit, running Console (shell)
Speed (9600):**

Enter the speed (in bits per second) at which you want the port to operate, or simply press **<Enter>** to accept the default in parentheses. Valid values are 1200, 9600, 19200, and 38400 bps.

5. The following prompt displays:

Data size {7/8} bits (8) :

Enter the data size in bits (7 or 8). The default is 8. Press **<Enter>** to accept the default in parentheses.

6. The following prompt displays:

Parity { (N)one/(E)ven/(O)dd } (None) :

Enter the parity (none, even, odd) and press **<Enter>**. The default is None.

7. The following prompt displays:

Stop bits {0/1/2} (1):

Enter the number of stop bits (0, 1, or 2) and press **<Enter>**. The default is 1.

8. The following prompt displays:

Mode {(D)own,(C)onsole,(A)uxConsole,(S)LIP} (C) :

Enter the port mode and press **<Enter>**. This option defaults to console for a console connection and down for a modem connection. You can also configure the port for SLIP. If you are configuring the modem port, you should plan the mode configuration carefully. See *Configuring the Modem Port* on page 6-3 for further information.

◆ Important Note ◆

You cannot configure the console port as an auxiliary port (**AuxConsole**).

9. The following prompt displays:

Set (and save) these settings {(S)ave/(Q)uit} (Save) :

Enter **save** to accept the parameters you entered and exit, or enter **quit** to exit this command without saving your changes.

Changing Port Speed When Communication With The Switch Lost

When you cannot communicate with the switch, there is an alternative method you can use to toggle through the various serial port speed options. The port defaults to 9600 bps. But if you send a Break signal (by pressing the **BREAK** key), the port speed will change to the next higher speed. When it reaches the highest speed (38400 bps), it toggles back to the lowest speed (1200 bps). You cycle through the port speeds in the following order: 9600–19200—38400–1200.

◆ Note ◆

On the MPX you must remove the default baud rate shunt (E1), which fixes the baud rate at 9600 bps, before you can change the baud rate. This shunt is located near the front end of the circuit board, just to the right of the Ethernet management port.

Configuring the Modem Port

If you plan to use the modem port as your main connection to User Interface software, then you need to make sure its mode and jumper settings are configured correctly.

Modem Port Mode

The **ser** command allows you to configure an active modem port to SLIP, console, or auxiliary console mode. When using a modem, it is recommended that you configure the two ports as follows:

```
modem port mode=SLIP
console port mode=console
```

This configuration allows you to use the modem port to access User Interface software through a SLIP connection. The console port is used as an optional way to access software.

◆ Please Note ◆

You need Release 3.2 or above to use the modem and console ports simultaneously.

Another valid configuration is as follows:

```
modem port mode=console
console port mode=down
```

This configuration does not allow you to use the console port as an optional access method since it is configured down. Using a cross-over cable, you could access the modem port through an attached PC. If you could not use the modem port for some reason, you would have to reboot the switch to get back, or—if the cable connection were the problem—use a cross-over cable to connect through a PC.

A third valid configuration that keeps both ports active is:

```
modem port mode=console
console port mode=SLIP
```

This configuration allow you to use the modem port regularly and use a SLIP connection to access switch software through the console port.

A fourth valid configuration that keeps both ports active is:

```
modem port mode=auxiliary
console port mode=console
```

This configuration allow you to use the console and modem ports simultaneously to access switch software.

Configuring SLIP

To configure SLIP, enter the **slipc** command. If you enter the command and SLIP is not running on any ports, the system displays the following message:

Current SLIP configuration

SLIP not running on any ports, do you want to configure it?

Yes, No {Y/N} (Y) :

Enter **y** to display current information. Enter **n** to skip the display. To configure the required SLIP parameters, complete the following steps:

1. Type **slipc** at the prompt and press **<Return>**.
2. Enter a valid IP address.
3. Enter a valid remote IP address.

You can use the **ping** command to validate the connection's integrity.

Configuring the Ethernet Management Port

To configure the Ethernet management port, you use the **ethernetc** command. To use this command, enter

```
ethernetc
```

at the system prompt. A screen similar to the following will be displayed.

Ethernet Port Configuration

```
1) Port Admin status UP : Yes
2) IP Address           : 198.206.184.175
3) Subnet Mask         : 255.255.255.0
4) Bcast Address       : 198.206.184.255
5) Gateway Address     : 198.206.184.254
6) Remote Host Address : UNSET
7) RIP Mode            : Inactive
```

```
Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :
```

The question mark option (?) and the **Help** option provide reference and instructional information on using this command. The **Redraw** option refreshes the screen.

You make changes by entering the line number for the option you want to change, an equal sign (=), and then the value for the new parameter. When you are done entering all new values, type **save** at the colon prompt (:), and all new parameters will be saved. If you do not want to save the changes enter **quit** or **Ctrl-D**.

◆ Important Note ◆

On some revisions of the MPX, you *must* configure the Ethernet management port with the boot prompt before you can use the **ethernetc** command. See Appendix A, “The Boot Prompt,” for more information on configuring the Ethernet management port with the boot prompt.

The configurable options displayed by the **ethernetc** command are described below.

1) Port Admin status UP

Enter **1=Yes** (the default) to enable the Ethernet management port or **1=No** to disable it.

2) IP Address

Enter an IP address for the Ethernet management port in dotted decimal or hexadecimal notation (the default is **192.168.11.1**). For example, to change the Ethernet management port's IP address to **198.206.184.170**, enter

2=198.206.184.170

at the prompt.

◆ Note ◆

This IP address *must* not be on the same subnet as any other IP router on the switch.

3) Subnet Mask

Enter an IP subnet mask in dotted decimal or hexadecimal notation (the default is **255.255.255.0**). If no mask is provided, the switch will try to determine the mask using Internet Control Message Protocol (ICMP) requests. For example, to change the subnet mask to **255.255.255.254**, enter

3=255.255.255.254

at the prompt.

4) Bcast Address

The default broadcast address is automatically derived from the default VLAN IP address class (the default is 192.255.255.255). You can enter a new address in dotted decimal or hexadecimal notation. For example, to change the broadcast address to **198.206.184.255**, enter

4=198.206.184.255

at the prompt.

5) Gateway Address

You can enter an IP address for the first hop router to a remote host (if the host is on a different IP net) in dotted decimal or hexadecimal notation. The default is 192.168.1.1. For example, to change this address to **198.206.184.170**, enter

5=198.206.184.170

at the prompt.

6) Remote Host Address

You can enter an IP address for a a remote host (if the host is on a different IP net) in dotted decimal or hexadecimal notation. The default is 192.168.1.1. For example, to change this address to **198.206.184.170**, enter

5=198.206.184.170

at the prompt.

7) RIP Mode

This parameter is an informational field, which shows that the RIP mode is inactive. You *cannot* modify this parameter.

Ethernet Management Ports and Redundant Management Processor Modules

If redundant MPXs both have Ethernet management ports (EMPs), both EMPs in the switch will have the same IP address if automatic file synchronization is enabled. If both EMPs are plugged into the same subnet, the UI will show that there are duplicate IP addresses on the network.

To get around this duplicate IP address problem, you must disable automatic file synchronization and then you must configure different IP addresses for the two EMPs. To do this, perform the following steps:

1. On the primary management module, enter

syncctl

at the system prompt. (See *Setting Automatic Config Synchronization* on page 6-15 for more information on the **syncctl** command.)

2. If automatic file synchronization is already disabled, simply press **<Enter>**. If it is enabled, enter **disable** at the prompt.

3. Enter

ethernetc

at the prompt. (See *Configuring the Ethernet Management Port* on page 6-5 for more information on the **ethernetc** command.)

4. Enter **2=** followed by the IP address for the EMP on the primary management module.

5. Enter

save

at the prompt to save the IP address.

6. Enter

renounce

at the prompt to make the primary management module the secondary module and the secondary module primary.

7. Log into the now primary management module.

8. On the now primary management module, enter

syncctl

at the system prompt.

9. If automatic file synchronization is already disabled, simply press **<Enter>**. If it is enabled, enter **disable** at the prompt.

10. Enter

ethernetc

at the prompt.

11. Enter **2=** followed by the IP address for the EMP on the management module.

12. Enter

save

at the prompt to save the IP address.

13. Enter

renounce

at the prompt to make the management module that was originally the primary one primary again.

The MPM Command/Menu

The **mpm** command has two functions: displaying the MPX redundancy configuration and entering the **mpm** menu. Displaying the MPX redundancy is described below and the **mpm** menu is described in *MPM Menu Commands* on page 6-9.

Displaying MPX Redundancy

You can display the number of MPXs, their location in the switch, and the MPX redundancy configuration of the switch by entering

mpm

at the system prompt. The following is a typical example of the message that displays when you enter **mpm** for a switch without a redundant MPX.

Currently this slot 1 holds the Primary MPM; there is no secondary MPM.

The following is a typical example of the message that displays when you enter **mpm** for a switch with redundant MPXs on the primary MPX.

Currently this slot 1 holds the Primary MPM and slot 2 holds the secondary.

The following is a typical example of the message that displays when you enter **mpm** for a switch with redundant MPXs on the secondary MPX.

Currently slot 1 holds the Primary MPM; this slot 2, holds the secondary MPM.

MPM Menu Commands

The **mpm** command also takes you to the **mpm** menu which contains the commands needed to configure single and redundant MPXs. With a serial or modem connection, you can communicate with either the primary or secondary MPX by connecting to the respective RS232 connectors. With a telnet connection, however, you can only communicate with the primary MPX.

Type a **?** to list the **mpm** commands. One set of commands will be displayed if you are connected to the primary MPX and another command will be displayed if you are connected to the secondary MPX. If you are connected to the primary MPX, you will see the following.

<u>Command</u>	<u>Redundancy Menu</u>
sls	List the contents of the Secondary /flash and /simm directories
mpmstore	Store file to Secondary /flash or /simm directory
mpmreplace	Replace file on Secondary /flash or /simm directory
mpmload	Load file from Secondary MPM
mpmrm	Remove file from Secondary MPM
renounce	Give up control to Secondary
nisuf	Set load suffix for NI image files
syncctl	Enable/Disable synchronization of configuration data
configsync	Synchronize configuration data
imgsync	Synchronize Image (Executable) files
secreset	Reset Secondary MPM
swap	Change swap status of chassis

All of the **mpm** menu commands, except for the **nisuf** and **swap** commands, function only if you have redundant MPXs. If you are connected to the secondary MPX, type a **?** to list the **mpm** commands shown below.

<u>Command</u>	<u>Redundancy Menu</u>
mpmget	Get file from Primary MPM
takeover	Become Primary

All of the **mpm** commands are described in the sections that follow.

Using MPM Commands with Software Release 3.2 and Later

In Release 3.2 and later, the commands in the **mpm** menu support the use of more than one flash directory. Since more than one flash directory can exist, you *must* indicate which flash directory you want to use when you access a secondary MPX from a primary MPX and when you access a primary MPX from a secondary MPX. All of these commands begin with the prefix **mpm** and are listed below.

- mpmstore**
- mpmreplace**
- mpmload**
- mpmrm**
- mpmget**

To indicate which flash directory you want to use, enter a slash (*/*), the name of the directory, and another slash (*/*) before the file name in all commands that begin with the prefix **mpm**. For example, to transfer the **asm.img** file from the **/simm** directory on the secondary MPX to the primary MPX when you have logged into the secondary MPX, enter

```
mpmget /simm/asm.img
```

at the system prompt.

◆ Important Note ◆

In the current release, you *must* indicate the name of the flash directory in commands that begin with the prefix **mpm** even if you have just one flash directory on both MPXs.

Listing the Secondary MPX Files

The **sls** command lists the files in the secondary MPX module. This is similar to the **ls** command; however, it lists files in the secondary MPX. To list files in the secondary MPX, enter

```
sls
```

at the system prompt. The following is a typical example.

```

/flash/esm.img          27204    7/14/99  11:39
/flash/mesm.img        27561    7/14/99  11:39
/flash/mpm.img         1790889  7/14/99  11:39
/flash/rav.img         83588    7/14/9   11:39
/flash/mpm.cnf         32768    1/ 1/70  00:00
/flash/mpm.log         18072    7/30/99  13:51
/flash/mpm.cfg         32768    7/30/99  14:40
/flash/mpm.cmd         32       1/ 1/70  00:00
/flash/gated.img       547041   8/27/9   16:01

/flash has          1071449 bytes free.
/simm Not present.
```

The **sls** command lists every file in the secondary MPX's flash memory followed by its size (in bytes), creation date, and creation time. The three-letter file name suffix indicates the type of file which includes configuration (**cnf** and **cfg**), command (**cmd**), and image (**img**). The image file suffix can be changed for both the primary and secondary MPXs with the **nisuf** command, which is described in *Setting the Load Suffix* on page 6-14.

Transferring a File to the Secondary MPX

The **mpmstore** command transfers a file in the flash memory of the primary MPX to the flash memory of the secondary MPX. To use this command, enter **mpmstore**, followed by a space, a slash (*/*), the name of the flash directory, another slash (*/*), and the name of the file you want to transfer.

For example, to transfer the file **mpm.log** from the **/flash** directory on the primary MPX to the secondary MPX, for example, you would enter

```
mpmstore /flash/mpm.log
```

at the system prompt. The following will be displayed.

```
Transferring...
```

If the file already exists on the target MPX, something similar to the following message will be displayed.

```
File mpm.log exists on slot 2
```

Use the **mpmreplace** command, which is described in *Replacing a File on the Secondary MPX* on page 6-12, to replace a file that already exists.

Replacing a File on the Secondary MPX

The **mpmreplace** command replaces a file on the secondary MPX. It works like a combination of **mpmrm**, which is described in *Removing a File from the Secondary MPX* on page 6-13, and **mpmstore**, which is described in *Transferring a File to the Secondary MPX* on page 6-11. To use this command, enter **mpmreplace**, followed by a space, a slash (/), the name of the flash directory, another slash (/), and the name of the file you want to replace.

For example, to replace the file **mpm.log** on the secondary MPX with the file **mpm.log** from the **/flash** directory on the primary MPX, for example, you would enter

```
mpmreplace /flash/mpm.log
```

at the system prompt. The following will be displayed.

```
Deleting.  
Transferring
```

If the file already exists on the target MPX and it is identical to the one you are transferring, something similar to the following message.

```
File mpm.log is identical on Primary and Secondary 2
```

If the files are identical, the **mpmreplace** command will terminate and the file will not be replaced.

Loading a File from the Secondary MPX

The **mpmload** command loads a file from the flash memory of the secondary MPX into the flash memory of the primary MPX. To use this command, enter **mpmload**, followed by a space, a slash (/), the name of the flash directory, another slash (/), and the name of the file you want to load.

For example, to load the file **mpm.log** from the **/flash** directory on the secondary MPX into the primary MPX, for example, you would enter

```
mpmload /flash/mpm.log
```

at the system prompt.

Removing a File from the Secondary MPX

The **mpmrm** command removes (deletes) a file from the flash memory of the secondary MPX. To use this command, enter **mpmrm**, followed by a space, a slash (/), the name of the flash directory, another slash (/), and the name of the file you want to remove.

◆ Note ◆

You can only remove a single file with the **mpmrm** command. You *cannot* use wildcards to remove multiple files.

For example, to remove the file **mpm.log** from the **/flash** directory on the secondary MPX in slot 2, for example, you would enter

```
mpmrm /flash/mpm.log
```

at the system prompt. Something similar to the following will be displayed.

```
Checking for /flash/mpm.log on slot 2
```

After a brief moment, the file will be deleted from the secondary MPX and something similar to the following will be displayed.

```
Deleting /flash/mpm.log on slot 2 . Done.
```

◆ Warning ◆

You *cannot* recover a file once it has been deleted with the **mpmrm** command.

Giving Up Control to the Secondary MPX

The **renounce** command tells the primary MPX to give up control and become the secondary MPX. It does this by issuing a request to the secondary MPX to take control. You *must* be logged into the primary MPX to use this command. If you are logged into the secondary MPX, use the **takeover** command, which is described in *Gaining Control from the Primary MPX* on page 6-18.

◆ Warning ◆

The **renounce** command should only be used during network down times since it could cause network interruptions.

To transfer control from primary MPX to the secondary MPX, enter

```
renounce
```

at the system prompt. The following prompt will display.

```
Confirm? (n):
```

Press **y** to transfer control to the secondary MPX or press **n** to cancel the command (the default is **n**). If you enter **y**, the switch will reset after displaying the following message.

```
System going down immediately...
```

The switch will reboot and the original secondary MPX will be the primary once the switch comes back up.

Setting the Load Suffix

The **nisuf** command sets the load suffix for the switch's executable image files. (The factory default suffix is **img**.)

◆ Warning ◆

The **nisuf** command should only be used when it is necessary to have two versions of the software on the switch at the same time and the user is directly connected to the console for reboot.

You can change it by typing the **nisuf** command followed by the new suffix. For example, to change the load suffix from **img** to **bin**, enter

```
nisuf bin
```

at the system prompt. The following message will then be displayed.

```
Changing load suffix from img to bin
```

You should create or load new image files with the new suffix as soon as possible because the switch will not recognize the files with the old suffix as image files. See Chapter 5, "Installing Switch Software," and Chapter 7, "Managing Files," for information on loading and creating files.

Setting Automatic Config Synchronization

The **syncctl** command sets the automatic configuration synchronization to Enabled or Disabled. If it is Enabled, then the MPX primary/secondary pair will continue to maintain synchronization automatically. This means that when the configuration file (**mpm.cfg**) is updated in the primary MPX, it will automatically be updated in the secondary MPX, keeping the two MPXs in sync.

Enabling Automatic Config Synchronization

To enable synchronization between the primary and secondary MPXs, enter

```
syncctl
```

at the system prompt. The following prompt will then be displayed if synchronization is not enabled.

Desired state (enable):

Press **<Enter>** to enable synchronization or enter **disable** to cancel. If you enabled synchronization, the following will be displayed.

Configuration synchronization is now Enabled

Note that automatic configuration synchronization is disabled unless all image (**img**) and Programmable Gate Array (PGA) files in the switch are synchronized first. See *Synchronizing Image Files* on page 6-16 for information on the **imgsync** command, which synchronizes image and PGA files.

The interval between updates is 5 minutes. The primary MPX will copy any changes to the secondary MPX after 5 minutes have elapsed since the last update.

Disabling Automatic Config Synchronization

To disable synchronization between the primary and secondary MPXs, enter

```
syncctl
```

at the system prompt. The following prompt will then be displayed if synchronization is enabled.

Desired state (disable):

Press **<Enter>** to disable synchronization or enter **enable** to cancel. If you disabled synchronization, the following will be displayed.

Configuration synchronization is now Disabled

If automatic config synchronization is Disabled, the configuration file in the secondary MPX will be unaffected if you change the configuration file in the primary MPX.

Synchronizing Configuration Data

The **configsync** command copies the configuration files (**mpm.cnf** and **mpm.cfg**) in the primary MPX to the secondary MPX. You can run this command whether or not automatic config synchronization is on. For example, to copy the configuration file from the primary MPX to the secondary MPX, you would enter

```
configsync
```

at the system prompt. Something similar to the following will be displayed.

```
Syncing Config file  
Config files are currently synchronized.
```

See *Setting Automatic Config Synchronization* on page 6-15 for information on setting automatic config synchronization.

Synchronizing Image Files

The **imgsync** command copies all of the image (executable) files in the primary MPX to the secondary MPX. When used in conjunction with the **configsync** command, it ensures that the two MPXs are running exactly the same versions of software and are in sync (i.e., have the same configuration). To synchronize all the image files, enter

```
imgsync
```

at the system prompt. When you run **imgsync** you will be asked if you want to synchronize the **cmd** file and/or PGA files if they are found to be different.

◆ **Note** ◆

If any PGA file is being used by a Token Ring module and you choose to sync the cmd file, then the PGA file that is in use will be synced even if you do not choose to synchronize PGA files.

Something similar to the following prompt will be displayed.

```
Sync cmd file (y) :
```

Press **y** to sync the **cmd** file or press **n** to skip this file (the default is **y**). If you have any PGA files, you will be asked if you want to sync those files. In addition, if the secondary MPX has any additional image, then the following prompt will be displayed.

```
Remove Additional images from Secondary (n) :
```

Press **y** to remove any extra image on the secondary MPX or press **n** to keep these files (the default is **n**). After you answer all the prompts, something similar to the following will be displayed.

```
8 files to be synchronized  
1 file to be synchronized  
Syncing  
Deleting /flash/mpx.cmd.....  
Replacing /flash/mpx.cmd.....
```

Loading a File From the Primary MPX

The **mpmget** command loads a file from the primary MPX and copies it into the secondary MPX. This command is only available and can only be run from a secondary MPX. To use this command, enter **mpmget**, followed by a space, a slash (/), the name of the flash directory, another slash (/), and the name of the file you want to transfer.

For example, to load the file **mpm.log** from the **/flash** directory on the primary MPX to the secondary MPX you would enter

```
mpmget /flash/mpm.log
```

at the system prompt. After a brief moment, the file will be transferred into the secondary MPX. The following would then be displayed.

```
Transferring .. Complete
```

Gaining Control from the Primary MPX

The **takeover** command tells the secondary MPX to take control and become the primary MPX. It does this by issuing a request to the primary MPX to relinquish control. You *must* be logged into the secondary MPX to use this command. If you are logged into the primary MPX, use the **renounce** command, which is described in *Giving Up Control to the Secondary MPX* on page 6-14.

◆ Warning ◆

The **takeover** command should only be used during network down times since it could cause network interruptions.

To transfer control from primary MPX to the secondary MPX, enter

takeover

at the system prompt. The following prompt will display.

Confirm? (n):

Press **y** to transfer control to the secondary MPX or press **n** to cancel the command (the default is **n**). If you enter **y**, the switch will reset after displaying the messages similar to the following.

System going down immediately...

**Please standby, chassis configuration changing (Hit ^C to abort).....Taking over
as Primary**

... Alcatel SNMP Agent Operational.

The switch will reboot and the original secondary MPX will be the primary once the switch comes back up.

Resetting a Secondary MPX

The **secreset** command initiates a soft reset on the secondary MPX. Conceptually, resetting a secondary MPX with this command is similar to switching off power to the module; the MPX will be in the same state after a reset as it is after a power on.

To reset a secondary MPX, enter

```
secreset
```

at the system prompt. Messages similar to the following will display:

```
Module 1 changed while Swap OFF
```

```
Syncing configuration data with secondary 1 .. complete
```

◆ **Note** ◆

To reset a switching module, use the **reset** command, which is described in Chapter 36, “Running Hardware Diagnostics.”

Displaying and Setting the Swap State

The **swap** command displays or alters the swap state of the chassis. The swap state must be on in order to hot swap modules. If not, the system may halt or restart. While the swap state is on, performance may decrease. Therefore, the swap state should only be turned on when you want to hot swap modules. See Chapter 3, “Omni Switch/Router Switching Modules,” for instructions on hot swapping a switching module.

Displaying the Swap State

To display the current swap state of the chassis, enter

```
swap
```

at the system prompt. If the swap mode is **OFF** (the default for the switch), something similar to the following will be displayed.

```
Swap is OFF, timeout is 5 minutes  
usage swap { ON [ minutes ] | OFF [ minutes ] }
```

If the swap mode is **ON**, something similar to the following will be displayed.

```
Swap is ON, expires in 4 minutes  
usage swap { ON [ minutes ] | OFF [ minutes ] }
```

The swap mode *must* be enabled (**ON**) to hot swap a switching module. If not, the system may halt or restart. See the subsection below for instructions on enabling the swap mode.

Enabling the Swap Mode

To turn the swap mode **ON**, enter

```
swap on
```

at the system prompt. (The default for swap mode is 5 minutes). Something similar to the following will be displayed.

```
Swap is ON for 5 minutes
```

When you turn the swap state on, you set a timer which determines how long the system will remain in swap state. After the timer expires, the system will automatically turn off the swap state.

If you want to vary the amount of time that the swap mode is enabled, enter **swap on** followed by the number of minutes you want the swap mode enabled. You can set the swap state from 1 to 227,055 minutes. To set the swap mode on for 10 minutes, for example, enter

```
swap on 10
```

at the system prompt. The following will then be displayed.

```
Swap is ON for 10 minutes  
Save minutes value {Y/N}? (N) :
```

Press **y** and then press **<Enter>** to save the new value. If you don't want save, just press **<Enter>** and the default value will not change. You can also turn off the swap immediately as shown in *Disabling the Swap Mode* on page 6-21.

Disabling the Swap Mode

Normally, the swap mode will timeout and no user intervention is required. However, you can manually turn the swap mode off. This function is particularly useful since the performance of the switch can be adversely affected if the swap mode is enabled. To turn the swap mode off immediately, enter

swap off

at the system prompt. The swap mode will be disabled and something similar to the following will be displayed.

Swap is OFF, timeout is 5 minutes

7 Managing Files

Depending on the model type and configuration, an Alcatel switch has anywhere from 8 or 16 MB of usable flash memory. This memory is used to store files, including executable files (used to operate switching modules), configuration files, and switch usage log files. Through the User Interface (UI), you can load, copy, and delete any of these files types. In addition, the UI has commands for displaying, creating, and editing ASCII (text-based) files.

All commands described in this chapter will work with files located in the **/flash** directory on either the primary or secondary MPX. However, these commands work only with the files that reside on the MPX to which you are connected. See Chapter 6, “Configuring Management Processor Modules,” for more information on commands for working with redundant MPXs.

UI commands for file maintenance are grouped into two menus: the File menu and System menu. File menu commands are listed below. For a list of System menu commands, see *System Menu* on page 7-13.

File Menu

The File menu contains commands for loading, listing, copying, and deleting individual switch files. To access the File menu, enter

file

at the UI prompt.

If verbose mode is enabled, the following list of commands will be displayed automatically.

If verbose mode is disabled, press the question mark (?) to display the following list of commands. (For information on enabling verbose mode, refer to the **uic** command description in Chapter 4, “The User Interface.”)

Command	File Menu
load	Download system software using the serial interface
ftp	Download from an FTP server
pwd	Display the current working directory
ls	List the contents of the current working directory (default working directory is /flash)
rm	Remove a file
cp	Copy a file
view	View an ASCII file
edit	Edit buffer locally
imgcl	Remove all image files

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

All commands in the File menu, except for the **load** and **ftp** commands, are described in the following sections. For instructions on using the **ftp** and **load** commands, refer to Chapter 5, “Installing Switch Software.”

◆ Note ◆

If you want to use the **rm**, **cp**, **imgcl**, and the **edit** sub-menu commands, you must be logged in as **admin** or **diag**. See Chapter 4, “The User Interface,” for more information on login accounts.

Displaying the Current Directory

To display the switch's current directory, enter

```
pwd
```

at the system prompt. The working directory will be the **/flash** memory system and the corresponding directory information will be displayed:

```
/flash
```

Configuration and Log File Generation

The **mpm.cnf**, **mpm.cfg**, and **mpm.log** files are generated automatically by the switch and placed in flash memory during the boot process; you do not have to load them.

◆ Important ◆

If you remove the configuration files (**mpm.cnf** and **mpm.cfg**) from your switch, all of your switch's non-default configuration settings will be deleted at the next boot sequence. Use caution when removing configuration files and be sure to create backup copies if you want to safeguard your current configuration.

Changing Directories

You can change the working directory with the **cd** command. For example:

```
cd test
```

at the system prompt. To change the working directory back to **/flash** file system, enter

```
cd flash
```

at the system prompt.

Listing Switch Files

You can use the **ls** command to list the files in the primary MPX's flash memory. To use this command, enter

```
ls
```

at the system prompt. A screen similar to the following will be displayed.

```
mpx.cmd          18      05/30/98  13:04
mpm.log          18072   06/15/98  17:57
mpx.img         1573617 06/18/98  12:16
esx.img          24289   06/18/98  12:18
mpm.cfg           1024    01/01/70  00:00
mpm.cnf          32768   06/18/98  12:27
```

```
1858057 bytes free.
```

The **ls** command lists all the files in the current working directory of the primary MPX's flash memory, followed by its size (in bytes), creation date, and creation time. The three-letter file extension indicates the type of file. Examples include configuration (**cnf** and **cfg**), command (**cmd**), image (**img**), Programmable Gate Array (**.pga**), etc. The **ls** command also lists the total number of bytes of free memory in flash memory.

◆ Note ◆

If you are connected to the primary MPX and you want to display the files in a secondary MPX, use the **sls** command, which is further detailed in Chapter 6, "Configuring Management Processor Modules."

Deleting Switch Files

You can use the **rm** command to delete files in the primary MPX's flash memory. To use this command, enter **rm**, followed by the name of the file you want to delete. For example, to delete the file **mpm.log**, you would enter

```
rm mpm.log
```

at the UI prompt. The following screen will be displayed:

```
File system compaction in progress...
```

The switch will take a few seconds to delete the file and compact the flash memory.

◆ Note ◆

If you are connected to the primary MPX and you want to remove files from a secondary MPX, use the **mpmrm** command, which is described in Chapter 6, “Configuring Management Processor Modules.”

Deleting Multiple Files

You can remove multiple files either by entering multiple file names in the command line or by using wildcards.

When entering multiple file names, be sure to include a space between each file name you want to delete. For example, to remove both the **mpm.cfg** and **mpm.cnf** files, you would enter the following:

```
rm mpm.cfg mpm.cnf
```

Wildcards let you substitute an asterisk (*) for file name text. You can remove all files with the same extension by entering **rm**, followed by an asterisk (*), a period (.), and the file extension. For example, if you want to delete all the files with the **log** extension, enter

```
*.old
```

at the UI prompt. The following message will be displayed:

```
Remove the following?  
  /flash/mpm.log.old  
  /flash/mpm.old  
Are you sure you want to remove this? (n)
```

Press the **y** key to delete the selected files or press **<Enter>** to cancel. If you press the **y** key, the following will be displayed:

```
...2 files removed
```

The switch will take a few seconds to delete the file and compact the flash memory.

◆ Note ◆

If you want to delete all the image files (i.e., files with the **img** extension), you can use the **imgcl** command, which is described in *Deleting All Image Files* on page 7-5.

Deleting All Image Files

You can use the **imgcl** command to delete all executable (image) files. The files deleted by the **imgcl** command include the MPX boot file (**mpx.img**), and all executable switching module files (the factory default is all files ending with the **.img** extension).

◆ Important ◆

You should only use the **imgcl** command during network down times and when you are connected to the switch through the serial port.

To use this command, enter

```
imgcl
```

at the system prompt. A screen similar to the one shown below will be displayed.

```
Remove the following?  
/flash/esx.img  
/flash/mpx.img  
Are you sure you want to remove them? (n)
```

Press the **y** key to delete all the image files or press **<Enter>** to cancel. If you press the **y** key, the switch will spend several minutes deleting the image files.

◆ Note ◆

If you want to delete *all* files in flash memory, you can use the **newfs** command, which is described in *Creating a New File System* on page 7-15.

After you have deleted all the old image files, you must load new image files using FTP or ZMODEM so the switch can function. See Chapter 5, “Installing Switch Software,” for instructions on using the **ftp** and **load** commands.

Copying System Files

You can use the **cp** command to copy files. This is particularly useful if you want to make backups of important files. To use this command, enter **cp**, followed by the name of the original file you want to copy, and then by the name that you wish to give the duplicate file. For example, to make a duplicate of the file **mpx.cmd** that is to be called **mpx.bak**, enter

```
cp mpx.cmd mpx.bak
```

at the system prompt. The following information will be displayed:

```
/flash/mpx.cmd -> /flash/mpx.bak : 100%
```

Displaying Text Files

You can use the **view** command to display the contents of ASCII (text-based) files. To use this command, enter **view**, followed by the name of the file you want to display. To display the **mpx.cmd** file, for example, enter

```
view mpx.cmd
```

at the system prompt. A screen similar to the one shown below will be displayed.

```
cmDoDump=1  
cmInit
```

Note that if you try to view a file with non-ASCII characters, an error message will be displayed. For example, if you use the **view** command on the file **mpm.cfg**, the following error message will appear:

```
The file mpm.cfg has non-printable characters, can't view
```

◆ Note ◆

You can edit text files with the **edit** sub-menu commands, which are described in *Editing Text Files* on page 7-7.

Editing Text Files

The commands in the Edit sub-menu (also called the Text Buffer or Edit Buffer) are used to create new text files and to modify existing text files. To enter the edit sub-menu, enter

edit

at the system prompt.

If verbose mode is enabled, the following list of commands will be displayed automatically.

If verbose mode is disabled, press the question mark (?) to display the following list of commands. (For information on enabling verbose mode, refer to the **uic** command description in Chapter 4, “The User Interface.”)

Command	Edit Menu
ab	Append line(s) to the buffer
cb	Clear the buffer
db	Delete line from the buffer
eb	Edit a buffer line
ib	Insert buffer line
lb	List contents of the buffer
nb	Name file for buffer
rb	Read file into buffer
wb	Write buffer to file

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The Edit sub-menu commands are outlined in the following sections. You can edit up to 100 lines of text. Each line of text can be up to 97 characters long.

◆ Note ◆

When you edit text files, you will normally use several of the Edit sub-menu commands to produce the results you want. See *Real-World Example 1* on page 7-11 or *Real-World Example 2* on page 7-12 for examples of how to use multiple commands from the Edit sub-menu.

Clearing the Text Buffer

You can use the **cb** command to clear the Edit buffer’s memory so you can create a new text file. To use the **cb** command, enter

cb

at the system prompt.

Loading an ASCII File into the Text Buffer

You can use the **rb** command to load—or *read*—an existing ASCII file in flash memory to the Edit buffer's memory. To use this command, enter **rb**, followed by the file you wish to edit. For example, to edit the **mpx.cmd** file, enter

```
rb mpx.cmd
```

at the system prompt.

◆ Loading Binary Files ◆

You can load a binary file into the Edit buffer but you will not be able to edit it.

Listing the Contents of the Text Buffer

The **lb** command is used to list the contents of the Edit buffer's memory. To use this command, enter

```
lb
```

at the system prompt. If there is something in the buffer, the system will display the contents numbered from the zero. The following display is a typical example:

```
00: cmDoDump=1  
01: cmlnit
```

If there is nothing in the buffer, nothing will be displayed.

Adding Lines of Text to the Text Buffer

You can use the **ab** command to manually add lines of text to the Edit sub-menu. Note that the lines you enter are appended at the end of the buffer. For example, if there are 10 lines of text in the buffer, you will begin entering text at the 11th line. If the buffer is empty, the line of text you enter will be the first line of text in the buffer.

To add text to the buffer, enter

```
ab
```

at the system prompt. A screen similar to the one shown below will be displayed:

```
02 :
```

Enter your text and press the **<Enter>** key to add the text to the buffer. If the buffer is not full, the system will prompt you to enter another line of text. If the buffer is full (i.e., there are 100 lines in the text buffer), the following message will be displayed.

```
Buffer Full!
```

To exit the **ab** command, type a period (.) and press **<Enter>**.

Deleting a Line of Text from the Text Buffer

You can use the **db** command to delete a specific line in the text buffer. To use this command, enter **db**, followed by line number of the line of text you want delete, which is shown by the **lb** command. For example, to delete the third line of text in the text buffer, enter

```
db 3
```

at the system prompt.

Enter the **lb** command again to view the contents of the buffer. Note that the text that appeared at line 3 has been deleted.

Inserting a Line of Text into the Text Buffer

You can use the **ib** command to insert a line of text between two existing lines in the buffer. To use this command, enter **ib**, followed by the number of the line where you want the new text to appear. For example, if you want to add the text, **atm_use_mbus=3**, between lines **00** and **01** in the buffer, enter

```
ib 1
```

at the system prompt. The following screen will be displayed:

```
01:
```

Enter the line of text, **atm_use_mbus=3**.

At the system prompt, enter the **lb** command to view the contents of the buffer. If the original text buffer looked like this,

```
00: cmDoDump=1
01: cmlnit
```

the revised text buffer, with the inserted text, will now appear as follows:

```
00: cmDoDump=1
01: atm_use_mbus=3
02: cmlnit
```

Editing a Line Name of Text in the Text Buffer

You can use the **eb** command to edit an existing line of text in the buffer. To use this command, enter **eb**, followed by the line number of the text you want to edit. For example, if you want to edit the text at line 01, enter

```
eb 1
```

at the system prompt. The following screen will be displayed:

```
01:
```

Enter the text as you want it to appear and press **<Enter>**.

Enter the **lb** command again to list the contents of the text buffer. Note that the buffer now reflects the edited line of text.

Creating a File Name for the Text Buffer

If no file name has been created for the text buffer, the following message is displayed whenever the **lb** command is executed:

Work buffer is unnamed

Use the **nb** command to create a name for the text buffer. To use this command, enter **nb**, followed by the name you wish to give the text buffer. For example, if you want to name the buffer **mpx.cmd**, enter

nb mpx.cmd

at the system prompt. The following screen is displayed, showing the current working directory (**/flash**), followed by the new name for the text buffer (**/mpx.cmd**):

Work buffer name is: /flash/mpx.cmd

Creating a Text File from the Text Buffer

The **wb** command is used to create—or *write*—a text file from the text buffer. To use this command, enter **wb** followed by the name of the output file. For example, if you want to create the file **switch.txt**, enter

wb switch.txt

at the system prompt. The following screen is displayed:

Work buffer name is: /flash/switch.txt

Writing Changes to Existing Files

You can also use the **wb** command to overwrite changes to an existing file. For example, if you want to overwrite changes to the file **mpx.cmd**, enter

wb mpx.cmd

at the system prompt. The following screen is displayed:

/flash/mpx.cmd exists in /flash. Overwrite it? (y)

Press **<Enter>** to create the text file from the text buffer. The computer will take a few seconds as it overwrites the file, and the following information is displayed:

File system compaction in progress...

At the system prompt, enter the **lb** command to view the name of the buffer. Note that the work buffer is now named **/flash/mpx.cmd**.

Real-World Examples

As noted on page 10-7, when you edit text files, you will normally use several of the Edit sub-menu commands to produce the results you want. The following two examples, *Real-World Example 1* and *Real-World Example 2*, are actual multi-command procedures that you may encounter as you work with your switch.

Real-World Example 1

```
cp mpx.cmd mpx.bak
rb mpx.cmd
lb
00: cmDoDump=1
01: cmlnit
nb mpx.cmd
Work buffer name is: /flash/mpx.cmd
ab
02 :
02 : reg_port_rule=1
03 :
No line 3 inserted
lb
00: cmDoDump=1
01: cmlnit
02: reg_port_rule=1
Work buffer name is: /flash/mpx.cmd
wb
/flash/mpx.cmd exists in /flash. Overwrite it? (y)
File system compaction in progress...
view mpx.cmd
cmDoDump=1
cmlnit
reg_port_rule=1
```

Real-World Example 2

```
cp mpx.cmd mpx.bak
rb mpx.cmd
lb
00: cmDoDump=1
01: cmlnit
02: reg_port_rule=1
nb mpx.cmd
Work buffer name is: /flash/mpx.cmd
db 2
lb
00: cmDoDump=1
01: cmlnit
ib 1
01 :
01 : rifStripping=1
lb
00: cmDoDump=1
01: rifStripping=1
02: cmlnit
Work buffer name is: /flash/mpx.cmd
wb
/flash/mpx.cmd exists in /flash. Overwrite it? (y)
File system compaction in progress...
view mpx.cmd
cmDoDump=1
cmlnit
rifStripping=1
```

System Menu

The System menu contains two commands, **fsck** and **newfs**, for checking and deleting all files in the flash memory. To access the System menu, enter

system

at the UI prompt.

If verbose mode is enabled, the following list of commands will be displayed automatically.

If verbose mode is disabled, press the question mark (?) to display the following list of commands. (For information on enabling verbose mode, refer to the **uic** command description in Chapter 4, “The User Interface.”)

Command	System Menu
info	Basic info on this system
dt	Set system date and time
ser	View or configure the DTE or DCE port
mpm	Configure a Management Processor Module
slot	View Slot Table information
sysstat	View system stats related to system, power and environment
taskstat	View task utilization stats
memstat	View memory use statistics
fsck	Perform a file system check on the flash file system
newfs	Erase all file from /flash and create a new file system
syscfg	Configure info related to this system
uic	UI configuration; change - prompt, timeout, more, verbose.
camstat	View CAM info and usage
camcfg	Configure CAM info and usage
hrex	Enter HRE-X management command sub-menu
ver/ter	Enables/disables automatic display of menus on entry (obsolete)
echo/noecho	Enable/disable character echo
chpr	Change the prompt for the system (obsolete, use ‘uic’ command)
logging	View system logs.
health	Set health parameters or view health statistics
cli	Enter command line interface
saveconfig	Dump the cache configuration content to the mpm.cnf file.
cacheconfig	Set the flag to use cache configuration only.

Main Interface	File Security	Summary System	VLAN Services	Networking Help
-----------------------	----------------------	-----------------------	----------------------	------------------------

Checking the Flash File System

The **fsck** command performs a file system check of flash memory, which consists of the flash file system. All image files are stored in flash memory and loaded into system memory when the switch boots up.

The command also provides diagnostic information in the event of file corruption. To perform a file system check of flash memory, enter

fsck

at the system prompt. A screen similar to the following will be displayed:

Your bootroms support Flash File System Version 2 and greater.

Out of 16 file descriptors in use, 0 of these are opened on the /flash device.

Performing a file system check using manual mode. If a file is encountered with a potential problem, you may wish to consider preserving it for technical support analysis...

**Flash file system check in progress...
Checking root file system... OK
Performing file consistency check...
Done.**

There doesn't appear to be a system problem related to the Flash File system or kernel file system data structures. If you are experiencing problems with the flash file system, perhaps try using the "info", "systat", or "memstat" commands. They may indicate some other condition (such as low memory) which could prohibit correct operation of the file system.

If the **fsck** command detects a problem with the flash file system, a message will be displayed indicating the problem, along with any steps needed to resolve it.

Each logical file system must be checked independently.

Creating a New File System

The **newfs** command removes a complete flash file system and all files within it, replacing it with a new empty flash file system. Use this command when you want to reload all files in the file system, or in the unlikely event that the flash file system becomes corrupted.

To create a new file system and re-initialize the flash memory, enter

```
newfs
```

at the system prompt. The following will be displayed.

```
You are about to destroy all files on file system /flash. If you  
are experiencing problems with the flash file system, you might  
want to use the "fsck" command to help determine where problems  
may exist.
```

```
Are you absolutely sure you want to strip the current file  
system and create a new one? (n)
```

Press **<Enter>** to cancel, or enter **y** to create a new file system. If you enter **y**, you will have to load new software into the switch.

◆ Warning ◆

Do not power-down the switch after running the **newfs** command until you reload your image and configuration files. Otherwise, you will have to reload the image files at the boot monitor prompt using the serial interface (e.g., ZMODEM), which can take several minutes. Also, before you execute the **newfs** command, you may also want to preserve your configuration file by saving it to another host.

You can now download new files via FTP or ZMODEM.

8 Switch Security

Commands listed in the Security menu are for configuring system security parameters such as the password and logout time. The menu also provides a command for rebooting the switch. Enter

security

at the prompt to enter the Security menu. Press ? to see the following list of commands:

<u>Command</u>	<u>Security Menu</u>
pw	Set a new password for a login account
reboot	Reboot this system (allowed if the user is "admin")
timeout	Configure Auto Logout Time (obsolete, use "uic" command)
layer2auth	Enable/Disable layer2 user authentication
seclog	Display Secure Access log file entries
secdefine	Define Secure Access filter(s)
secapply	Apply Secure Access filter(s)
useradd	Create a new user for a login account
usermod	Modify a user's privileges
userdel	Remove a user
asacfg	Configure Authenticated Switch Access
userview	View the users in the local user database
auth	Enter the Authentication menu

Main File Summary VLAN Networking
Interface Security System Services Help

The **pw**, **reboot**, **seclog**, **secdefine**, and **secapply** commands are described in this chapter. The **useradd**, **usermod**, **userview** and **userdel** commands are also described in this chapter.

For information about the **layer2auth** and **asacfg** command as well as the authentication (**auth**) submenu, see the *Switched Network Services User Manual*.

Changing Passwords

The switch provides three types of login accounts by default—Administrator, User and Diagnostics. The Administrator login provides full READ/WRITE access to all command families. The login name for the Administrator account is **admin**. The login name for the default User account is **user** and provides READ ONLY access to the switch's command families except for the global family, and NO WRITE privileges. The Diagnostics login has full READ/WRITE access to all command families plus a command for running switching module tests. The login name for Diagnostics is **diag**.

The initial password for all three accounts is **switch**. If you log in as **diag** you can change the passwords for the **diag** and **admin** login accounts. If you log in as **admin**, however, you can only change the password for the **admin** login account. To change the password, complete the following steps. Remember that the User Interface does not echo (display) the password characters.

1. From the prompt, type

pw <account-name>

The **<account-name>** is the user login name (**diag**, **admin**) for which you want to change the password. The following prompt displays:

**Changing password for account:<account-name>
Old password:**

2. Enter the old password and press **<Enter>**. If you enter the old password incorrectly, the following message displays:

Authentication failure

and the command will terminate. You will then need to start over from **Step 1** above.

If you answered the old password correctly, the following prompt displays:

New password:

3. Enter the new password (you are allowed up to 18 characters) and press **<Enter>**. The following prompt displays:

Retype new:

4. Re-enter the new password to confirm it and press **<Enter>**.

◆ Note◆

It is recommended that you change the password from the default for all login accounts.

The passwords are stored encrypted in the **mpm.cnf** file. If you forget your password, you will have to delete the **mpm.cnf** file which will cause the passwords to revert to the default.

◆ Caution ◆

Deleting the **mpm.cnf** file will also remove all of your configuration data and restore everything back to factory settings.

Rebooting the Switch

The **reboot** command should only be executed during network down time and when no data is being transmitted across the network. Also, you should ensure that all configuration information has been saved first. Note that the **reboot** command is only available to the **admin** and the **diag** logins.

◆ **Caution** ◆

Rebooting the switch will disconnect a Telnet connection to the User Interface and will interrupt the network connections on the switching modules.

To reboot the switch from the command line, enter

reboot

at the prompt and press **<Enter>**. The following prompt will display:

Confirm? (n) :

Enter **Y**. The following message displays:

```
Locking file system...locked  
System going down immediately...  
switch[489917b0]: System rebooted by admin
```

The switch will now take at least a minute to start up again. (If you are connected to the User Interface with a serial connection, the console displays start-up related information.) The login message displays when the reboot is complete:

```
Welcome to the Alcatel Omni Switch/Router! (Serial # xxxx)  
login :
```

Secure Switch Access

Secure Switch Access is a filtering program that prevents unauthorized access to the switch by allowing you to define a list of *filters* and *filter points*. For Secure Switch Access, filters are lists of source traffic that are allowed onto the switch. Filter points operate on IP protocols that include FTP, Telnet, SNMP, TFTP, HTTP, and a custom IP protocol. Whenever any of these filter points is enabled, all filters configured for that protocol are applied to incoming traffic using the filter point protocol.

All access violations are logged. If a filtering point is not enabled, it is accessible to all users.

Configuring the Secure Switch Access Filter Database

Use the **secdefine** command to view and configure the database of secure access filters. This database includes information on filter names, source IP addresses, source MAC addresses, and the physical ports receiving data.

The following is a sample **secdefine** display:

```

Secure Access Filter Database

List      (l) :
Create    (c) :
Delete    (d) :
Modify    (m) :
Find      (f) :
Help      (h) :
Quit      (q) :
Enter selection:

```

Select an option by entering the relevant letter at the selection prompt. To exit this menu, enter **q** (quit). Descriptions and sample displays for each of the options are as follows:

List

This is a list of all defined filters. A filter determines what traffic is allowed on the switch. The list includes information on the filter's name, IP Address, MAC Address, and physical port receiving the user's data. The following is a sample display:

Filter Name	Source IP Address	Source MAC Address	Slot #	Port #
Engineering	198.34.56.10	0:23:da:67:97:e4	4	1
Test	ANY	ANY	7	3
Accounting	172.14.25.13	0:32:e4:a3:6f:e4	2	1
HR	198.34.56.15	ANY	ANY	ANY

The value **ANY** displays if a field is left blank when configuring filter information through the **Create (c)** option. The **ANY** value signifies a "don't care" condition. When an inbound packet is checked against a Filter Name to establish authorized access, the **ANY** fields are not checked.

Create

This option allows you to create a new filter in the secure access database. The following is a sample display:

```

Create Filter
-----
Enter Filter Name:

Enter IP Address ( [a.b.c.d] ) :
Enter MAC Address ( [XXYYZZ: AABBC] ) :
Is this MAC in Canonical or Non-Canonical (C or N) [C] :
Enter Slot :
Enter Port :

```

After you have created a filter, the information is automatically saved in the secure access database, and the **secdefine** submenu re-displays. To review your new configuration, simply select the list (I) option. Descriptions of the fields are as follows:

Enter Filter Name: The name of the new filter. The name is required and must be at least one character long and no more than 25 characters.

Enter IP Address ([a.b.c.d]): The allowed IP address. The address must be in the displayed format ([a.b.c.d]). If you enter a value here, the user may access the switch only from this IP address. If you leave this field blank, a value of **ANY** will display in the secure access list, allowing access to the switch from any IP address.

Enter MAC Address (([XXYYZZ: AABBC])): The allowed MAC address. The address must be in the displayed format (([XXYYZZ: AABBC])). If you enter a value here, a user may access the switch only from this source MAC address. If you leave this field blank, a value of **ANY** will display in the secure access list, allowing this user access to the switch from any MAC address.

Is this MAC in Canonical or Noncanonical (C or N) [C] : The format of the specified MAC address. Typically, ethernet MAC addresses are in canonical format while token ring and addresses are in noncanonical format. The default is canonical (C). This parameter is not required.

Enter Slot: The module on the switch receiving data from the specified IP or MAC address. If you leave this field blank, a value of **ANY** will display in the secure access list, allowing data from the specified IP or MAC address to be sent through any module on the switch.

Enter Port: The port on the module receiving data from the specified IP or MAC address. If you enter a value here, you should also specify a slot in the above field. If you leave this field blank, a value of **ANY** will display in the secure access list, allowing data from the specified IP or MAC address to be sent through any port on the module (if one is specified) or on the switch (if no slot is specified).

Delete

This option allows you to delete a filter from the secure access list. The screen displays similar to the following:

```

Delete Filter
-----
Enter Filter Name:

```

If you enter a filter name here, that filter will be immediately deleted from the secure access database.

Modify

This option allows you to modify information about an existing secured access filter. Enter the name of the filter you wish to modify, as follows:

Modify Filter

Filter Name: Test

The filter's existing information will display. For example:

Filter Name	Source IP Address	Source MAC Address	Slot #	Port #
Test	ANY	10.2.8.13	5	2

Enter IP Address ([a.b.c.d]) :
Enter MAC Address ([XXYYZZ: AABBC]) :
Is this MAC in Canonical or Non-Canonical (C or N) [C] :
Enter Slot :
Enter Port :

To change a value, type in the new value at the prompt. If you do not wish to modify a particular field, press **Enter** and the existing user information will remain unchanged. To change a field to **ANY** privilege, enter a value of **0**, an asterisk (*), or **ANY** at the prompt. Descriptions of the fields in the above display are provided earlier under the option "List" on page 8-4.

Find

This option allows you to find information about a specified filter in the secured access database. You must know the filter's name in order to use this search feature. The following is a sample display:

Find Filter

Filter Name: Test

To find a filter in the database, enter the name of the filter at the prompt. If the filter you enter is a valid one, information on that filter will display similar to the following:

Filter Name	Source IP Address	Source MAC Address	Slot #	Port #
Test	ANY	10.2.8.13	5	2

Configuring Secure Access Filter Points

The **secapply** command allows you to view the list of secure access filter points, to enable/disable security globally or for a specific IP protocol filter point, and to define a filter list for each filter point. To use this command, enter:

```
secapply
```

A screen similar to the following displays:

```

                Secure Access Filter Points

1) FTP Security           : Enabled
   11) Filter List       : Test, Engineering
2) Telnet Security       : Disabled
   21) Filter List       : Test
3) SNMP Security         : Enabled
   31) Filter List       :
4) TFTP Security         : Enabled
   41) Filter List       : Manufacturing
5) HTTP Security         : Disabled
   51) Filter List       :
6) Custom Security       : Enabled
   61) Filter List       : HR
   62) Protocol          :
   63) Port Service      :
7) One-touch Global Security :
   71) One-touch Filter List :

Command { Item=Value/?/Help?Quit/Redraw/Save} (Redraw) :
```

◆ Note ◆

If security is enabled for a filter point and there are no names defined on its list, then the filter point is essentially inaccessible to all users. For example, in the above sample display, SNMP is not accessible to any user.

You can enter commands by entering just the first letter of the command. For example, select **Quit** by entering **q** and pressing **<Enter>**. The question mark option (?) and the **Help** option provide reference and instructional information on using this command. The **Quit** option exits this command without saving configuration changes. The **Redraw** option refreshes the screen.

When you are done entering new values, type **save** at the prompt and all new settings will be saved.

The following option is available for all filter points:

Filter List

Applies the filter name(s) defined through the **secdefine** command for this filter point.

Filter points are disabled by default. The different filter points are defined as follows:

1) FTP Security

Indicates whether or not secure access is enabled for File Transfer Protocol (FTP) on the switch. **Enabled** means secure access is enabled for FTP services, and only filters on FTP's filter list have authorization. **Disabled** indicates that secure access is not enabled for FTP services, and all users can access the switch through FTP.

2) Telnet Security

Indicates whether or not secure access is enabled for Telnet service on the switch. **Enabled** means secure access is enabled, and only filters on Telnet's filter list have authorization. **Disabled** indicates that secure access is not enabled for Telnet service, and all users can access the switch through Telnet.

3) SNMP Security

Indicates whether or not security is enabled for Simple Network Management Protocol (SNMP) on the switch. **Enabled** means security is enabled for SNMP services, and only filters on SNMP's filter list are authorized. **Disabled** indicates that secure access is not enabled for SNMP services, and all users can access the switch through SNMP.

4) TFTP Security

Indicates whether or not security is enabled for Trivial File Transfer Protocol on the switch. **Enabled** means security is enabled for TFTP services, and only users on TFTP's filter list are authorized. **Disabled** indicates that security is not enabled for TFTP services, and all users can access the switch through TFTP.

5) HTTP Security

Indicates whether or not security is enabled for HyperText Transfer Protocol (HTTP) on the switch. **Enabled** means that security is enabled for HTTP, and only filters on HTTP's filter list are authorized. **Disabled** indicates that security is not enabled for HTTP, and all users can access the switch through HTTP.

6) Custom Security

Configures whether or not security is enabled for the custom IP protocol specified in line 62. **Enabled** means that security is enabled for the custom IP protocol, and only filters on that protocol's filter list are authorized. **Disabled** indicates that security is not enabled for the custom IP protocol, allowing all users access to the switch through that protocol.

62) Protocol

(Available for Custom Security only.) The IP protocol number to be included as a secured access protocol (IP protocol field in the IP header). You may define only one custom IP protocol.

63) Port Service

(Available for Custom Security only.) The Custom IP protocol's destination port (port field in the IP header)

7) One-touch Security

Configures the same **Security** value for all secure access protocols. **Enabled** enables security for all secure access filter points. **Disabled** disables security for all secure access filter points. Any value configured for individual security parameters overrides the global setting. If you wish to set a different value for **Telnet Security**, for example, enter the line number for Telnet, followed by an equal sign (=) and the new value.

71) One-touch Filter List

Configures a single filter list for all security filter points.

Enabling/Disabling Security Parameters

To change any of the **Security** values, enter the line number for the parameter, followed by an equal sign (=), and then **enabled** or **e** for enable or **disabled** or **d** for disable at the prompt. For example, to enable security for Telnet, enter the following:

```
2=e
```

Adding Filters

To add a filter, at the command prompt, enter the line number for the parameter, followed by an equal sign (=), and then the filter's name at the prompt. For example:

```
21=Test
```

◆ Note ◆

If the filter does not exist in the secure access database, the system prompts you to create the filter. To view the list of secure access filters, use the **secdefine** command. For more information, see “Configuring the Secure Switch Access Filter Database” on page 8-4.

Enter **save** to save the new filter.

Deleting Filters

To remove an existing filter from a filter list, at the command prompt, enter the line number for the parameter, followed by an equal sign (=), a negative sign (-), and then the filter's name as follows:

```
11=-Engineering
```

To remove all filters in a list, include an asterisk after the negative sign. For example:

```
4=-*
```

Enter **save** to save the change.

Viewing Secure Access Violations Log

The **seclog** command displays a log of all secure access violations.

◆ **Note** ◆

To log access violations on the switch, use the **swlogc** command. For more information on the **swlogc** command, see Chapter 10, “Switch Logging.”

To view the secure access violations log, enter

seclog

The following is a sample display:

Secure Access Violations Log					
Time	Protocol	Source IP	Attempts	Slot/ Intf	Elapsed Time (secs)
12:49:02	FTP	172.23.8.801	1	5/1	23
03:15:34	Telnet	198.20.2.101	10	2/3	240

Descriptions of the fields are as follows:

Time. The first time the access violation occurred.

Protocol. The IP protocol for which the violation occurred.

Source IP. The source IP address of the unauthorized user.

Attempts. The number of access attempts made by this user within the sample period (5 minutes).

Slot/Intf. The physical port that received the unauthorized user information.

Elapsed Time (secs). The duration (in seconds) from the first unauthorized access to the end of the sampling period. Secure access violations will take 5 minutes to display in the log file.

Managing User Login Accounts

Prior to software release 4.4, the switch provided security in the form of privilege control for individual login accounts by allocating each user accounts READ or WRITE privileges. Software release 4.4 contains a partition management feature that enhances the privilege capability with an authorization scheme based on the functional capacity assigned to each user.

The purpose of partition management is to provide a mechanism in the switch operating system for system administrators to control access while maintaining enough flexibility to use the switch's full range of services. This is normally done for security reasons. System administrators can partition access to the switch by restricting a user's ability to perform certain switch commands or to use certain command groups.

◆ Terminology Notes◆

A *user account* refers to the user's ability to log onto the switch and perform certain functions. From the user's perspective, it consists of the login name and a password.

A *privilege* refers to the user's ability or permission from the system administrator to execute a command.

Partition Management Requirements

Partition management is available *only* for user login accounts that have *no* permission to use the UI command mode. Where a user account has permission to use the UI mode, partition management is effectively destroyed for that user account. To maintain partition management capability for a user account, that account must be restricted to using the CLI mode only. Refer to "Assigning Account Privileges Using the UI Command Mode" on page 8-16 or "Assigning Account Privileges Using the CLI Command Mode" on page 8-13 for information on restricting use UI commands.

◆ Note◆

Not all UI commands have CLI equivalents. Also, not all CLI commands support partition management. For detailed information, refer to the UI to CLI Cross Reference Tables in Chapter 4 of this manual.

Default Accounts

Initially each switch is preconfigured with three default logins (**admin**, **user** and **diag**). See Chapter 4, “The User Interface,” for more information about login accounts. If you are logged into an account with the WRITE privilege to the USER command you may create or delete login accounts as described in this section. You may also create new user accounts.

◆ Note◆

At least one **user** account with WRITE privileges to use the USER family of commands is required on the switch at all times. If you attempt to remove or modify the only user account to READ-ONLY privilege, the switch will reject the modification command.

There are several commands available for modifying the user login accounts on the switch. To see a list of all user accounts currently available on the switch, use the **userview** command in the UI mode.

Adding a User Account Using the UI Command Mode

To add a user account you must be logged into an account with administrative privileges.

1. At the system prompt enter the **useradd** command. The following prompt displays:

Enter Username: () :

2. Enter the desired user name. The following prompt displays:

Force Password change on next login [y/n] ? (y) :

3. Press **<Enter>** to force a password change at the next login for this user, or enter **n** to keep the configured password. The following prompt displays:

Enter password: () :

4. Enter the desired password. The following prompt displays:

Enter new password again: () :

5. Enter the desired password again. In this example, the username “TechPubs1” is entered. A message similar to the following displays:

User TechPubs1 user privileges (0:0:0) :

The user login account “TechPubs1” is now active on the switch.

At this point the new account has permission to log onto and off of the switch. To add other privileges refer to “Assigning Account Privileges Using the UI Command Mode” on page 8-16 or to “Assigning Account Privileges Using the CLI Command Mode” on page 8-13.

Adding a User Account Using the CLI Command Mode

To add a user account from the CLI mode, you must be logged into an account with administrative privileges. Enter the following at the command prompt.

```
user user_name <password user_password>
```

where *user_name* is the new user login account name and *user_password* is the new user login account password. Both these values are specified by the user. For the user name “Techpubs1”, the following message is displayed:

```
User Techpubs1 created.
```

If you do not specify a password when you create the new account, **switch** becomes the default password.

◆ Note◆

It is recommended that you change the password from the default for all login accounts.

Both the user account name and the password are limited to 16 text characters. The new login account and password will take effect at the user’s next login session.

Assigning Account Privileges Using the CLI Command Mode

A user account’s READ and WRITE privileges can be assigned for all commands or for various subsets of commands. The command subsets referred to as command families are shown here:

config, vlan, iprout, ipxrout, bridge, snmp, xswitch, hrefilter, atmser, atmup, cem, csm, pnni, atmacct, voip, mpoa, mpls and **user**.

In addition to assigning privileges according to command families, an administrator can restrict the user account’s ability to execute specific commands. Here is a list of commands that can be restricted from a user account.

system, status, slot, timeout, prompt, define, prefix, reboot, telnet, ftp, ping, swap, reset, cd, ls, rm, file, interface, ethernet, gated, and **ui**.

◆ Warning◆

If partition management is intended for a user account, that account *cannot* have permission to use the UI command or the UI mode.

User Write Privileges

To assign privileges to a user account, you must be logged into an account with WRITE privileges to the USER family of commands. Enter the following command at the system prompt.

```
user userId [write list-of-families]
```

where *userId* indicates the name assigned to the user account for which you want to assign READ and WRITE privileges. The *list-of-families* parameter indicates the switch command families and the specific commands for which the user account will receive READ and WRITE privileges. Command families must be separated by commas.

User Read Privileges

To assign READ-ONLY privileges to a user account, you must be logged into an account with WRITE privileges to the USER family of commands. Enter the following command at the system prompt.

```
user userId [read list-of-families]
```

where *userId* indicates the name assigned to the new login account for which you want to assign READ-ONLY privileges. The *list-of-families* parameter indicates the switch command families and the specific commands for which the user account will receive READ-ONLY privileges. For a list of command families and specific commands, refer to the “Assigning Account Privileges Using the UI Command Mode” section on page 8-16 or to “Adding a User Account Using the CLI Command Mode” on page 8-13.

Removing Privileges

You can remove READ and WRITE privileges from a user created login account if you are logged into an account with WRITE privileges to the USER command family. Use the following command:

```
user userId no write list-of-families
```

You can remove READ-ONLY privileges from a user created login account by using the following command:

```
user userId no read list-of-families
```

For both these commands, the *userId* parameter indicates the name assigned to the user created login account for which you want to remove privileges. The *list-of-families* parameter indicates the switch command families and the specific commands from which you want to remove READ or WRITE privileges.

Miscellaneous CLI Privileges Commands

The following is a list of privileges-related CLI commands. For more details on these commands and other CLI commands, refer to the *Text-Based Configuration CLI Reference Guide*.

- To create a new user login account, use the following command:

```
user user_name [password user-password]
```

where *user_name* is the new user login account name and *user-password* is the new user password. Both these values are defined by the user.

- To set or change the password of the current user account, use the following command:

```
password password
```

Where *password* is the new *password* for this user account.

- To delete a login account, use the following command:

```
no user user_name
```

where *user_name* is the current login you want to delete.

- To view user privileges for a specific user login account, use the following command:

```
view user [user_name]
```

where *user_name* is the name of the user login account for which you will view privileges.

Assigning Account Privileges Using the UI Command Mode

When you add a new user login account, the account has permission to log in and to log out. If you want the new account to have additional privileges you must add them separately. To add privileges to a user account, you must be logged into an account with administrative privileges. From the system prompt enter the **usermod** command. The following prompt displays:

Enter Username : () :

Enter the login name of the user account you are modifying. The following screen will display.

```
- CONFIG      : NO
- GROUP       : NO
- IPROUT      : NO
- IPXROUT     : NO
- BRIDGE      : NO
- SNMP        : NO
- XSWITCH     : NO
- HREFILTER   : NO
- ATMSEAR     : NO
- ATMUP       : NO
- CEM         : NO
- CSM         : NO
- PNNI        : NO
- ATMACCT     : NO
- VOIP        : NO
- MPOA        : NO
- MPLS        : NO
- USER       : NO
Subsets of the global family:
- SYSTEM      : NO
- STATUS      : NO
- SLOT        : NO
- TIMEOUT     : NO
- PROMPT      : NO
- DEFINE      : NO
- PREFIX      : NO
- REBOOT      : NO
- TELNET      : NO
- FTP         : NO
- PING        : NO
- SWAP        : NO
- RESET       : NO
- CD          : NO
- LS          : NO
- FM          : NO
- FILE        : NO
- INTERFACE   : NO
- ETHERNET    : NO
- GATED       : NO
- UI          : NO
1. MODIFY ONE FAMILY RIGHTS
2. SET ALL READ RIGHTS
3. SET ALL WRITE RIGHTS
4. SET NO READ RIGHTS
5. SET NO WRITE RIGHTS
6. MODIFY ONE GLOBAL SUBSET
7. SET NO GLOBAL SUBSET
8. SET ALL GLOBAL SUBSET
```

[1 TO 8, (c)ancel or (s)sav] () :

This screen displays the default privileges for a new user login account. Note that the default privileges give the new user neither read nor write permission. To grant privileges to the user account, enter a number from 1 to 5 as indicated in the display. To set WRITE privileges for a single family of commands, enter **1** and press **<Enter>**. The display will prompt you for the family number as shown here:

Give the family number : () :

Enter the number of the command family for which you want to set WRITE privileges. Refer to the “Command Family Table” on page 8-18 for the number.

For example, if you wanted to enable WRITE privileges for the Bridge command family, enter the number **5** as shown here.

Give the family number : () : 5

The following will display.

```
Give rights on family BRIDGE
  0.      NO
  1.      READ
  2.      WRITE
  3.      READ&WRITE
( ) :
```

Enter the number **2** at the prompt to assign WRITE privileges. The following shows a portion of the display.

```
User 'TechPubs1' user privileges (0:0X20:0) :
- CONFIG      : NO
- GROUP       : NO
- IPRROUT     : NO
- IPXROUT     : NO
- BRIDGE      : READ & WRITE
- SNMP        : NO
- XSWITCH     : NO
(Continued)
```

The privilege listed next to Bridge shows WRITE. This indicates that the user “TechPubs1” now has WRITE privileges for the Bridge family of commands.

Command Family Table

Number	Command Family
1	Configuration
2	Group
3	IP Routing
4	IPX Routing
5	Bridge
6	SNMP
7	QOS Policy
8	HRE Filter
9	ATM Service
10	WAN
11	CSM
12	PNNI
13	ATM Accounting
14	Voice Over IP
15	MPOA
16	MPLS
17	(unsupported)
18	User

The global family contains commands that apply globally to the switch rather than to individual applications or services. Privileges for global family commands can be set on an individual command basis or altogether so the privilege applies to the whole global family. If you want to set privileges for the global commands, you must enter 6, 7 or 8 when the screen prompt displays the following:

1. MODIFY ONE FAMILY RIGHTS
2. SET ALL READ RIGHTS
3. SET ALL WRITE RIGHTS
4. SET NO READ RIGHTS
5. SET NO WRITE RIGHTS
6. MODIFY ONE GLOBAL SUBSET
7. SET NO GLOBAL SUBSET
8. SET ALL GLOBAL SUBSET

[1 TO 8, (c)ancel or (s)ave] () :

To give the user account the privilege to set all global commands, enter the numeral 8. To deny the user the privilege to set any of the global commands, enter the numeral 7. To set individual global commands, enter the number 6. If you are assigning privileges on an individual command basis the display will look like this:

[1 TO 8, (c)ancel or (s)sav] () : 6
Give the subset number : () :

Enter the number of the command for which you want to set WRITE privileges. Refer to the “Global Family Table” on page 8-19 for the number.

Global Family Table

Number	Global Family
1	System
2	Status
3	Slot
4	Timeout
5	Prompt
6	Define
7	Prefix
8	Reboot
9	Telnet
10	FTP
11	Ping
12	Swap
13	Reset
14	CD
15	LS
16	RM
17	File
18	Interface
19	Ethernet
20	Gated
21	UI

For example, if you wanted to assign the user account the privilege to use the define command, enter the number 6 as shown here.

Give the family number : () : 6

The following will display.

Give rights on subset DEFINE

- 0. NO
- 1. YES

() :

If you enter 1, all the command families will display and the DEFINE command under the global family will be shown as follows:

- DEFINE : YES

After you set the user account privileges, the switch displays the current configuration. At this point you may enter **s** to save your configuration or **c** to cancel.

◆ Warning◆

If partition management is implemented on a user account, that account *must* have the UI command family set to NO privilege. If an account has the privilege to use the UI command, partition management is effectively destroyed for that account.

Modifying a User Account

You can use the **usermod** command to modify account privileges as shown here. You must be logged into a user account with administrative privileges.

1. At the system prompt enter the **usermod** command. A prompt similar to the following displays:

```
Enter Username: ( ) :
```

2. Enter the name assigned to the user account you want to modify. A screen similar to the following displays where the account name is **TechPubs1**.

```
User 'TechPubs1' is configured with the following privileges:  
READ
```

1. READ
2. WRITE
3. ADMIN
4. FORCE new password

```
Select the privilege(s) number to add/remove.  
[ 1, 2, 3 (c)ancel or (s)ave] (c) :
```

◆ Note ◆

See “Managing User Login Accounts” on page 8-11 for definitions of the privileges.

3. Enter the number for the privilege you want to add or remove. The entry acts as a toggle to turn the privilege on or off for the user. In the current example, if you enter **2** at the prompt, a screen similar to the following displays:

```
User 'TechPubs1' is configured with the following privileges:  
READ WRITE
```

4. After modifying the privileges for the user, enter **s** at the selection prompt to save the change(s).

Deleting a User

To delete a user from the user database, you must be logged into an account with administrative privileges.

1. At the system prompt, enter the **userdel** command. The following prompt displays:

```
Enter Username to remove: ( ) :
```

2. Enter the username for the user you want to delete. A message similar to the following displays:

```
User 'TechPubs1' was removed.
```

◆ Note ◆

All users but one may be deleted from the switch, provided that the one remaining user is configured with all privileges.

9 Configuring Switch-Wide Parameters

The switch provides commands to display and configure parameters on a switch-wide basis. These commands are grouped into two menus: the Summary menu and the System menu. Descriptions for commands in the Summary menu begin below; descriptions for commands in the System menu begin on page 9-5.

In addition, this chapter contains documentation for configuring HRE-X ports (described in *Configuring the HRE-X Router Port* on page 9-27) duplicate MAC address support (described in *Duplicate MAC Address Support* on page 9-30), multicast claiming (described in *Multicast Claiming* on page 9-32), disabling flood limits (described in *Disabling Flood Limits* on page 9-32), and saving configurations (described in *Saving Configurations* on page 9-33).

Summary Menu

The Summary menu consists of commands for displaying summary switch information. To access this menu, enter

summary

at the UI prompt. Type the question mark (?) to see the following list of commands.

<u>Command</u>	<u>Summary Menu</u>
ss	Display MIB-II System group variables
sc	Display a summary of the chassis (type, id, serial no., base mac, etc.)
si	Current interface status

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The Summary menu commands are described in the sections that follow.

Displaying the MIB-II System Group Variables

MIB-II is a core set of definitions created to define the SNMP-based management framework. This MIB module contains definitions for both end systems and routers using the Internet protocol suite. To display the MIB-II system group variables, enter

```
ss
```

at the system prompt. A screen similar to the following will be displayed.

```
System description:   Alcatel Omni Switch/Router
System Object ID:    1.3.6.1.4.1.800.3.1.1.2.
Agent Up Time:       5 days, 00:28:14.38
Contact:             Administrator
Name:                TechWrite
Location:            Bldg 46
Device Services:
  DataLink/Subnetwork Layer
  Internetwork Layer
  Host Layernetwork Layer
  Application Layer (Rlogin, Telnet, FTP)
```

The fields displayed by the **ss** command are described below.

System description. The specific type of chassis, which can be an OmniSwitch, OmniAccess, or Omni Switch/Router. This field is set by the **syscfg** command, which is described in *Configuring System Information* on page 9-23.

System Object ID. The MIB entry for the switch (where the object ID starts). This is read only. This value helps you locate Alcatel-specific variables in the MIB tree.

Agent Up Time. The time (in days, hours, minutes, and seconds) since the switch was re-initialized.

Contact. The name of a person to contact about this switch. This field is set by the **syscfg** command, which is described in *Configuring System Information* on page 9-23.

Name. The name the system administrator assigned to this switch (the node's fully qualified domain name, by convention). This field is set by the **syscfg** command, which is described in *Configuring System Information* on page 9-23.

Location. The physical location of the switch. This field is set by the **syscfg** command, which is described in *Configuring System Information* on page 9-23.

Device Services. The type of services provided by the switch. Supported service types are listed below:

- **Data Link /Subnetwork Layer**
- **Internetwork Layer**
- **Host Layer**
- **Application Layer (Rlogin Telnet, FTP)**

Displaying the Chassis Summary

To display the chassis summary information, enter

```
sc
```

at the system prompt. A screen similar to the following will be displayed.

```

Type:                Omni Switch/Router XFRAME 9-slot
Chassis ID:         Alcatel
Description:        DESCRIPTION NOT SET.
Backplane:          5 SLOT
Master MPM Serial No.: 52601675
Physical Changes:   7
Logical Changes:    0
Number of Resets:   26
Base MAC Address:   00:20:da:02:04:80
Free Slots:         0

```

The fields displayed by the **sc** command are described below.

Type. The description of the specific type of chassis or device.

Chassis ID. The chassis ID for this switch.

Description. The description of this chassis. This field is set by the **syscfg** command, which is described in *Configuring System Information* on page 9-23.

Backplane. The style of backplane in this chassis.

Master MPM Serial No. The serial number for the primary MPX.

Physical Changes. The number of physical changes that has occurred since the last reset or power-on.

Logical Changes. The number of logical changes that has occurred since the last reset or power-on.

Number of Resets. The number of times this switch has been reset since the configuration file (**mpm.cnf**) was first removed.

Base MAC Address. The base MAC address for the primary MPX.

Free Slots. The number of front panel slots not occupied by a switching module.

Displaying Current Router Interface Status

To display current interface status information, enter

si

at the system prompt. A screen similar to the following will be displayed.

Interface Summary Status			
4 Interfaces			
Logical Interface	Interface Type	Administrative Status	Operational Status
1	Slip	Enabled	Enabled
2	Virtual Router	Enabled	Active
3	Virtual Router	Enabled	Active
4	SoftwareLoopback	Enabled	Enabled

The fields displayed by the **si** command are described below.

Logical Interface. A number, in sequence, that has been assigned to the virtual router port.

Interface Type. The type of interface, which can be virtual router (the standard interface type), SLIP, and software loopback.

Administrative Status. Whether the administrator has enabled or disabled the port. The port can be enabled by the administrator but still be made inactive by the system.

Operational Status. Whether the port is active (operational) or inactive. This status is set by the system software.

System Menu

The System menu contains commands to view or set system-specific parameters. To access this menu, enter

system

at the UI prompt to enter the System menu. If you are not in verbose mode, press a question mark (?) and then press **<Enter>** to display the commands in the system menu, as shown below.

<u>Command</u>	<u>System Menu</u>
info	Basic info on this system
dt	Set system date and time
ser	View or configure the DTE or DCE port
mpm	Configure a Management Processor Module
slot	View Slot Table information
systat	View system stats related to system, power and environment
taskstat	View task utilization stats
taskshow	View detailed task information
memstat	View memory use statistics
fsck	Perform a file system check on the flash file system
newfs	Erase all file from /flash and create a new file system
syscfg	View/Configure info related to this system
uic	UI configuration; change - prompt, timeout, more, verbose.
camstat	View CAM info and usage
camcfg	Configure CAM info and usage
hrex	Enter HRE-X management command sub-menu
ver/ter	Enables/disables automatic display of menus on entry (obsolete)
echo/noecho	Enable/disable character echo
chpr	Change the prompt for the system (obsolete, use 'uic' command)
logging	View system logs.
health	Set health parameters or view health statistics
cli	Enter command line interface
saveconfig	Dump the cache configuration content to the mpm.cnf file.
cacheconfig	Set the flag to use cache configuration only.

Main File Summary VLAN Networking
Interface Security System Services Help

All of the System menu commands—except for the **mpm**, **ver**, **ter**, **echo**, **noecho**, **chpr**, **logging**, **health**, and **cli** commands—are described in the following sections. The **uic**, **ver/ter**, **echo**, **noecho**, **chpr**, and **cli** commands are described in Chapter 4, “The User Interface.” The **mpm** command is described in Chapter 6, “Configuring Management Processor Modules.”

◆ Note ◆

The **ver**, **ter**, and **chpr** commands now appear as items in the UI Configuration menu (displayed through the **uic** command). If you enter the **ver/ter** and **chpr** commands, a message will advise you to use the **uic** command, and the UI Configuration menu will automatically display. For more information on the UI Configuration menu, refer to Chapter 4, “The User Interface.”

Displaying Basic System Information

To display basic information on the switch, enter

```
info
```

at the system prompt. The following display is a typical example.

```
System Make: Alcatel OmniSwitch
System Type: 5-slot OmniSwitch
Description: DESCRIPTION NOT SET.

Backplane: 9 SLOT                Bus Speed: 1200 XFRAME

Physical changes to the system since power-up or reset: 2
Logical changes to the system since power-up or reset: 0
Number of Resets to this system: 8

The attached MPM, slot 1, is the Primary
Automatic configuration synchronization is enabled

System base MAC Address: 00:20:da:04:21:f0
Number of Free Slots: 0
Action on Cold Start: Load & go
Action on Reset: Restart

VBus Mode : Mode 1

Script File: /flash/mpx.cmd
Boot File: /flash/mpx.img
Ni Image Suffix: img
```

The fields displayed by the **info** command are described below.

System Make. The description of the specific type of chassis or device.

System Type. The OmniSwitch type.

Description. A description of the chassis and product. This field is set by the **syscfg** command, which is described in *Configuring System Information* on page 9-23.

Backplane. The style of backplane used in this chassis.

Bus Speed. The speed of backplane, in Mbs, used in this chassis.

Physical Changes to the system since power-up or reset. The number of physical changes that has occurred since the last reset or power-on.

Logical Changes to the system since power-up or reset. The number of logical changes that has occurred since the last reset or power-on.

No. of Resets to the System. The number of times this switch has been reset since the last cold start.

◆ **Note** ◆

The **info** command will also display the number of MPXs, their location in chassis, and which one is the primary and which one is the secondary. In addition, it also displays whether automatic configuration synchronization is enabled. See Chapter 6, “Configuring Management Processor Modules,” for more information on redundant MPXs and automatic configuration synchronization.

System Base MAC Address. The base MAC address for the primary MPX in chassis.

Number of Free Slots. The number of slots not occupied by a module.

Action on Cold Start. The action taken when you switch the power on.

Action on Reset. The action taken when you reboot.

Script File. The name of the command file (**mpx.cmd** is the default) containing user-configurable commands.

Boot File. The boot file (**mpx.img** is the default) used by the switch when it boots up or reboots.

Ni Image Suffix. The name of the file extension (**img** is the default) indicating that the file is an executable binary file. See Chapter 6, “Configuring Management Processor Modules,” to change this suffix.

Setting the System Date and Time

The **dt** command allows you to set the local date, time, and time zone. Additionally, you can set the system clock to run on Universal Time Coordinate (UTC or GMT). If applicable, you can also configure Daylight Savings Time (DST) parameters. To view or make changes to date, time, time zone, and DST for the switch, enter

dt

at the System prompt. This command displays a screen similar to the following:

Modify Date and Time Configuration

```
1) Local time                : 1:45:41
2) Local date                : 01/15/01
3) Timezone (-13 . . 12, name) : MST   UTC-7 hrs
4) Daylight Savings Time active : DisabledCommand
{Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :
```

To use the **dt** command, you must have UI write privileges. Enter the line number for the variable that you would like to change, an equal sign (=), and then the new value for the variable. For example, to set a new date, you would enter:

2=4/20/99

After you have made changes, enter

save

to save your changes and to exit the **dt** menu. If you do not wish to make any changes, enter

quit

at the system prompt. The following sections describe the variables on this screen.

1) Local time

Indicates the current and local time. To set the time, enter the line number for **Local Time (1)** followed by the new time. The time format is as follows:

HH:MM:SS

where **HH** is the hour to be set based on a 24 hour (military) clock, **MM** is the minutes to be set, and **SS** is the seconds to be set. For example, if you wanted to set the time to 3:15 p.m., you would enter:

1=15:15:00

2) Local date

The current and local date. To set the date, enter the line number for **Local Date (2)** followed by the new date. The date format is as follows:

MM/DD/YY

where **MM** is the month to be set, **DD** is the day to be set, and **YY** is the last two digits of the year to be set. Remember to include a slash (/) between the month and the day and between the day and the year. For example, if you wanted to set the date to January 15, 2001, you would enter:

2=01/15/01

3) Timezone

This parameter specifies the time zone for the switch and sets the system clock to run on UTC time (or Greenwich Mean Time). Additionally, if Daylight Savings Time is enabled (see option 4 below), the clock automatically sets up default DST parameters (if applicable) for the local time zone. The local time remains active for all User Interface commands and other subsystems that require the local time. To set the time zone for the switch, you may use one of two methods:

- a. Enter the line number for **Timezone (3)** followed by the hour(s) offset from UTC. This can be a number from -13 to +12. The number you enter will set the system clock x hours from the local time. For example, if the local time, 1:45:00, is seven hours behind UTC time, you would enter:

3=-7

This specification sets the UTC time to 8:45:00, seven hours ahead of the local time, 1:45:00.

- b. Enter the line number for **Timezone (3)** followed by the time zone name. There is a limited number of time zone names available. For example, if the local time zone name is Mountain Standard Time (MST), you would enter:

3=MST

This specification automatically sets the switch to -7 hours, the number of hours MST is offset from UTC.

Daylight Savings Time. The software will automatically configure DST values for a specified time zone. However, the user can manually modify DST values.

Non-integer Offsets. Non-integer offsets are acceptable for **Timezone**. Some parts of the world are offset from UTC by increments of 15, 30, or 45 minutes. India, for example, is offset from UTC by 5 hours and 30 minutes. If you wanted to enter the time zone offset for India, for example, you would type the line number for Timezone (3), followed by the non-integer hour offset in the **HH:MM** format, as follows:

3=05:30

where the value of **05:30** is five hours and thirty minutes offset from UTC.

◆ Note ◆

The switch automatically enables UTC. However, if you do not want your system clock to run on UTC, simply enter the offset **+0** for the **Timezone** parameter. This sets UTC to run on local time.

The table on the following page lists the options available for **Timezone** names:

Timezone and DST Parameters

Abbr.	Name	Hours from UTC	DST Start	DST End	DST Change
NZST	New Zealand	+12:00	1st Sunday in Oct. at 2:00 a.m.	3rd Sunday in March at 3:00 a.m.	1:00
ZP11	No standard name	+11:00	No default	No default	No default
AEST	Australia East	+10:00	Last Sunday in Oct. at 2:00 a.m.	Last Sunday in March at 3:00 a.m.	1:00
GST	Guam	+10:00	No default	No default	No default
ACST	Australia Central Time	+9:30	Last Sunday in Oct. at 2:00 a.m.	Last Sunday in March at 3:00 a.m.	1:00
JST	Japan	+9:00	No default	No default	No default
KST	Korea	+9:00	No default	No default	No default
AWST	Australia West Time	+8:00	No default	No default	No default
ZP8	China, Manila, Philippines	+8:00	No default	No default	No default
ZP7	Bangkok	+7:00	No default	No default	No default
ZP6	No standard name	+6:00	No default	No default	No default
ZP5	No standard name	+5:00	No default	No default	No default
ZP4	No standard name	+4:00	No default	No default	No default
MSK	Moscow	+3:00	Last Sunday in March at 2:00 a.m.	Last Sunday in Oct. at 3:00 a.m.	1:00
EET	Eastern Europe	+2:00	Last Sunday in March at 2:00 a.m.	Last Sunday in Oct. at 3:00 a.m.	1:00
CET	Central Europe	+1:00	Last Sunday in March at 2:00 a.m.	Last Sunday in Oct. at 3:00 a.m.	1:00
MET	Middle European Time	+1:00	Last Sunday in March at 2:00 a.m.	Last Sunday in Oct. at 3:00 a.m.	1:00
BST	British Standard Time	+0:00	Last Sunday in March at 1:00 a.m.	Last Sunday in Oct. at 3:00 a.m.	1:00
WET	Western Europe	+0:00	Last Sunday in March at 1:00 a.m.	Last Sunday in Oct. at 3:00 a.m.	1:00

Timezone and DST Parameters Con't

Abbr.	Name	Hours from UTC	DST Start	DST End	DST Change
GMT	Greenwich Mean Time	+0:00	No default	No default	No default
WAT	West Africa	-1:00	No default	No default	No default
ZM2	No standard name	-2:00	No default	No default	No default
ZM3	No standard name	-3:00	No default	No default	No default
NST	Newfoundland	-3:30	1st Sunday in April at 2:00 a.m.	Last Sunday in Oct. at 2:00 a.m.	1:00
AST	Atlantic Standard Time	-4:00	1st Sunday in April at 2:00 a.m.	Last Sunday in Oct. at 2:00 a.m.	1:00
EST	Eastern Standard Time	-5:00	1st Sunday in April at 2:00 a.m.	Last Sunday in Oct. at 2:00 a.m.	1:00
CST	Central Standard Time	-6:00	1st Sunday in April at 2:00 a.m.	Last Sunday in Oct. at 2:00 a.m.	1:00
MST	Mountain Standard Time	-7:00	1st Sunday in April at 2:00 a.m.	Last Sunday in Oct. at 2:00 a.m.	1:00
PST	Pacific Standard Time	-8:00	1st Sunday in April at 2:00 a.m.	Last Sunday in Oct. at 2:00 a.m.	1:00
AKST	Alaska	-9:00	1st Sunday in April at 2:00 a.m.	Last Sunday in Oct. at 2:00 a.m.	1:00
HST	Hawaii	-10:00	No default	No default	No default
ZM11	No standard name	-11:00	No default	No default	No default

4) Daylight Savings Time active

Enables and disables DST (Daylight Savings Time). To enable DST, enter:

4=Enable

To disable DST, enter:

4=Disable

If DST is disabled, options 41-49 will not be displayed.

41) DST Start Month

Indicates which month of the year DST starts. To set the month when DST should start, enter the sequential number of the month (January=1, February=2, . . . December=12). For example, if you want DST to begin in April, you would enter the line number for **DST Start Month (41)** and the month, as follows:

41=4

42) DST Start Week

Indicates which week in a month DST starts. To set the week DST should start, enter the sequential number of the week. The possible values are 1st (1), 2nd (2), 3rd (3), 4th (4), and Last. For example, if you want DST to start on the 3rd Tuesday of a month, you would enter the line number for **DST Start Week (42)** and the week, as follows:

42=3

43) DST Start Day

Indicates which day of the week DST starts. To set the day DST should start, enter the sequential number of the day (Sunday=1, Monday=2, . . . Saturday=7). For example, if you want DST to begin on Friday, you would enter the line number for **DST Start Day (43)** and the day, as follows:

43=6

44) DST Start Time

Indicates what time of day (in local time) DST starts. To set the time DST should start, enter the time in the form **HH:MM**, where **HH** is the clock hours of a 24 hour (military) clock and **MM** is the clock minutes that DST should start. For example, if you want DST to start at 1:00 a.m., you would enter the line number for **DST Start Time (44)** and the time, as follows:

44=1:00

45) DST End Month

Indicates which month of the year DST ends. To set the month DST should end, enter the sequential number of the month (January=1, February=2, . . . December=12). For example, if you want DST to end in April, you would enter the line number for **DST End Month (45)** and the month, as follows:

45=4

46) DST End Week

Indicates which week in a month DST ends. To set the week DST should end, enter the sequential number of the week. The possible values are 1st (1), 2nd (2), 3rd (3), 4th (4), and Last. For example, if you want DST to end on the last Tuesday of a month, you would enter the line number for **DST End Week (46)** and the week, as follows:

46=Last

47) DST End Day

Indicates which day of the week DST ends. To set the day DST should end, enter the sequential number of the day (Sunday=1, Monday=2, . . . Saturday=7). For example, if you want DST to end on Wednesday, you would enter the line number for **DST End Day (47)** and the day, as follows:

47=4

48) DST End Time

Indicates what time of day (in local time) DST ends. To set the time DST should end, enter the time in the form of **HH:MM**, where **HH** is the clock hours of a 24 hour (military) clock and **MM** is the clock minutes that DST should end. For example, if you want DST to end at 2:00 a.m., you would enter the line number for **DST End Time (48)** and the time, as follows:

48=2:00

49) DST Offset

Indicates the amount of time to change the local time when DST changes. To set how much time DST should change, enter the change in the form of **HH:MM**, where **HH** is the clock hours and **MM** is the clock minutes that DST should change. For example, if you want the local time to move 1 hour when **DST** changes, you would enter the line number for **DST Offset** and the hour, as follows:

49=1:00

Viewing Slot Data

You can view slot table information by entering the **slot** command. To view information on a particular slot, enter the **slot** command together with the slot number. For example, to view information for slot 1, enter

```
slot 1
```

at the system prompt. You can also view information on all slots in the switch at the same time in a table. To view data, for all slots in the switch, enter

```
slot
```

at the system prompt. A table similar to the following will be displayed.

Slot	Module-Type Part-Number	Adm-Status Oper-Status	HW Rev	Board Serial #	Mfg Date	Firmware-Version Base-MAC-Address
1*	MPM	Enabled	L3	52601675	01/05/01	4.305002600 Operational 00:20:da:04:21:f0
2	HSM	Enabled	B11	53404264	01/19/01	4.3 05003106 Operational 00:20:da:02:28:60
2-1	FDDI		D	53404104	01/24/01	05003706
3	HSM Enabled	Enabled	L	53404645	01/21/01	4.3 05003106 Operational 00:20:da:04:87:30
3-1	ATM		B	53404116	01/11/01	05004400
4	Ether/8	Enabled	D	53404229	01/07/01	4.3 05000014 Operational 00:20:da:03:09:90
5	F-Ether/M	Enabled	A5	73250839	01/07/01	4.3 05015906 Operational 00:20:da:85:40:50

The fields display by the **slot** command are described below.

Slot. The slot number for the MPX or switching module.

Module-Type. The type of module in this slot.

Part-Number. The factory-assigned part number.

Adm-Status. The administration status. This can be enabled or disabled by the operator through the **reset** command, which is described in Chapter 36, "Running Hardware Diagnostics."

Oper-Status. The operational status. Whether the port is Up (Operational), Down, or Unknown. (Unknown means uninitialized or that the module is in a transitional state.)

HW Rev. The revision number for this module. This number may be helpful when troubleshooting.

Board Serial #. Serial number for this module.

Mfg Date. The manufacturing date for this module.

Firmware-Version. The version of the module's firmware. All modules should use the same version of software.

Base-MAC-Address. The base MAC address(es) of this module.

Viewing System Statistics

The **systat** command displays statistics related to system, power, and environment. To view these parameters, enter

```
systat
```

at the system prompt. A screen similar to the following will be displayed.

```

System Uptime                1 days, 12:09:22.64
MPM Transmit Overruns       : 0
MPM Receive Overruns       : 22
MPM total memory            : 16 MB
MPM free memory             : 6522536 bytes
MPM CPU Utilization ( 5 sec) : 5% ( 0% intr 0% kernel 3% task 95% idle)
MPM CPU Utilization ( 60 sec) : 5% ( 0% intr 0% kernel 3% task 96% idle)
Power Supply 1 State        : OK
Power Supply 2 State        : Not Present
Temperature Sensor          : OK - Under Threshold

Temperature                  : 37:00c 98.60f
Temperature Alarm Masking    : Disabled

```

The fields displayed by the **systat** command are described below.

System Uptime. The time since the last boot that the system has been running, displayed in days, hours, minutes, and seconds (to the nearest hundredth).

MPM Transmit Overruns. The number of times a VSE transmit buffer could not be allocated by a task on the MPX.

MPM Receive Overruns. The number of times packets were dropped because the bus had more packets to deliver than the MPX could handle. This is a “receive overrun” condition which can happen when a storm occurs or when the switch is first powered up and many unknown MAC frames are being forwarded to the MPX.

MPM total memory. The amount of total memory installed on the MPX.

MPM Free Memory. The amount of free, or unused, memory available in the MPX. This data is also displayed by the **memstat** command, which is described in *Viewing MPX Memory Statistics* on page 9-20.

MPM CPU Utilization (5 seconds). The amount of time, by percent, the MPX processor actually worked during the last 5 seconds.

MPM CPU Utilization (60 sec). The amount of time, by percent, that the MPX processor actually did work during the last minute.

Power Supply 1 State. Valid states are **OK**, **Not Present**, and **Bad**. A power supply that has been turned off will be in the **Bad** state. If not installed, it will be in the **Not Present** state.

Power Supply 2 State. Valid states are **OK**, **Not Present**, and **Bad**. A power supply that has been turned off will be in the **Bad** state. If not installed, it will be in the **Not Present** state.

Temperature Sensor. Indicates whether the MPX temperature sensor detects overheating. Valid states are **Under Threshold**, **Over Threshold**, and **Not Present**.

Temperature. Indicates the switch temperature Celsius and Fahrenheit.

Temperature Alarm Masking. Indicates whether temperature alarm masking is Enabled or Disabled. You enable masking through the **maskta** command, which is described in Chapter 36, “Running Hardware Diagnostics.”

Clearing System Statistics

You may want to clear statistics for a specific module, port or service for dialogistic or accounting purposes. To clear switch statistics enter

clearstat

at the system prompt. A screen similar to the following will display.

Usage: clearstat slot [,port] [,service] [,instance]

As indicated in the prompt, you can clear all statistics from a module by entering the slot number as shown here:

clearstat 3

This entry will clear all statistics for the module located in slot 3. If you want to clear statistics for a specific port, service or instance, enter the **clearstat** command followed by the appropriate numbers. You must use a comma (,) to separate the slot number from the port, service and instance numbers. The following command will clear all statistics for port 1 of the module located in slot 3.

clearstat 3,1

◆ Caution◆

When the **clearstat** command is used, no notification is sent to the SNMP manager about the cleared statistics. Use of this command can cause unpredictable results with your NMS statistics.

Viewing Task Utilization Statistics

The **taskstat** command displays the task utilization statistics of the switch. To display the task utilization statistics, enter

```
taskstat <task-number> <sample-period>
```

at the system prompt. The **<task-number>** is an optional number of tasks and the **<sample-period>** is an optional sample period of 1 to 60 seconds. You must enter the **<task-number>** if you want to enter the **<sample-period>**.

The default number for **<task-number>** is 5 and the default sample period for **<sample-period>** is 5 seconds. To display the task utilizations statistics for 10 tasks over a 20-second period, for example, enter

```
taskstat 10 20
```

at the system prompt. A screen similar to the following will display.

Task Name	Utilization (20 secs)
tUi_shellt0	0.76%
tCMProber	0.70%
tUi_shellC	0.60%
tSnmp_agent	0.34%
tNetTask	0.32%
tTelnetOut0	0.19%
tif_vblInput	0.19%
vseReceive	0.11%
tTelnetIn0	0.08%
bslMgr	0.07%
All Other Tasks:	0.68%
Total Task Utilization:	4.04%

The **taskstat** command displays the tasks in descending order in terms of the switch's CPU utilization. You may use the **taskstat 0** command if you want to list utilization statistics for all the tasks executed by the switch.

The **taskshow** command displays a table listing all tasks and their priority, status and memory allocation. A partial table is shown here.

NAME	ENTRY	TID	PRI	STATUS	PC	SP	ERRNO	DELAY
tExcTask	_excTask	499f7f20	0	PEND	4892067c	499f7d38	9	0
tLogTask	_logTask	499f5598	0	PEND	4892067c	499f53b0	0	0
tCMWatcher	_cmWatchdogK	4999f108	0	DELAY	4893c028	4999efb8	0	5
tHelperTask	_exc2Task	499fc018	2	PEND	4892067c	499fbc30	0	0
tAscSTimer	_ascSessTime	49a53498	10	DELAY	4893c028	49a53348	0	170
bpeMgr	_bpm_initial	46037630	20	PEND	4892a41c	46037430	3d0002	0
ipxTimer	_ipxTimerTas	49a83168	49	DELAY	4893c028	49a83010	0	26
ipxGapper	_ipxGapperTa	49a7cdc0	49	PEND	4892067c	49a7cb70	0	0
tNetTask	_netTask	499eee40	50	PEND	4892a0a4	499eec68	0	0
ipx	_ipxMain	49fe0350	50	PEND	4892a41c	49fe0168	3d0002	0

The fields displayed by the **taskshow** command are described below.

NAME. Name of the task whose statistics are being shown.

ENTRY. Shows the routines that are currently being executed by the specified task.

TID. Address of the task listed in this row.

Viewing Task Utilization Statistics

PRI. Priority of the specified task.

STATUS. Current status of the specified task.

PC. Program Counter. The program counter identifies the routing code as it enters the stack.

SP. Stack pointer. The stack pointer points to the code being loaded when the status is taken.

ERRNO. Error number indicator.

DELAY. The time elapsed between task routines.

Viewing Memory Utilization

The leak monitor diagnostic utility is used to display information about memory utilization. This utility requires the use of three UI commands: **leakstart**, **leakstop** and **leakdumpall**.

◆ Note◆

You may want to log this operation to a text file to make it easier to view the data.

To start the utility, enter

```
leakstart
```

at the system prompt. This command starts a leak monitor daemon that gathers memory information in the background until you stop it by using the **leakstop** command. The **leakstop** command stops the leak monitor daemon from recording data and preserves the data already recorded. To view the memory utilization information enter the following command

```
leakdumpall
```

at the system prompt. This command dumps all memory recorded by the leak daemon. A screen similar to the following will display.

```

Outstanding Memory - at TUE  APR  24  19:00:29  2001

Task ID   Name   Functi 1  Functi 2   Functi 3   Address  Len   Time
=====  =====  =====  =====  =====  =====  ==  =====
49a69a58  tUi_she 484fe4do 484f1284 484ffbc8 4800ef28  9 TUE APR 24 18:06:4 7 2001
49559bb8  t_AtMg 49db6e90 49d6a780 49d4c3bd 4800ef88 16 TUE APR 24 18:06:4 6 2001
49559bb8  t_AtMg 49db6e90 49d4be4c 49d8639c 4800efb8 64 TUE APR 24 18:06:4 6 2001
49559bb8  tUi_she 49db6e90 49d9cce4 49d9c910 4800f050  4 TUE APR 24 18:06:4 6 2001

```

End of memory report.

The length of the display shown will vary depending on the length of time between use of the **leakmon** command and the **leakstop** command. The fields displayed by the **leakdumpall** command are described below.

Task ID. The address of the task that is allocating the block of memory.

Name. Name of the task that is allocating the block of memory.

Functi 1, 2, 3. These three columns indicate functions entered above the *malloc* package. Function 1 is the function that called *malloc*. Function 2 is the function that called Function 1. Function 3 is the function that called Function 2.

Address. The starting address space for the memory that was allocated.

Length. The length of the block requested on the *alloc()* call

Time. The timestamp taken when the *alloc* call occurred.

Viewing MPX Memory Statistics

The **memstat** command displays the MPX's memory statistics. The statistics will tell you how memory is currently being used and help determine if memory problems exist, such as memory exhaustion. To view the MPX's memory statistics, enter

memstat

at the system prompt. A screen similar to the following will be displayed.

Summary of Memory Usage

<u>status</u>	<u>bytes</u>	<u>blocks</u>	<u>avg block</u>	<u>max block</u>
current				
free	4761672	64	74401	4719704
alloc	6429088	9114	705	-
cumulative				
alloc	24942880	148235	168	-
MPM total memory			: 16MB	

The fields displayed by the **memstat** command are described below.

status. The statistics appear in two groups: **current** and **cumulative**. The current status shows free and allocated memory. The cumulative status shows only allocated memory. Cumulative memory is the total amount of memory that has been allocated since the switch was started up. This value increases each time a memory allocation takes place. It can never decrease.

bytes. The number of bytes for free and allocated memory.

blocks. Block size is dynamic and depends upon memory usage and the amount of fragmentation.

avg block. The average block indicates the average size of all the memory blocks.

max block. The maximum block indicates the largest free memory block available. When this value drops to around 10K it usually indicates that the free memory is highly fragmented and probably near exhaustion.

MPM total memory. The total number of megabytes available in the MPX's memory.

Checking the Flash File System

The **fsck** command performs a file system check of flash memory, which consists of the flash file system. Image files are stored in flash memory and loaded into system memory when the switch boots up. It also provides diagnostics in the case of file corruption. To perform a file system check of flash memory, enter

```
fsck
```

at the system prompt. A screen similar to the following will be displayed.

```
Your bootroms support Flash File System Version 2 and greater.
```

```
Out of 16 file descriptors in use, 0 of these are opened on the /flash device.
```

```
Performing a file system check using manual mode. If a file is encountered  
with a potential problem, you may wish to consider preserving it for technical  
support analysis...
```

```
Flash file system check in progress...
```

```
Checking root file system... OK
```

```
Performing file consistency check...
```

```
Done.
```

```
There doesn't appear to be a system problem related to the Flash File  
system or kernel file system data structures. If you are experiencing  
problems with the flash file system, perhaps try using the "info",  
"systat", or "memstat" commands. They may indicate some other condition  
(such as low memory) which could prohibit correct operation of the  
file system.
```

If the **fsck** command finds a problem with the flash file system, a message will be displayed detailing the problems found and/or actions taken to correct those problems.

Checking the SIMM Files

Each logical file system (**/flash** and **/simm**) must be checked independently. If you have installed the 32 or 56 Mb SIMM upgrade and you want to check the SIMM's memory, enter

```
cd /simm
```

at the system prompt before you execute the **fsck** command.

Creating a New File System

The **newfs** command removes a complete flash file system and all files within it. It then creates a new flash file system, which is empty. You can use this command when you want to reload all files in the file system from a readily-accessible backup device or in the unlikely event that the flash file system becomes corrupted.

◆ Important Note◆

Before you execute the **newfs** command you should preserve your configuration file by saving it to another host.

To re-initialize the flash memory, enter

```
newfs
```

at the system prompt. The following screen will display.

```
You are about to destroy all files on file system /flash. If you  
are experiencing problems with the flash file system, you might  
want to use the "fsck" command to help determine where problems  
may exist.
```

```
Are you absolutely sure you want to strip the current file  
system and create a new one? (n)
```

Enter **y** to re-initialize the flash memory or **n** to cancel (the default is **n**). If you enter **y**, you will have to load new software into the switch.

◆ Warning ◆

Do not power-down the switch after running the **newfs** command until you reload your image and configuration files. If you do, you will have to reload the image files at the boot monitor prompt using the serial interface (e.g., ZMODEM), which can take several minutes.

You can then download new files via FTP or ZMODEM.

Creating a SIMM File System

If you have installed the 32 or 56 Mb SIMM upgrade and you want to create a new file system in the SIMM's memory, enter

```
cd /simm
```

at the system prompt before you execute the **newfs** command.

Configuring System Information

You can enter or modify a description of a switch, its location, and a contact person. Although this information is not required, you may find it helpful in managing the switch. To enter or modify the switch descriptions, perform the following steps.

1. At the system prompt, enter

```
syscfg
```

The current system information will appear with a prompt asking if you want to change any of the information; for example:

```
System Contact           : Usenet
System Name             : Testnet4
System Location         : Calabasas
System Description      : Marketing_testnet
Duplicate MAC Aging Timer : 0 (not configured)
Change any of the above {Y/N}? (N) :
```

If you enter **n**, the **syscfg** command will exit and no changes will be made (the default is **n**). If you enter **y**, the current system information will be displayed line by line. To keep the current value (shown in brackets) for a line, press **<Enter>**. To change a value, enter the new value and press **<Enter>**.

◆ Important Note ◆

Except for the **Duplicate MAC Aging Timer** field, all changes you make take place immediately.

If you entered **y**, something similar to the following will be displayed.

```
System Contact (Usenet) :
```

2. Enter the new system contact or just press **<Enter>** to accept the default. A screen similar to the following will be displayed.

```
System Name (no_name) :
```

3. Enter the new system name or just press **<Enter>** to accept the default. A screen similar to the following will be displayed.

```
System Location (Unset) :
```

4. Enter the new system location or just press **<Enter>** to accept the default. A screen similar to the following will be displayed.

```
System Description (DESCRIPTION NOT SET.) :
```

5. Enter the new system description or just press **<Enter>** to accept the default. A screen similar to the following will be displayed.

```
Duplicate Mac Aging Timer :
```

The **Duplicate MAC Aging Timer** indicates the time, in seconds, duplicate MACs remain in CAM if there is no traffic from those MACs. After this time, inactive MACs will age out of the CAM. You must reset the switch before this parameter takes effect. Duplicate MAC addresses will display as normal MAC addresses in other software commands, such as **fw** and **macinfo**. See *Duplicate MAC Address Support* on page 9-30 for further discussion.

6. Enter a new duplicate MAC aging timer value (the valid range is from 10 to 1000000) or just press **<Enter>** to accept the default.

Viewing CAM Information

The **camstat** command displays information and usage about the content addressable memory (CAM) on each switching module in the chassis. To view this CAM information, enter

```
camstat
```

at the system prompt. Something similar to the following will be displayed.

Slot	# of CAMs	Cfg Usage	Max Avail	Actual Usage
MPM	1	NA	NA	NA
2	4 (2 + 2)	0	3966	0
3	1 (1 + 0)	0	1008	0
4	1 (1 + 0)	0	1004	0
5	4 (2 + 2)	0	4093	0

The fields displayed by the **camstat** command are described below.

Slot. The slot number of the switching module for which CAM information is provided.

of CAMs. The number of CAM chips installed on the switching module.

Cfg Usage. The number of CAM entries this module is configured to support. By default a module will use the maximum amount of entries supported by on-board CAM. However, you can alter this default through the **camcfg** command (described in *Configuring CAM Distribution* on page 9-25) to make the most efficient use of the CAM distributed among all switching modules in the chassis. Up to 31.25 K of CAM is supported over all modules in an OmniSwitch/Router.

Max Avail. The number of CAM entries available. This number will be less than the number of CAM entries configured because some entries will be used by learned MAC addresses (shown in the **Actual Usage** column) and others are used internally by the OmniSwitch.

Actual Usage. The number of MAC addresses learned by the module in this slot.

◆ Note ◆

For CAM statistics for an entire chassis, use the **hdstat** command, which is described in Chapter 11, "Health Statistics."

Configuring CAM Distribution

CAM (Content Addressable Memory) on switching modules is used to look up the MAC address of endstations attached to the modules. You can use the **camstat** command to display each module's CAM usage. See *Viewing CAM Information* on page 9-24 for more information on the **camstat** command.

The Omni Switch/Router supports approximately 31.25 K of usable CAM among all the switching modules in a chassis. (A small amount of CAM memory is reserved by the Omni Switch/Router for its processing.)

When each switching module in a 9-slot chassis has 1 K of CAM, the 31.25 K limitation is not reached since only 8 K (assuming 8 switching modules) is used. However, when some switching modules use 4 K or 8 K of CAM the 31.25 K limitation could be reached quickly.

For example, if *all* the switching modules in a fully-loaded 9-slot chassis have 4 K CAMs you would exceed the 31.25 K limit. In this configuration, the Omni Switch/Router would subtract 256 K of available CAM memory from the first switching module to initialize and 512 K of available CAM memory from the last switching module to initialize. If you need to configure CAM usage use the **camcfg** command, which is described below.

◆ Important Note ◆

If you use a configuration file (e.g., **mpm.cfg**) from an OmniSwitch on an Omni Switch/Router, any CAM configuration settings will be ignored.

The **camcfg** command allows you to individually allocate CAM space to switching modules. This command configures the maximum entries a switching module may use, freeing up overall CAM space in the chassis so that some modules can use more of their on-board CAM. Follow these two additional rules:

- The CAM memory size for a switching module must be configured to at least one-half of the total memory available on the switching module. For example, if your switching module has 2 K of CAM memory, you must allocate at least 1 K of CAM to that switching module.
- The amount of CAM memory allocated for a switching module must be a whole-number multiple of 1024 (e.g., 1024, 2048, etc.).

Follow these steps to configure the number of CAM entries used by a switching module:

1. Enter **camcfg** followed by the slot number for the module that you want to configure. You can configure the CAM on switching modules only, not on the MPX. For example, to configure CAM for the module in slot 3, enter

```
camcfg 3
```

2. The system displays a prompt asking for the number of CAM entries to use for this module.

Enter maximum number of CAM entries for slot 3 (1024):

Enter the number of CAM entries to use for this module. The current value is listed in parentheses. The value you enter must be equal to or less than the total number of entries available on board this module. For example, you could not configure 2048 entries for a switching module with only 1K of CAM.

A message similar to the following will display:

**Slot 3 Configured to learn 256 MACs will round up to 256 MACs
This configuration will take effect only after system reboot**

3. The new CAM configuration will take effect after you reboot the system. For this reason, you may want to configure the CAM for all modules in this system. Reboot the system and check the updated CAM configurations through the **camstat** command.

Configuring the HRE-X Router Port

Various services in the switch use the HRE-X router port MAC registers. The registers are allocated as the services are loaded at startup. The **hrex** submenu contains five commands for use with the Hardware Routing Engines (HREs). The **hrexassign** command allows you to configure the switch so that registers are reserved for particular services. The **hrexdisplay** command allows you to view your current configuration. To display the **hrex** submenu, enter

```
hrex
```

at the system prompt. A screen similar to the following is displayed.

Command	HRE-X Management Menu
hrexassign	Assign an HRE-X router port MAC register to a service
hrexdisplay	Display HRE-X router port MAC register assignments
hrexutil	Display HRE-X Pseudo CAM and cache utilization
hrexhashopt	Optimize HRE-X Pseudo CAM hash function for current data
hrexhashdflt	Restore default HRE-X Pseudo CAM hash function

To view the current HRE-X configuration enter

```
hrexdisplay
```

at the system prompt. A screen similar to the following is displayed.

Reg	Configured	Actual
1	Any	Routing
2	Any	Unused
3	Any	Unused

The fields displayed by the **hrexdisplay** command are described below:

Reg. The number of the MAC registers.

Configured. The service type assigned to the register.

Actual. The service that is actively using the register.

To reserve a register for a particular service, you can assign the registers to the service. To assign the registers on the HRE-X router port, enter

```
hrexassign
```

at the system prompt. A screen similar to the following is displayed.

```
hrexassign <register number> <service type>
```

The **<register number>** is either 1, 2 or 3 referring to the MAC register. The **<service type>** parameter specifies the service configured to the registers. The service types are shown on the screen display are defined here.

any. This register is not reserved to a particular service.

routing. This register is assigned to standard routing.

cip. This register is assigned to Classical IP

m013. This register is assigned to Channelized DS-3 module (WSX-M013).

mpoa. This register is assigned to Multiprotocol Over ATM

vrrp. This register is assigned to Virtual Router Redundancy Protocol.

Configuring the HRE-X Router Port

For example, to assign register 3 to the Classical IP service enter

```
hrexassign 3 cip
```

at the system prompt. A screen similar to the following is displayed.

```
HRE-X RPM 3 configured for "CIP"; reboot to make effective.
```

As indicated on the screen, the register assignment will not take effect until the switch is rebooted. If you use the **hrexdisplay** command after making a the register assignment shown in the above example, a screen similar to the following is displayed.

Reg	Configured	Actual
1	Any	Routing
2	Any	Unused
3	CIP	Routing

Configuration changed since last reboot.

This indicates that register 3 is assigned to the CIP service but is actually using the Routing service. Also, the message at the bottom of the table indicates that the HRE-X configuration has changed since the last reboot of the switch. After a reboot, the **hrexdisplay** command will display the following screen.

Reg	Configured	Actual
1	Any	Routing
2	Any	Unused
3	Routing	Routing

Configuring and Displaying the HRE-X Hash Table

The HRE-Xs use a hardware implemented hash table to route packets for transmission. The switch employs a default hash function that works well in a broad range of data environments. In rare cases, you may want to change the hash table configuration to optimize it for your particular data flow. This should be done with care because the data population will change over time. A hash function that works well for one set of data may not work as well for another. Also, note that optimizing the hash function will cause all of the current entries in the HRE-X to be cleared and then relearned; therefore, this should be done with extreme caution.

Two HRE-X commands are used to optimize the hash function. They are the **hrexutil** and the **hrexhashopt** commands. The **hrexutil** command displays the current utilization of the hash table. To view the HRE-X Utilization table, enter

```
hrexutil
```

at the system prompt. A screen similar to the following is displayed.

```
HRE-X Utilization
-----
Hash - Total: 65536 Free: 65528
Collisions - Total: 131072 Free: 131069
Cache - Total: 40960 Free: 40949
Collision Length - Max: 3 Avg: 1
```

The fields displayed by the **hrexutil** command are described below:

Hash. The number of entries in the hash table.

Total. The total number of units available.

Free. The number of units that are not yet used.

Collisions. The number of entries that have hashed to the same index in the hash table.

Cache. The number of modifications required to route a packet.

Collision Length. The length of the longest (**Max**) collision list and the average length (**Avg**) of the collision lists.

The **hrexhashopt** command causes the switch to compute an optimized hash function based on the data currently in the HRE-X. This function is saved in the configuration file so it will be present after a reboot.

To use the **hrexhashopt** command, enter

```
hrexhashopt
```

at the system prompt. The screen does not display a confirmation message after this command. You can verify optimization by observing the changes in the HRE-X Utilization. After using **hrexhashopt**, the maximum and average collision lengths should be reduced as shown in the HRE-X Utilization table shown above. If they are not, you should consider returning to the default hash function by using the **hrexhashdfit** command.

To use the **hrexhashdfit** command, enter

```
hrexhashdfit
```

at the system prompt. The screen does not display a confirmation message after this command. The **hrexhashdfit** command will return the hash function back to the default value.

Duplicate MAC Address Support

When the switch sees the same MAC address sending traffic on a different switch port (a Duplicate MAC Address), it assumes the original network device moved. The switch sends a trap notifying network management of this station move event. It sends one trap for a device move within the same Group and another trap for a device move outside of the home Group.

A station move trap is normally sent after an actual station move. However, certain network configurations assign the same MAC address to different network devices (physical and virtual) as standard practice. In these situations, the duplicate MAC address appears as a station move when it is really a normal occurrence in these network configurations. These network configurations that use the same MAC address for different devices include:

- LAN Emulation under Cisco routers. Cisco routers use the same MAC address for each LAN Emulation Client (LEC). In LAN Emulation, each ELAN needs to be treated as a separate LAN and should therefore have a separate MAC address.
- IBM Front End Processor (FEP). Many IBM FEPs use the same MAC address assigned to the connecting devices for the purpose of redundancy.
- DECnet networks. The DECnet protocol assigns the special MAC address, AA000400XXYY (XXYY is an internal protocol ID) to each DECnet station or routing device regardless of the number of physical interfaces.

Initially, duplicate MAC addresses in these special situations may be no more of a problem than extra traps being sent for an event (station move) that did not really happen. However, when a large number of these network devices send the same MAC address out the same port, flooding can occur and the switch will eventually shut the port down.

To prevent a port from being shut down, the switch needs some way of knowing the duplicate MAC addresses originating from the port are not an error condition.

The switch will treat duplicate MAC addresses as separate addresses as long as they are learned from a different Group as the original MAC. Each duplicate MAC address will use one entry in the CAM. Up to 32 duplications of the same MAC address are supported. Duplicate MAC addresses learned from virtual ports within the same Group are treated as station moves and will generate corresponding traps. If the MAC address moves from one VLAN to another VLAN within the same Group, the switch will not treat the MAC addresses as separate.

If your network supports duplicate MAC addresses, there may be a significant performance impact due to the following reasons:

- A MAC address is usually stored only in the CAM of the switching module where its destination address is located. If duplicate MAC addresses are treated as separate addresses, then the same MAC address may have to be stored in the CAM of multiple switching modules, not just the module that originally learned the address.
- Every duplicate MAC address becomes a CAM table entry, so there will be less room in the CAM for other entries to be learned. Since up to 32 duplications of a single MAC address are possible, this CAM can become crowded with these duplicate entries.

You can reduce the impact of a crowded CAM by configuring the **Duplicate MAC Aging Timer** in the **syscfg** command, which is described in *Configuring System Information* on page 9-23. This timer allows you to age out Duplicate MAC CAM entries from devices that are inactive for the time period you specify.

- Extra search time will be required for each lookup of the same MAC address since it is treated as a separate entry in the CAM.

In addition to these performance impacts, you will lose the tracking of legitimate station moves. No traps will be sent for Duplicate MAC addresses that appear in different Groups.

Multicast Claiming

Multicast claiming can be enabled for networks with heavy multicast traffic. When enabled, multicast claiming frees the MPX from processing multicast packets by off-loading this traffic to the switching modules. When multicast claiming is enabled, the switch “claims” destination multicast addresses and places them in the CAMs of all switching modules in the switch.

You can enable multicast claiming by adding the following line to the **mpx.cmd** file:

```
bsiLearnMcPkt=1
```

You can use the **edit** command to make this change. (See Chapter 7, “Managing Files,” for instructions on using the **edit** command.) You will need to reboot the switch for this parameter to take effect. Multicast claiming can later be disabled by changing the setting for this parameter to zero (0), as follows:

```
bsiLearnMcPkt=0
```

An alternative method for managing multicast traffic is through the use of Multicast VLANs. See Chapter 27, “Managing AutoTracker” and Chapter 28, “Managing Multicast VLANs” for further information.

Disabling Flood Limits

Two UI commands are available for controlling flood limits for individual ports and Groups. The **modvp** command (described in Chapter 24, “Managing Groups and Ports”) allows you to control the flood limits for a specific port. The **flc** command (described in Chapter 22, “Configuring Bridging Parameters”) allows you to configure flood limits for all ports in a group.

You can also disable flood limits on a switch-wide basis by adding the following line to the **mpx.cmd** file:

```
disableFloodLimiting=1
```

You can use the **edit** command to make this change. See Chapter 11, “Managing Files,” for instructions on using the **edit** command. You will need to reboot the switch for this parameter to take effect.

Saving Configurations

Under normal conditions, configurations you make using the UI are written into cache and automatically saved into the switch's flash memory. In this case, it is not necessary to issue a special command to save your configurations. When you use the UI to enter multiple configurations, periodically the switch will display the following message.

File system compaction in progress . . .

This message indicates that the switch is compacting data in the cache buffer before writing it into the mpm.cnf file. This message normally disappears after a few seconds.

◆ Warning ◆

It is highly recommended that you use the default setting and allow the switch's save function to operate automatically.

You can change the switch's save function so that the cache is not saved automatically by executing the **cacheconfig** command. To turn off the switch's automatic save function, enter

cacheconfig on

at the system prompt. The following message will display.

Cache Configuration is now on

◆ Warning ◆

Any configurations you enter before executing the **saveconfig** command will not be saved in case of system failure or reboot.

Once **cacheconfig** is implemented, you must use the **saveconfig** command to manually synchronize your configurations into flash memory. When you execute the **saveconfig** command at the system prompt, the following message will display.

File system compaction in progress . . .

The UI does not indicate when the **cacheconfig** function is in operation. However, if you attempt a reboot the following message will display if you are in the cache configuration mode.

**!!!Warning!!! You are in the cache configuration mode.
Please enter 'n'/N' to the following confirm prompt.
Then enter the UI command "saveconfig", or
enter the CLI command "dump configuration cache" to
save the current configuration to mpm.cnf in the flash.**

Otherwise, all/some your configuration changes will be lost!

Confirm? (n) :

This message gives you the opportunity to execute the **saveconfig** command prior to the reboot.

Saving Configurations

To determine whether you are in the cache configuration mode, enter the **cacheconfig** command. If cache config is operational the following message will display one of the following messages.

Cache Configuration is currently on.

or

Cache Configuration is currently off.

To turn off the cache configuration mode, enter the following command at the system prompt.

cacheconfig off

The following message will display.

**File system compaction in progress . . .
Cache Configuration is now off**

10 Switch Logging

Logging Overview

Whether you are troubleshooting, configuring, or simply monitoring the switch, you may find it useful to view a history of various switch activities. The Logging submenu contains a list of commands for viewing and configuring logging on the system. To enter the logging submenu, enter

logging

at the system prompt. Enter a question mark (?) and then press **<Enter>** to display the following list of commands:

<u>Command</u>	<u>Logging Menu</u>
syslog	Change the syslog parameters (not part of Switch Logging feature).
swlogc	Configure Switch Logging source/destination mapping and priority levels.
cmdlog	Show UI Command entries in the mpm.log file
conlog	Show Connection entries (logins/logouts) entries in the mpm.log file
caplog	Show Screen Capture entries in the mpm.log file.
debuglog	Show Debug message entries in the mpm.log file
seclog	Display Secure Access log file entries.

Commands in the submenu are described here.

System Log Messages

The **syslog** command is used to configure how system log messages, like diagnostic and error messages, are handled on the switch. See *Configuring the Syslog Parameters* on page 10-2.

Switch Logging Parameters

The **swlogc** and remaining commands in the submenu are part of the Switch Logging feature, which is a separate logging mechanism. The **swlogc** command is used for configuring the logging parameters of various switch activities such as FTP and Telnet, and is described in *Configuring Switch Logging* on page 10-6.

The other commands listed in the submenu above are support commands for Switch Logging.

- **cmdlog** command—displays the UI command entries in the mpm.log file, which is one of the possible destinations for Switch Logging data. See *Displaying the Command History Entries in the MPM Log* on page 10-9.
- **conlog** command—displays the connection entries in the mpm.log file. See *Displaying the Connection Entries in the MPM Log* on page 10-10.
- **caplog** command—displays the screen capture entries in the mpm.log file. See *Displaying Screen (Console) Capture Entries in the MPM Log* on page 10-11.
- **debuglog** command—shows the debug entries in the mpm.log file. See *Displaying Debug Entries in the MPM Log* on page 10-13.
- **seclog** command—shows the Secure Access violation event entries in the mpm.log file. See *Displaying Secure Access Entries in the MPM Log* on page 10-13.

Configuring the Syslog Parameters

Syslog messages are messages generated by individual processes in the switch. These messages contain information for conditions that range from debugging to emergency error conditions.

The **syslog** command allows you to control how these messages will be handled. You can designate what kinds of messages you will see and where the messages will be sent. This syslog implementation is compatible with the standard BSD UNIX implementation for syslog services.

To see the current syslog configuration, enter

```
syslog
```

at the system prompt. A screen similar to the following will be displayed.

```
SYSLOG current configuration:
```

```
1) Log host           - UNDEFINED
2) Log host IP       -
3) Syslog port (514) - 514
4) Default facility code - local0
  41) Override internals - no
5) Default priority mask - emerg
  51) Override internals - no
  52) Display internals - no
6) Console logging   - yes
7) Log Task ID       - yes
  71) Use Task Name   - no
8) Message tag       - switch
```

```
(save/quit/cancel)
```

```
:
```

Select the number of the item you want to change. To change any of the values on the previous page, enter the line number, followed by an equal sign (=), and then the new value. For example, to turn off console logging, enter:

```
6=no
```

The question mark (?) option refreshes the screen. To update the values you have changed, enter **save**. If you do not want to save the changes enter **quit** or **cancel**, or press **Ctrl-D**.

The parameters displayed by the **syslog** command are described below.

Log host

The name of the host where you want the syslog messages sent. The Domain Name Server (DNS) must be configured for this to work. Use the **res** command to configure the DNS. (The **res** command is described in Chapter 14, "RMON and DNS Resolver.")

Log host IP

The IP address of the host where you want the syslog messages sent. If the IP address and the Log host name disagree, the IP address takes precedence.

Syslog port (514)

The port to which the syslog messages will be sent on the specified host. Port 514 is the normal port number used and is the default.

Default facility code

The facility code is used to identify which sub-system generated the syslog message. Note that this code is used only as a default for tasks that do not have a facility code. See the table below for a list of the facility codes. The default is **local0**.

Syslog Facility Codes

Facility	Source
LOG_KERN	Messages generated by the kernel
LOG_USER	Message generated by random user processes
LOG_MAIL	The mail system
LOG_DAEMON	System daemons
LOG_AUTH	The authorization system
LOG_LPR	The line printer spooling system
LOG_NEWS	Reserved for the USENET system
LOG_UUCP	Reserved for the UUCP system
LOG_CRON	The cron/at facility
LOG_LOCAL0-7	Reserved for local use

Override internals

This setting will force all syslog messages to use the default facility code specified in **Default facility code** instead of their own predefined facility codes.

Default priority mask

The mask for the priority code. Indicates the type of syslog message. Note that this mask is used only as a default for tasks that do not have a priority code. Priority codes for syslog messages are usually hardcoded. The following table is a list of priority codes.

Syslog Priority Codes

Level	Value	Meaning
LOG_EMERG	0	FATAL system event
LOG_ALERT	1	FATAL subsystem event
LOG_CRIT	2	Problem, subsystem unstable
LOG_ERR	3	Problem, bad event, recoverable
LOG_WARNING	4	Unexpected, non-fatal event
LOG_NOTICE	5	normal but significant condition
LOG_INFO	6	info
LOG_DEBUG	7	Internal debug messages

Override internals

This field will force all syslog messages to use the default priority mask specified instead of their own predefined priority masks.

Display internals

This field allows the user to display the task log level. Enter **52=yes** to display the sub-menu below. If, for example, you wanted to change the priority mask **CM via kern** from “warn” to “alert,” you would enter **4=alert**. Note that this change will take place immediately and you do not need to enter **save** for it to take effect. Type **save**, **quit**, or **cancel** and then press **<Enter>** to return to the main **syslog** menu.

Internal task syslog configuration:
(NOTE: changes take effect immediately and are NOT saved across reboots!)

- 0) PPM via kern - alert
- 1) LPM via kern - alert
- 2) VPM via kern - alert
- 3) SNMP via kern - alert
- 4) CM via kern - warn
- 5) ATMmgr via kern - alert
- 6) atmLANE via kern - alert
- 7) Q93bif via kern - alert
- 8) ILMlif via kern - alert
- 9) SSI0 via kern - alert
- 10) atmSNMP via kern - alert

Console logging

Determines whether or not you want to see syslog messages on your console (terminal). If set to yes, the messages will be displayed on either an **ASCII** terminal connected to the console port or via a Telnet session.

Log Task ID

Determines whether or not you want to see the task ID that can be included in the syslog message.

Use Task Name

This allows the user to display descriptive task names for syslog messages (see the **Display internals** sub-menu above) instead of numeric codes.

Message tag

Text of up to 10 characters that is added to every message leaving the switch. It is useful when multiple switches send messages to the same host.

Configuring Switch Logging

Switch logging is a feature that allows you to activate and configure the logging of various types of switch information. Once you activate logging for a specific facility through the switch logging command, you may also decide whether the log output should display on the console, be saved to a file, or be both displayed and saved to a file. To enter the switch logging submenu, enter

swlogc

at the system prompt. A screen similar to the following displays:

CONFIGURATION MENU FOR SWITCH LOGGING

```
1) Security Logging                : Disabled
   11) Output to File              : Yes
   12) Output to Console           : No
2) FTP Logging                    : Disabled
   21) Output to File              : Yes
   22) Output to Console           : No
3) Flash File Logging             : Disabled
   31) Output to Console           : Yes
4) Screen Capture                 : Disabled
   41) Output to File              : Yes
5) Console Event Logging          : Disabled
   51) Output to File              : Yes
   52) Output to Console           : No
6) User Interface Logging        : Disabled
   61) Output to File              : Yes
   62) Output to Console           : No
7) Telnet Logging                 : Disabled
   71) Output to File              : Yes
   72) Output to Console           : No
8) Log File (mpm.log) Size        : 20000 bytes
9) Return Logging to Default Configuration : No
```

Command {Item/ Item=Value/ ?/ Help/ Quit/ Cancel/ Save} (Redraw) :

The logging types are described here:

1) Security Logging

Enabling security logging allows you to view all security violations that occur within the switch. Set to **enable** to activate logging for any security violations that occur within the switch. Set to **disable** to de-activate logging for security violations.

◆ Note ◆

Security Logging must be enabled in order to display the Secure Switch Access violations log (**seclog**).

2) FTP Logging

FTP Session Events is a record of all FTP (File Transfer Protocol) activities since logging was activated. Once you enable FTP Logging by entering **2=enable**, you may view it through the **conlog** command (described in *Displaying the Connection Entries in the MPM Log* on page 10-10). To disable FTP Session Events logging, enter **2=disable**.

3) Flash File Logging

Flash file logging records debug information from the code that manages the switch logging feature itself (previously called “flash file system logging”). To enable flash file logging, enter **3=enable**. To disable flash file logging, enter **3=disable**. Flash file logging messages cannot be saved in the `mpm.log` file, but flash file logging messages may be displayed on the console by entering **31=yes**. To disable sending flash file logging messages to the console, enter **31=no**.

4) Screen Capture

Screen logging captures screen text for logging. To enable screen logging, enter **4=enable**. To disable screen logging, enter **4=disable**. Note that since screen text already goes to the screen, logging output to the screen is not permitted. If you want to display the screen capture entries for all logged users, use the **caplog** command (for more information, see *Displaying Screen (Console) Capture Entries in the MPM Log* on page 10-11).

◆ Note ◆

The screen capture feature has not yet been implemented.

5) Console Event Logging

Console Session Events is a record of all console login activities in the switch, including user names, and connection times. Once you enable Console Event logging by entering **5=enable**, you may view it through the **conlog** command (described in *Displaying the Connection Entries in the MPM Log* on page 10-10). To disable logging for Console Events, enter **5=disable**. Note that logging output to the console is not permitted.

6) User Interface Logging

User Interface Logging is executed on the switch since the UI log was activated. Once you enable UI logging by entering **6=enable**, you may view it through the **cmdlog** command (described in *Displaying the Command History Entries in the MPM Log* on page 10-9). To disable logging for the UI, enter **6=disable**.

7) Telnet Logging

Telnet Logging is a record of all Telnet activities since Telnet logging was activated. Once you enable Telnet logging by entering **7=enable**, you may view it through the **conlog** command (described in *Displaying the Connection Entries in the MPM Log* on page 10-10). To disable logging for Telnet, enter **7=disable**.

8) Log File Size

Use this parameter to set the `mpm.log` file size. The default is 20,000 bytes. The maximum number of bytes is dependent upon the available flash in your system. If you set a file that is too large, the command will tell you the maximum allowed size. (This is half of the remaining free space in your flash file system.) The minimum file size is 3,240 bytes.

9) Return Logging to Default Configuration

Use this parameter to return all of the switch logging options to their default values. Enter **9=yes** to reset the configuration at reboot. To keep the same logging configuration at the next reboot, make sure this parameter is set to **no**.

In addition to enabling or disabling each type of logging, you can also specify whether to output the log to a file or to the console:

Output to File

Set to **yes (y)** to store the log messages in the mpm.log file. Set to **no (n)** to disable sending log messages to this file. (This option is not available for flash file logging or screen capture.)

Output to Console

Set to **yes** to display the log messages on the console screen. Set to **no** to disable the screen as an output device for Security Logging.

Displaying the Command History Entries in the MPM Log

The **cmdlog** command displays a list commands executed since User Interface (UI) facility logging was activated by the **swlogc** command (described in *Configuring Switch Logging* on page 10-6). To display this data, enter

cmdlog

at the system prompt. The following is a sample display.

User	Line	Time	User Input
admin	198.206.187.113	08/14/00 16:42	cmdlog
admin	198.206.187.113	08/14/00 16:42	xlat
admin	198.206.187.113	08/14/00 16:43	conlog
admin	console	08/15/00 10:28	logging
admin	console	08/15/00 10:28	?
admin	198.206.187.113	08/15/00 14:03	taskstat
admin	198.206.187.113	08/15/00 14:05	taskstat

The fields displayed by the **cmdlog** command are described below.

User. The login name of the user who executed the command.

Line. The login type of the user who executed the command. If, for example, the user was connected through the console port, “console” will be displayed. If the user was connected through Telnet, on the other hand, then the IP address of that user will be displayed.

Time. The time that the command was executed.

User Input. The actual text (up to 32 characters) that the user entered at the system prompt.

◆ **Note** ◆

If you just want to display the commands executed during the current session you can use the **history** command, which is described in Chapter 4, “The User Interface.”

Displaying the Connection Entries in the MPM Log

The **conlog** command displays a list of connections made since console event, FTP, or Telnet logging was activated by the **swlogc** command (described in *Configuring Switch Logging* on page 10-6). To display this data, enter

```
conlog
```

at the system prompt. A screen similar to the following will be displayed.

User	Line	Peer	Start	Finish
-----	-----	-----	-----	-----
admin	Telnet	198.206.187.113	08/14/00 09:47 -	09:47 (00:00)
admin	Telnet	198.206.187.113	08/20/00 09:47 -	09:53 (00:05)
admin	Telnet	198.206.187.113	08/20/00 09:55 -	10:00 (00:05)
admin	console		08/20/00 10:35	logged in (00:27)
admin	Telnet	198.206.187.113	08/20/00 11:02	logged in (00:00)

The fields displayed by the **conlog** command are described below.

User. The name of the user who made the connection to the switch.

Line. The login type of connection to the switch (e.g., a Telnet or console port connection).

Peer. If the user was connected through Telnet, then the IP address of the user will be displayed. If the user was connected through the console port, then this field will be blank.

Start. The time that the connection started.

Finish. Displays the time the connection terminated or **logged in** for sessions that are still current. The value in parenthesis is the duration of the session, in minutes.

Displaying Screen (Console) Capture Entries in the MPM Log

The **caplog** command displays the screen capture entries in the mpm.log file. (*Note: This feature is not yet implemented.*) In order to view screen capture entries through this command, you must first enable the Screen Capture log facility through the **swlogc** command (see *Configuring Switch Logging* on page 10-6). To display screen capture entries in the log, enter

```
caplog
```

at the system prompt. A screen similar to the following will be displayed.

```

1) Console
2) Modem
3) Telnet (0)
4) Telnet (1)
5) Telnet (2)
6) Telnet (3)
    select ?

```

Select which user's screen entries you would like to view by entering the user's line number at the prompt. For example, if you enter **1** at the **select ?** prompt, a screen similar to the following displays:

```

=====Start Screen Capture Display for Console=====
/ % systat

System Uptime                : 0 days, 01:01:47.01
MPM Transmit Overruns       : 0
MPM Receive Overruns        : 0
MPM total memory             : 18548968 bytes
MPM CPU Utilization (5 sec)  : 3 % ( 0% kernel 1% task 97% idle)
MPM CPU Utilization (60 sec) : 4% ( 0% intr 0% kernel 2% task 96% idle)\
Power Supply 1 State         : OK
Power Supply 2 State         : Not Present
Temperature                  : 32.00c 89.60f
Temperature Sensor           : OF - Under Threshold
Temperature Alarm Masking    : Disabled
=====End Screen Capture Display for Console=====

```

The options displayed by the **caplog** command are described below.

- 1) **Console**. Displays screen capture entries for the user logged in from the console.
- 2) **Modem**. Displays screen capture entries for the user logged in from the modem.
- 3) **Telnet (0)**. Displays screen capture entries for the user logged in from the first telnet session.

Displaying Screen (Console) Capture Entries in the MPM Log

- 4) **Telnet (1)**. Displays screen capture entries for the user logged in from the second telnet session.
- 5) **Telnet (2)**. Displays screen capture entries for the user logged in from the third telnet session.
- 6) **Telnet (3)**. Displays screen capture entries for the user logged in from the fourth telnet session.

Displaying Debug Entries in the MPM Log

The **debuglog** command displays the debug entries in the mpm.log file. (*Note: Currently there are no facilities using debugging.*) Below is a sample display of the **debuglog** command.

Task Name	Time	Debug Message
tUdpRelay	14:33:36	Undersized DHCP req rcvd; discarding

The fields displayed by the **debuglog** command are described here.

Task Name. The task that generated the debug message.

Time. The time the message was generated by the task.

Debug Message. Information relevant to debugging.

Displaying Secure Access Entries in the MPM Log

The **seclog** command displays the secure access violation event entries in the mpm.log file. To display this data, enter

seclog

at the system prompt. A screen similar to the following will be displayed.

Secure Access Violations Log

Time	Protocol	Source IP	Attempts	Slot/ Intf	Elapsed Time (secs)
12:49:02	FTP	172.23.8.801	1	5/1	23
03:15:34	Telnet	198.20.2.101	10	2/3	240

Descriptions of the fields are as follows:

Time. The first time the access violation occurred.

Protocol. The IP protocol for which the violation occurred.

Source IP. The source IP address of the unauthorized user.

Attempts. The number of access attempts made by this user within the sample period (5 minutes).

Slot/Intf. The physical port that received the unauthorized user information.

Elapsed Time (secs). The duration (in seconds) from the first unauthorized access to the end of the sampling period (5 minutes). Secure access violations will take 5 minutes to display in the log file.

11 Health Statistics

The health statistics feature monitors the consumable resources of a switch, and provides a single integrated source for Network Management Software (NMS), such as X-Vision, to use in obtaining statistics on switch performance. With the health statistics, the user can set specific threshold levels for consumable resources in the switch. Such resources include bandwidth capacity, CAM and CPU usage, and RAM memory usage. If a threshold for a particular resource is exceeded, a notification is sent to the NMS via an SNMP trap.

◆ Important ◆

You must configure your NMS to accept traps from the monitored switch. X-Vision allows you to set which network management stations receive traps. For more information, see the X-Vision online help.

The health statistics software monitors the resource utilization levels and thresholds of a switch, and at fixed intervals collects the current values for each resource being monitored. After obtaining the statistics, the health statistics software checks to see if any rising or falling threshold crossings occurred since its last poll by comparing the current poll data with the previous poll data. If a threshold crossing has occurred, a trap is sent to NMS (such as X-Vision), allowing the system administrator to pinpoint possible performance issues.

Through the UI (user interface), threshold levels can be set, the sampling interval can be changed, and statistics (for a switch, module, or port) can be viewed or cleared.

The Health Statistics Management Menu

To access the Health menu, log on to a switch via a Telnet or console session, and type the following command:

```
health
```

If the session is in terse mode, you will need to type `?` to see the menu. If you are in verbose mode, the following screen is displayed:

Command	Health Menu
hdcfg	Set or view parameters
hdstat	View device-level statistics
hmstat	View module-level statistics
hpstat	View port-level statistics
hreset	Reset health statistics

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

```
/System/Health %
```

The **hdcfg** command allows you to set global thresholds for the switch. The **hdstat**, **hmstat**, **hpstat** commands allow you to view the statistics on a switch, module, or port level, respectively. The **hreset** command resets the statistics for this switch.

Setting Resource Thresholds

The health statistics software operates by monitoring set threshold levels on consumable resources. When a resource exceeds a set level, a trap is generated and sent. These threshold levels are set for the entire switch (or device) by using the **hdcfg** command. To set the threshold level for a switch's consumable resources, enter the **hdcfg** command at the system prompt. The following screen appears:

Device-level Resource Monitoring Configuration

- 1) **Set Bandwidth Thresholds** :
- 2) **Set Miscellaneous Thresholds** :
- 3) **Set Sampling Interval** :

There are three sets of resources that are configurable:

- **Bandwidth thresholds.** These settings allow you to set a percentage of available bandwidth for received traffic, sent traffic, and the backplane. For more information on setting bandwidth thresholds, see *Setting Bandwidth Thresholds* on page 11-3.
- **Miscellaneous thresholds.** These settings allow to set a percentage for memory usage, VCC usage, virtual port usage, and temperature. For more information on setting miscellaneous thresholds, see *Setting Miscellaneous Thresholds* on page 11-4.
- **Sampling interval.** The sampling interval is the number of seconds between health statistics checks. For information on how to set the sampling interval, see *Setting the Sampling Interval* on page 11-6.

Setting Bandwidth Thresholds

Bandwidth is a measure of the amount of traffic a switch can handle for receiving, sending, and on the backplane. The health statistics allow you to set a percentage of available bandwidth, at which an SNMP trap is generated to alert the network administrator that the threshold has been exceeded. To set the threshold levels for switch bandwidth:

1. Enter **health** at a system prompt. The health menu (described above) displays.
2. Enter a **1** at the health menu prompt. The following menu displays:

Bandwidth Resource Monitoring Configuration

```

1) Receive Threshold      : 80
2) Transmit/Receive Threshold : 80
3) Backplane Threshold    : 80

```

3. Threshold values are measured as a percentage of the total capacity of the resource. To change a threshold or sampling interval value, type the index for the field, followed by an equals sign, then the new value. For example, to change the **Receive Threshold** to 50 percent, you would type the following at the prompt:

```
1=50
```

The Receive Threshold would now be set to 50 percent of its total capacity (bandwidth).

4. When you have finished entering the new values, you must enter **save** to keep the new configuration settings.

◆ Note ◆

Changing a threshold value sets the value for all levels of the switch (switch, module, and port). You cannot set different threshold values for each level.

Below is a description of the fields in the **hdcfg** command menu. The default for all monitored resources is eighty (80) percent of the maximum capacity of the resource.

Receive Threshold

The receive threshold sets a percentage of total bandwidth of the switch, module, or port. When the amount of received data exceeds this percentage, an SNMP trap is sent.

Transmit/Receive Threshold

The transmit/receive threshold sets a percentage of the total bandwidth of the switch, module, or port. When the amount of transmitted and received data exceeds this percentage, an SNMP trap is sent.

Backplane Threshold

The backplane threshold sets a percentage of total backplane bandwidth of the switch, module, or port. When backplane usage exceeds this percentage, an SNMP trap is sent.

◆ **Note** ◆

When “U-turn” switching (i.e., data enters a module port and is transmitted from a port on the same module) is employed, the backplane threshold reading will not be correct. Switched frames are not transmitted over the backplane but are counted by health statistics, causing the backplane percentage reading to be higher than it should be.

Setting Miscellaneous Thresholds

The miscellaneous thresholds cover consumable resources such as memory, VCCs, temperature, and virtual ports. The health statistics allow you to set a percentage of the available resource, at which an SNMP trap is generated to alert the network administrator that the threshold has been exceeded. To set the threshold levels for switch bandwidth:

1. Enter **health** at a system prompt. The health menu (described above) displays.
2. Enter a **2** at the health menu prompt. The following menu displays:

Miscellaneous Resource Monitoring Configuration

1) CAM Threshold	: 80
2) CPU Threshold	: 80
3) Memory Threshold	: 80
4) VCC Threshold	: 80
5) Temperature Threshold	: 80
6) Virtual Port Threshold	: 80

3. Threshold values are measured as a percentage of the total capacity of the resource. To change a threshold or sampling interval value, type the index for the field, followed by an equals sign, then the new value. For example, to change the **CAM Threshold** to 50 percent, you would type the following at the prompt:

1=50

The CAM Threshold would now be set to 50 percent of its total capacity (memory).

4. When you have finished entering the new values, you must enter **save** to keep the new configuration settings.

◆ **Note** ◆

Changing a threshold value sets the value for all levels of the switch (switch, module, and port). You cannot set different threshold values for each level.

CAM Threshold (MPM/HRE or NI)

The CAM threshold sets a percentage of the total amount of space available for storing the cache tables. Cache tables maintain associations between received MAC addresses and the ports they were received on. For the switch level, the CAM threshold separately monitors the MPX and the HRE-X daughtercard (if it is installed) CAM tables. For the module level, it monitors the switching module CAM tables. CAM thresholds are not available on the port level.

When this percentage is exceeded, an SNMP trap is sent.

CPU Threshold

The CPU threshold sets a percentage of the total amount of processing ability for the MPX. When the CPU usage exceeds this percentage, an SNMP trap is sent. The CPU threshold is only used for the switch level.

Memory Threshold

The memory threshold sets a percentage of the total amount to MPX RAM memory for the switch. When RAM usage exceeds this percentage, an SNMP trap is sent. The memory threshold is only used for the switch level.

VCC Threshold

This value is a number set as a percent. VCC Threshold is equal to the total number of active VCCs divided by the switch VCC capacity. When this value is exceeded, an SNMP trap is sent.

Temperature Threshold

This threshold sets the number of degrees for the switch at which an SNMP trap is sent. This threshold is measured in degrees Celsius. The range is from 0 to 100.

Virtual Port Threshold

This threshold sets a percentage of the total number of available virtual ports for the switch. When the set percentage of available virtual ports is exceeded, an SNMP trap is sent.

Setting the Sampling Interval

The sampling interval is the time interval between polls of the switch's consumable resources to see if it is performing within the set thresholds. To set the amount of time between polls:

1. Enter **health** at a system prompt. The health menu (described above) displays.
2. Enter a **3** at the health menu prompt. The following menu displays:

Resource Monitoring Interval Configuration

1) Sampling Interval : 5

3. To change the sampling interval, enter a 1, and equal sign, and the new interval in seconds. For example, to change the sampling interval to 4 seconds, you would enter the following:

1=4

4. When you have finished entering the new value, you must enter **save** to keep the new configuration setting.

Sampling Interval

This sets the number of seconds between internal polling intervals. The health statistics compares the current poll statistics with the last poll statistics to determine whether or not to send a trap. The default for the **Sampling Interval** is five (5) seconds.

View Switch-Level Statistics

To view the statistics for the entire switch, enter the **hdstat** command at a system prompt. The following table is displayed:

Device Resources	Limit	Curr	1 Min Avg	1 Hr Avg	1 Hr Max
Receive	80	00	00	00	00
Transmit/Receive	80	00	00	00	00
Backplane	80	01	01	01	01
CAM [MPM]	80	00	00	00	00
CAM [HRE]	80	00	00	00	00
CPU	80	93*	13	13	22
Memory	80	50	50	50	50
Temperature	45	44	44	44	44
Virtual Ports	80	11	11	11	11

/System/Health %

Statistics are displayed as percentages of the total resource capacity, and represent data taken from the last sampling interval. If a threshold for a resource was exceeded, then that statistic is marked with an asterisk (*).

◆ Important Note ◆

The **hdstat** command displays CAM usage for the entire chassis. To see CAM usage for switching modules only, use the **camstat** command as described in Chapter 9, "Switch Wide Parameters."

For field descriptions of the device resources column, see *Setting Bandwidth Thresholds* on page 11-3 and *Setting Miscellaneous Thresholds* on page 11-4 above.

◆ **Note** ◆

When calculating percentages, the health statistics cannot display less than one percent. If a single packet is sent through a port, for example, the receive resource usage is represented as one percent.

The following section describes the statistics displayed using the **hdstat** command.

Limit

The set threshold for this resource. You can set the resource levels using the **hdcfg** command. See *Setting Resource Thresholds* on page 11-2 for specific procedures.

Current

The current resource usage. This number is a percentage of the total resource capacity.

1 Minute Average

The average percent of resource use for the last sixty seconds.

1 Hour Average

The average percent of resource use for the last sixty minutes.

1 Hour Maximum

The maximum percent of resource use for the last sixty minutes.

View Module-Level Statistics

To view module level statistics, type the **hmstat** command at a system prompt followed by the slot number. For example, to view the statistics for a module in slot three, type the following:

```
hmstat 3
```

The following screen is displayed:

Slot 3 Resources	Limit	Curr	1 Min Avg	1 Hr Avg	1 Hr Max
Receive	80	00	00	00	00
Transmit/Receive	80	00	00	00	00
Backplane	80	95*	00	00	00
CAM	80	00	00	00	00

/System/Health %

Statistics are displayed as percentages of the total resource capacity, and represent data taken from the last sampling interval. If a threshold for a resources was exceeded, then that statistic is marked with an asterisk (*). For descriptions of the monitored resources, see *Setting Bandwidth Thresholds* on page 11-3 and *Setting Miscellaneous Thresholds* on page 11-4 above.

For descriptions of the statistics, see *View Switch-Level Statistics* on page 11-6.

◆ **Note** ◆

The CPU and memory resources are not applicable to the module level statistics display, and therefore are not shown.

View Port-Level Statistics

To view port-level statistics, type the **hpstat** command at a system prompt as shown:

```
hpstat <slot>/<port>
```

where **<slot>** is the slot number and **<port>** is the port number. For example to view port 1 on slot 3, enter the following:

```
hpstat 3/1
```

The following screen is displayed:

Port 3/1 Resources	Limit	Curr	1 Min Avg	1 Hr Avg	1 Hr Max
Receive	80	00	00	00	00
Transmit/Receive	80	92*	00	00	00
Backplane	80	00	00	00	00

```
/System/Health %
```

Statistics are displayed as percentages of the total resource capacity, and represent data taken from the last sampling interval. If a threshold for a resource was exceeded, then that statistic is marked with an asterisk (*). For descriptions of the monitored resources, see *Setting Bandwidth Thresholds* on page 11-3 and *Setting Miscellaneous Thresholds* on page 11-4 above.

For descriptions of the statistics, see *View Switch-Level Statistics* on page 11-6.

Reset Health Statistics

To reset the health statistics for the switch, type the **hreset** command at a system prompt. The following message is displayed:

```
Are you sure you want to reset health statistics? (n) :
```

To confirm your choice to clear the switch health statistics, type **y** at the prompt. After you confirm your choice, the following confirmation notice is displayed:

```
RESET HEALTH STATISTICS
```

◆ **Note** ◆

The **hreset** command clears the statistics for the entire switch. You cannot clear statistics for the module or port level only.

12 Network Time Protocol

Introduction


The Network Time Protocol (NTP) is used to synchronize the time of a computer client or server to another server or reference time source, such as a radio or satellite receiver. It provides client time accuracies within a millisecond on LANs, and up to a few tens of milliseconds on WANs relative to a primary server synchronized to Coordinated Universal Time (UTC) (via a Global Positioning Service receiver, for example). Typical NTP configurations utilize multiple redundant servers and diverse network paths in order to achieve high accuracy and reliability. Some configurations include cryptographic authentication to prevent accidental or malicious protocol attacks.

It is important for networks to maintain accurate time synchronization between network nodes. The standard timescale used by most nations of the world is based on a combination of Universal Coordinated Time (UTC) (representing the Earth's rotation about its axis) and the Gregorian Calendar (representing the Earth's rotation about the Sun). The UTC timescale is disciplined with respect to International Atomic Time (TAI) by inserting leap seconds at intervals of about 18 months. UTC time is disseminated by various means, including radio and satellite navigation systems, telephone modems, and portable clocks.

Special purpose receivers are available for many time-dissemination services, including the Global Position System (GPS) and other services operated by various national governments. For reasons of cost and convenience, it is not possible to equip every computer with one of these receivers. However, it is possible to equip some computers with these clocks, which then act as primary time servers to synchronize a much larger number of secondary servers and clients connected by a common network. In order to do this, a distributed network clock synchronization protocol is required which can read a server clock, transmit the reading to one or more clients, and adjust each client clock as required. Protocols that do this include the Network Time Protocol (NTP).

Stratum

Stratum is the term used to define the relative proximity of a node in a network to a time source (such as a radio clock). Stratum 1 is the server connected to the time source itself. (In most cases the time source and the stratum 1 server are in the same physical location.) An NTP client or server connected to a stratum 1 source would be stratum 2. A client or server connected to a stratum 2 machine would be stratum 3, and so on, as demonstrated in the diagram below.



The farther away from stratum 1 a device is, the more likely there will be discrepancies or errors in the time adjustments done by NTP. A list of stratum 1 and 2 sources available to the public can be found on the Internet.

◆ Note ◆

It is not required that NTP be connected to an officially recognized time source (for example, a radio clock). NTP can use any time source to synchronize time in the network.

Using NTP in a Network


NTP operates on the premise that there is one true standard time (defined by UTC), and that if several servers claiming synchronization to the standard time are in disagreement, then one or more of them must be out of synchronization or not functioning correctly.

The stratum gradation is used to qualify the accuracy of a time source along with other factors such as advertised precision and the length of the network path between connections. NTP operates with a basic distrust of time information sent from other network entities, and is most effective when multiple NTP time sources are integrated together for checks and cross-checks.

To achieve this end, there are several modes of operation that an NTP entity can use when synchronizing time in a network. These modes help predict how the entity behaves when requesting or sending time information, listed below:

- A switch can be a client of an NTP server (usually of a lower stratum), receiving time information from the server but not passing it on to other switches.
- A switch can be a client of an NTP server, and in turn be a server to another switch or switches.
- A switch (regardless of its status as either a client or server) must be *peered* with another switch. Peering allows NTP entities in the network of the same stratum to regard each other as reliable sources of time and exchange time information.

Examples of these are shown in the simple network diagram on the following page:



Servers 1a and 1b receive time information from, or synchronize with, a UTC time source such as a radio clock. (In most cases, these servers would not be connected to the same UTC source, though it is shown this way for simplicity.) Servers 1a and 1b become stratum 1 NTP servers and are peered with each other, allowing them to check UTC time information against each other. These machines support machines 2a and 2b as clients, and these clients are synchronized to the higher stratum servers 1a and 1b.

Clients 2a and 2b are also peered with each other for time checks, and become stratum 2 NTP servers for more clients (3a and 3b, which are also peered).

In this hierarchy, the stratum 1 servers synchronize to the most accurate time source available, then check the time information with peers at the same stratum. The stratum 2 machines synchronize to the stratum 1 servers, but do not send time information to the stratum 1 machines. Machines 2a and 2b in turn provide time information to the stratum 3 machines.

It is important to consider the issue of robustness when selecting sources for time synchronization. It is suggested that at least three sources should be available, and at least one should be “close” to you in terms of network topology. It is also suggested that each NTP client is peered with at least three other same stratum clients, so that time information crosschecking will be performed.

When planning your network, it is helpful to use the following general rules:

- It is usually not a good idea to synchronize a local time server with a peer (in other words, a server at the same stratum), unless the latter is receiving time updates from a source that has a lower stratum than from where the former is receiving time updates. This minimizes common points of failure.
- Peer associations should only be configured between servers at the same stratum level. Higher Strata should configure lower Strata, not the reverse.
- It is inadvisable to configure time servers in a domain to a single time source. Doing so invites common points of failure.

NTP and Authentication

NTP is designed to use either DES or MD5 encryption authentication to prevent outside influence upon NTP timestamp information. This is done by using a key file. The key file is loaded into the switch memory, and consists of a text file that lists key identifiers that correspond to particular NTP entities.

If authentication is enabled on an NTP switch, any NTP message sent to the switch must contain the correct key ID in the message packet to use in decryption. Likewise, any message sent from the authentication enabled switch will not be readable unless the receiving NTP entity possesses the correct key ID.

Key files are created by a system administrator independent of the NTP protocol, and then placed in the switch memory. An example of a key file is shown below:

1	N	29233e0461ecd6ae	# des key in NTP format
2	M	Rlrop8KPPvQvYotM	# md5 key as an ASCII random string
14	M	sundial	# md5 key as an ASCII string
15	A	sundial	# des key as an ASCII string

In a key file, the first token is the key number ID, the second is the key format, and the third is the key itself. (The text following a “#” is not counted as part of the key, and is used merely for description.) There are 4 key formats:

N	Indicates a DES key written as a hex number, in NTP standard format with the high order bit of each octet being the odd parity bit.
M	Indicates an MD5 key written as a 1 to 31 character ASCII string with each character standing for a key octet.
A	Indicates a DES key written as a 1 to 8 character string in 7-bit ASCII format, where each character stands for a key octet string.
S	Indicates a DES key written as a hex number in the DES standard format, with the low order bit of each octet being the odd parity bit.

For information on activating authentication, specifying the location of a key file, and configuring key IDs for switches, see the following sections:

- *Configuring an NTP Client* on page 12-6
- *Configuring a New Peer Association* on page 12-12
- *Configuring a New Server* on page 12-13
- *Configuring a Broadcast Time Service* on page 12-13

Network Time Protocol Management Menu

To access the NTP management menu, connect to a switch via a console or telnet session and enter **NTP** at the system prompt. If you are in verbose mode, or enter a question mark (?) at the prompt, the following screen is displayed:

Command	NTP Management Menu
Ntconfig	Enter the NTP configuration menu
Ntinfo	Enter the NTP information menu
Ntstats	Enter the NTP statistics menu
Ntadmin	Enter the NTP administration menu
Ntaccess	Enter the NTP access control menu

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

Ntconfig. This command accesses the NTP configuration menu, which allows you to configure this NTP device, add or remove peer associations, add an NTP server, configure this NTP device's broadcast time, and set or change this NTP device's fudge factor. See *NTP Configuration Menu* on page 12-6 for more information on the NTP configuration menu.

Ntinfo. This command accesses the NTP information menu, which allows you to view a list of all peers for this NTP device, display a list of peers with summary information (in two different formats), display detailed information for one or more peers, and display local server information. See *NTP Information Menu* on page 12-15 for more information.

Ntstats. This command accesses the NTP statistics menu, which allows you to view the statistics for the loop filter, peer memory usage, I/O subsystem, local server, event time subsystem, packet counts, leap second state, clock status, monitoring routines data. See *NTP Statistics Menu* on page 12-23 for more information.

Ntadmin. This command accesses the NTP administration menu, which allows you to set the receive timeout, set an encryption delay, specify a remote NTP server, set a password and key ID for this NTP device, set and clear a system flag, and restart the NTP software. See *NTP Administration Menu* on page 12-33 for more information.

Ntaccess. This command accesses the NTP access control menu, which allows you to change the authentication key ID for request and control messages, reinitialize the key ID list, add a key ID to or remove a key ID from the trusted list, display the state of the authentication code, create or remove restrict and add flags to an entry, view a servers restriction list, remove a restriction entry from this NTP device, and configure, remove or view traps set in the server. See *NTP Access Control Menu* on page 12-36 for more information.

NTP Configuration Menu

To view the NTP configuration menu, enter the **ntpconfig** command at the system prompt. If you are in verbose mode the NTP configuration menu is displayed. Otherwise, enter a question mark (?) at the prompt to display this menu:

Command	NTP Configuration Menu
ntpconfig	Initial NTP configuration
ntpaddpeer	configure a new peer association
ntpaddserv	configure a new server
ntpbcast	configure broadcasting time service
ntpuncfig	unconfigure existing peer associations
ntpprec	set the server's advertised precision
ntpfudge	set/change one of a clock's fudge factors

Related Menus:

Ntconfig Ntinfo Ntstats Ntadmin Ntaccess

The main menu options are shown in the **Related Menus** list for quick access if you need to change menus.

A switch can be configured to act as an NTP client, or an NTP client/server. An NTP client receives updates from an NTP server without passing on time information to other clients, while an NTP client/server receives time information from a server, and acts as a server for other clients in a higher stratum.

Configuring an NTP Client

To set up the NTP client, use the **ntpconfig** command as follows:

1. Enter the command as shown, at the system prompt:

```
ntpconfig
```

The following menu appears:

NTP Startup Configuration

```

1) Response timeout           : 0
2) Authentication delay      : No
3) Authentication key file name : UNSET
4) NTP client mode           : Ucast
5) Enable monitor            : No
6) Enable NTP server         : No

```

2. Adjust the configurable variables for this NTP client as needed by entering the line number, and equal sign, and a new value at the system prompt, as shown:

```
<lineNumber>=<value>
```

For example, to change the **Response timeout** to 10, you would enter **1** (the line number for **Response timeout**), an equal sign (=), and the number **10** (the new value), as shown:

```
1=10
```

After enabling NTP for this switch, you need to configure at least one peer association, unless you will be supplying time synchronization. In that case, you need to configure a reference clock.

For information on adding a peer association, see *Configuring a New Peer Association* on page 12-12.

Field Descriptions

The following section describes the fields displayed using the `ntpconfig` command.

1) Response timeout

This field sets the timeout period for responses to server queries. Server queries come from the server responsible for providing this client with NTP time information. The default is 8000 milliseconds.

2) Authentication delay

This field sets a specified time interval that is added to timestamps included in requests to the server that required authentication. Typically this delay is needed in cases of long delay paths, or of servers whose clocks are unsynchronized.

3) Authentication key file name

The key file is a file that holds the NTP authentication keys used during remote access or configuration of the server responsible for this client. This field allows you to specify the name of the key file. The key file should be kept in the `/flash` directory of the switch.

Specifying a key file expands the NTP Startup Configuration menu. For more information on configuring authentication, see *Configuring Client/Server Authentication* on page 12-9.

4) NTP client mode

This field allows you to set how the client mode of this device sends its server queries. The options are **U** (for unicast), **B** (for broadcast), or **M** (for multicast).

Setting the NTP client mode to broadcast or multicast expands the NTP Startup Configuration menu. A suboption for the NTP client mode appears, allowing you to specify the broadcast or multicast address, as shown:

41) NTP multicast address :

Enter the broadcast or multicast address at the prompt by typing line number **41**, an equal sign (=), and the IP address. For example, to specify a multicast address of 204.0.1.1, you would enter the following:

41=204.0.1.1

5) Enable monitor

This field turns NTP monitoring on or off. Entering **yes** activates NTP monitoring, while entering **no** deactivates this function. The statistics for monitoring can be viewed using the `ntpmon` command in the statistics menu. See *NTP Statistics Menu* on page 12-23 for more information.

6) Enable NTP server

This field allows you to enable the server portion of the NTP software for this NTP device. When set to **yes**, this device can act as an NTP server for other clients. When set to **no**, this device is only a client of another NTP server.

Configuring an NTP Client/Server

A switch can be configured to act both as a client and a server. If you want to run both the client and server portions of the NTP software, follow the steps below:

1. Enter the command as shown, at the system prompt:

```
ntpconfig
```

The following menu appears:

NTP Startup Configuration

```
1) Response timeout           : 0
2) Authentication delay       : No
3) Authentication key file name : UNSET
4) NTP client mode            : Ucast
5) Enable monitor             : No
6) Enable NTP server          : No
```

2. Adjust the configurable variables for this NTP client as needed by entering the line number, and equal sign, and a new value at the system prompt, as shown:

```
<lineNumber>=<value>
```

For example, to change the **Response timeout** to 10, you would enter **1** (the line number for **Response timeout**), an equal sign (=), and the number **10** (the new value), as shown:

```
1=10
```

3. Enable the NTP server by entering a **6**, an equal sign (=), and **yes** at the prompt, as shown:

```
6=yes
```

The NTP Startup Configuration menu expands to display new options. The menu now appears similar to the following:

NTP Startup Configuration

```
1) Response timeout           : 0
2) Authentication delay       : No
3) Authentication key file name : UNSET
4) NTP client mode            : Ucast
5) Enable monitor             : No
6) Enable NTP server          : No
  61) Client limit             : 3
  62) Client limit period      : 3600
  63) Enable server authentication : No
  64) Advertised precision     : -7
  65) Broadcast delay         : 0
```

4. Adjust the configurable variables for this NTP server as needed by entering the line number, and equal sign, and a new value at the system prompt, as shown:

```
<lineNumber>=<value>
```

For example, to change the **Client limit** to 10, you would enter **61** (the line number for **Client limit**), an equal sign (=), and the number **10** (the new value), as shown:

```
61=10
```

Field Descriptions

The following section describes the expanded menu options.

61) Client limit

This field allows you to set a specific number of clients that are allowed to make requests of the server during a specified time period. Setting this field to **0** allows an unlimited number of clients to connect to the server.

62) Client limit period

This field allows you to set the client limit time period (in seconds). This along with the **client limit** field above determine how many clients are allowed to make requests of this server.

63) Enable server authentication

This field enables the authentication of unsynchronized peers. If set to **yes**, NTP only synchronizes with peers that has been authenticated with the correct key ID.

64) Advertised precision

Sets the precision which the server advertises to the specified value. This should be a negative integer in the range -4 through -20.

65) Broadcast delay

This fields allows you to set a specified network delay time. Normally, NTP automatically compensates for the network delay between the broadcast/multicast server and the client. If this calibration fails, the delay set here is used instead.

Configuring Client/Server Authentication

In order to use authentication, you must specify a key file. A key file contains the keys necessary for NTP to decode encrypted NTP messages. To specify a key file, follow the steps below:

1. Enter the command as shown, at the system prompt:

```
ntpiconfig
```

The following menu appears:

NTP Startup Configuration

1) Response timeout	: 0
2) Authentication delay	: No
3) Authentication key file name	: UNSET
4) NTP client mode	: Ucast
5) Enable monitor	: No
6) Enable NTP server	: No

2. Adjust the configurable variables for this NTP client as needed by entering the line number, and equal sign, and a new value at the system prompt, as shown:

<lineNumber>=<value>

For example, to change the **Response timeout** to 10, you would enter **1** (the line number for **Response timeout**), an equal sign (=), and the number **10** (the new value), as shown:

1=10

3. Enable authentication by entering a **3**, and equal sign (=), and a key file name at the prompt, as shown:

3=ntp.keys

The NTP Startup Configuration menu expands to display new options. The menu now appears similar to the following:

NTP Startup Configuration

1) Response timeout	: 0
2) Authentication delay	: No
3) Authentication key file name	: ntp.keys
31) Configuration info authentication key	:
32) Control request authentication key	:
33) Configuration change authentication key	:
4) NTP client mode	: Ucast
5) Enable monitor	: No
6) Enable NTP server	: No

4. Adjust the configurable variables for authentication as needed by entering the line number, and equal sign, and a new value at the system prompt, as shown:

<lineNumber>=<value>

For example, to change the **Configuration info authentication key** to 10, you would enter **1** (the line number for **Configuration info authentication key**), an equal sign (=), and the number **10** (the new value), as shown:

1=10

Field Descriptions

The following section describes the expanded menu options.

31) Configuration info authentication key

The number of the key in the key file used to authenticate configuration information. Configuration information sets configuration parameters. For more information on the key file, see *NTP and Authentication* on page 12-4.

32) Control request authentication key

The number of the key in the key file used to authenticate control requests. Control requests come from other NTP clients and servers. For more information on the key file, see *NTP and Authentication* on page 12-4.

33) Configuration change authentication key

The number of the key in the key file used to authenticate configuration change requests. Configuration change requests come from other NTP clients and servers. For more information on the key file, see *NTP and Authentication* on page 12-4.

Configuring a New Peer Association

When you have configured the NTP client and/or server, you will need to set at least one peer association for the switch. An NTP peer is a machine of the same stratum that will compare and check time information sent from the switch, and in turn send time information to the switch.

To configure a new peer, enter the **ntpaddpeer** command in the following manner:

```
ntpaddpeer <address> [<keyld> <version> <minpol>] [prefer]
```

where **<address>** is either the domain name or IP address of the peer machine. The optional configuration items are described below:

<keyld>. An unsigned 32-bit integer key identifier for encryption authentication. The default is for no key ID.

<version>. The version of NTP being used. The options are versions 1, 2, or 3. If no number is entered, it is assumed that version 3 is being used.

<minpol>. The minimum poll interval for time checks to this peer. The number entered is seconds raised to the power of 2.

prefer. An identifier that marks this peer as a preferred source of time information. In a situation where multiple peers could provide time information to this client, the preferred peer is the one that is used.

For example, to add a peer with an address of 1.1.1.1, a key identifier of 5, using version 3 of NTP, minimum poll of 16 seconds, and marked as a preferred server, you would enter the following:

```
ntpaddpeer 1.1.1.1 5 3 4 prefer
```

When you have finished press **<return>**. A brief message appears confirming the addition of a new peer.

Configuring a New Server

For the switch to synchronize its time, you must specify a server, or servers, from which the switch receives time information. This is done with the **ntpaddserv** command.

To add a synchronization server to a switch, use the command that follows:

```
ntpaddserv <address> [<keyId><version><minpol>] [prefer]
```

where **<address>** is either the domain name or IP address of the server. The optional configuration items are described below:

<keyId>. An unsigned 32-bit integer key identifier for encryption authentication. The default is no key ID.

<version>. The version of NTP being used. The options are versions 1, 2, or 3. If no number is entered, it is assumed that version 3 is being used.

<minpol>. The minimum poll interval for time checks to this server. The number entered is seconds raised to the power of 2.

prefer. An identifier that marks this peer as a preferred source of time information. In a situation where multiple peers could provide time information to this client, the preferred peer is the one that is used.

For example, to add a peer with an address of 1.1.1.1, a key identifier of 5, using version 3 of NTP, with a poll time of 16, and marked as a preferred server, you would enter the following:

```
ntpaddpeer 1.1.1.1 5 3 4 prefer
```

When you have finished press **<return>**. A brief message appears confirming the addition of a new server.

Configuring a Broadcast Time Service

The NTP server can be configured to operate in broadcast mode, where the server sends periodic broadcast messages to a client population by using the broadcast or multicast address specified. To configure the server to use a broadcast or multicast address, enter the **ntpbcast** command as shown:

```
ntpbcast <address> [<keyId>] [<version>] [<minpol>]
```

where **<address>** is either the domain name or the broadcast or multicast address.

◆ Important Note ◆

A multicast address of 224.0.1.1 has been assigned to NTP. Presently, this is the only address that should be used for multicast messages.

The optional configuration items are described below:

<keyId>. An unsigned 32-bit integer key identifier for encryption authentication. The default is no key ID.

<version>. The version of NTP being used. The options are versions 1, 2, or 3. If no number is entered, it is assumed that version 3 is being used.

<minpol>. The minimum poll interval for time checks to this server. The number entered is in seconds raised to the power of 2.

For example, to add broadcast address 1.1.1.1 with a key identifier of 5, using version 3 of NTP, and a minimum poll time of 16 seconds, you would enter the following:

```
ntpbcast 1.1.1.1 5 3 4
```

When you have finished press **<return>**. A brief message appears confirming the addition of a new server.

Unconfigure Existing Peer Associations

You can remove server, peer, or reference clock associations for this switch using the **ntpunconfig** command. This will remove a selected address from this switch's list of configured addresses. To do this, enter the **ntpunconfig** command as follows:

```
ntpunconfig <address>
```

where **<address>** is either the domain name or IP address of the association. For example, to remove a peer association with address 1.1.1.1, enter the following:

```
ntpunconfig 1.1.1.1
```

When you have finished press **<return>**. A brief message appears confirming the addition of a new server.

You can remove multiple addresses at one time by adding additional addresses to the command. For example, to remove a peer association with address 1.1.1.1 and a reference clock association with address 1.1.1.2, enter:

```
ntpunconfig 1.1.1.1 1.1.1.2
```

When you have finished press **<return>**. A brief message appears confirming the removal of the association.

Set the Server's Advertised Precision

If necessary, you can adjust the server's advertised precision. The precision of a server is a signed integer indicating the precision of the clocks in seconds to the nearest power of 2. It determines how accurate the clock is under normal circumstances, and allows NTP to determine which is the best time source for synchronization. To set the server's advertised precision, enter the **ntpprec** command as shown:

```
ntpprec <interval>
```

where **<interval>** is the signed integer in seconds. This number must be between -4 and -20. For example, to set the server's advertised precision to -5, you would enter the following:

```
ntpprec -5
```

When you have finished press **<return>**. A brief message appears confirming the change of the advertised precision.

◆ Note ◆

The determination of a server's advertised precision is based largely on the clock type used as the ultimate time source (stratum 1).

NTP Information Menu

To view the NTP configuration menu, enter the **ntinfo** command at the system prompt. If you are using verbose mode, the NTP configuration menu is displayed. Otherwise, enter a question mark (?) at the prompt to display this menu:

Command	NTP Information Menu
ntplpeers	display list of peers the server knows about
ntppeers	display peer summary information
ntpdmpeers	display peer summary info the way Dave Mills likes it
ntpshowpeer	display detailed information for one or more peers
ntpvers	print version number
ntpinfo	display local server information

Related Menus:

Ntconfig Ntinfo Ntstats Ntadmin Ntaccess

The main menu options are shown in the **Related Menus** list for quick access if you need to change menus.

Display List of Peers the Server Knows About

The **ntplpeers** command is used to display a brief list of all NTP associations related to this switch (servers, peers, etc.).

To display a list of NTP associations, enter the **ntplpeers** command at the system prompt. A display similar to the following is shown:

```
client 1.1.1.1
client 1.1.1.2
sym_active 1.1.1.3
```

The list shows the mode this switch is using in relation to the association, and the address of the remote association. The address is either a domain name or an IP address. The available modes are as follows:

- Symmetric Active (1)** A host in this mode sends periodic messages regardless of the reachability state of stratum of its peer. By operating in this mode the host announces its willingness to synchronize and be synchronized by the peer.
- Symmetric Passive (2)** This type of association is ordinarily created upon the arrival of a message from a peer operating in the symmetric active mode and persists only as long as the peer is reachable and operating at a stratum level less than or equal to the host; otherwise the association is dissolved. The association will always persist until at least one message has been sent in reply. By operating in this mode the host announces its willingness to synchronize and be synchronized by the peer.
- Client (3)** A host operating in this mode sends periodic messages regardless of the reachability state of stratum of its peer. By operating in this mode the host, usually a LAN workstation, announces its willingness to be synchronized, but not to synchronize the peer.

- Server (4)** This type of association is ordinarily created upon arrival of a client request message and exists only in order to reply to that request, after which the association is dissolved. By operating in this mode the host, usually a LAN time server, announces its willingness to synchronize, but not be synchronized by the peer.
- Broadcast (5)** A host operating in this mode sends periodic messages regardless of the reachability state or stratum of the peers. By operating in this mode, the host, usually a LAN time server operating on a high-speed broadcast medium, announces its willingness to synchronize all peers, but not be synchronized by any of them.

◆ **Note** ◆

The mode of the switch in relation to the remote association is determined when you create the association. See *NTP Configuration Menu* on page 12-6 for more information on creating NTP associations.

Display Peer Summary Information

The **ntppeers** command displays a more detailed version of the **ntplpeers** command. To display a list of peers that includes summary information, enter the **ntppeers** command at the system prompt. A screen similar to the following appears:

	remote	local	st	poll	reach	delay	offset	disp
=	1.1.1.1	0.0.0.5	16	64	0	0.00000	0.00000	16.0000
+	1.1.1.2	0.0.0.5	1	64	0	0.00000	0.00000	16.0000
=	1.1.1.3	0.0.0.5	2	64	0	0.00000	0.00000	16.0000

The symbols at the very left of this table note the relationship (mode) of the switch to the remote association. The section below is a key for interpreting these symbols:

- + The switch is in symmetric active mode.
- The switch is in symmetric passive mode.
- = The switch is in client mode.
- ^ The switch is broadcasting to this address.
- ~ The switch is receiving broadcasts from this address.
- * The switch is currently synchronizing with this address.

Field Descriptions

The following sections describe the fields displayed using the **ntppeers** command

Remote. The IP address of the remote association.

Local. The local interface address assigned by NTP to the remote association. If this address is **0.0.0.0**, then the local address has yet to be determined.

St. The stratum level of the remote peer. If this number is **16**, the remote peer has not been synchronized.

Poll. The polling interval, in seconds.

Reach. The reachability register of the remote association, in octal format. This number is determined by the NTP algorithm.

Delay. The currently estimated delay of this remote association, in seconds. This time is determined by the NTP algorithm.

Offset. The currently estimated offset of this remote association, in seconds. This time is determined by the NTP algorithm.

Disp. The currently estimated dispersion of this remote association, in seconds. This time is determined by the NTP algorithm.

Display Alternate Peer Summary Information

The **ntpdmpeers** command displays a more detailed version of the **ntpshowpeer** command with a slightly different output than the **ntppeers** command. To display a list of peers that includes summary information, enter the **ntpdmpeers** command at the system prompt. A screen similar to the following appears:

	remote	local	st	poll	reach	delay	offset	disp
+	1.1.1.1	0.0.0.5	16	64	0	0.00000	0.00000	16.0000
+	1.1.1.2	0.0.0.5	1	64	0	0.00000	0.00000	16.0000
*	1.1.1.3	0.0.0.5	2	64	0	0.00000	0.00000	16.0000

This table is identical to the **ntppeers** command except for the symbols displayed on the far left side. A key for the symbols is provided below:

- . Indicates that the remote association was cast aside during the false ticker detection.
- +
- Indicates that the remote association was accepted and not discarded by the false ticker detection.
- *
- Indicates the remote association the switch is currently synchronizing with.

Display Detailed Information for One or More Peers

The `ntpshowpeer` command allows you to view detailed NTP information about any remote associations of this switch. To view detailed NTP information about a remote association enter the `ntpshowpeer` command in the following manner:

```
ntpshowpeer <address>
```

where `<address>` is either the domain name or IP address of the remote association. For example, to show information for a peer with IP address 1.1.1.4, enter:

```
ntpshowpeer 1.1.1.4
```

A screen similar to the following is displayed:

```
remote 1.1.1.4, local 0.0.0.6
hmode sym_active, pmode server, stratum 16, precision -7
leap 11, refid [0.0.0.0], rootdistance 0.00000, rootdispersion 0.00000
ppoll 6, hpoll 6, keyid 0, version 3, association 41807
valid 0, reach 000, unreachable 0, flash 000, boffset 0.00391, ttl/mode 0
timer 32s, flags config, bclient
reference time:      00000000.00000000 Thu, Feb 7 1936 6:28:16.000
originate timestamp: 00000000.00000000 Thu, Feb 7 1936 6:28:16.000
receive timestamp:  00000000.00000000 Thu, Feb 7 1936 6:28:16.000
transmit timestamp:  00000000.00000000 Thu, Feb 7 1936 6:28:16.000
filter delay:        0.00000 0.00000 0.00000 0.00000
                    0.00000 0.00000 0.00000 0.00000
filter offset:        0.000000 0.000000 0.000000 0.000000
                    0.000000 0.000000 0.000000 0.000000
filter order:         7   6   5   4
                    3   2   1   0
offset 0.000000, delay 0.000000, dispersion 16.000000, selectdisp 0.000000
```

It is possible to display information from more than one remote association by adding more addresses when entering the `ntpshowpeer` command. For example, to display information on a peer with IP address 1.1.1.4 and a peer with IP address 1.1.1.5, enter:

```
ntpshowpeer 1.1.1.4 1.1.1.5
```

Field Descriptions

The following section describes the fields displayed using the `ntpshowpeer` command.

Remote. The IP address of the remote association.

Local. The local interface address assigned by NTP to the remote association. If this address is `0.0.0.0`, then the local address has yet to be determined.

Hmode. The host mode of this remote association. There are five possible modes: symmetric active, symmetric passive, client, server, and broadcast. The displayed mode is assumed if this association becomes the switch's host NTP server. For a description of the modes, see *Display List of Peers the Server Knows About* on page 12-15. For a description of how to set a switch host NTP server, see *Specify the Host Whose NTP Server We Talk To* on page 12-34.

Pmode. The peer mode of this remote association. There are five possible modes: symmetric active, symmetric passive, client, server, and broadcast. The displayed mode is assumed if this association becomes the switch's host NTP server. For a description of the modes, see *Display List of Peers the Server Knows About* on page 12-15. For a description of how to configure a peer, see *Configuring a New Peer Association* on page 12-12.

Stratum. The stratum level of the remote peer. If this number is `16`, the remote peer has not been synchronized.

Precision. The advertised precision of this association, which is a number from -4 to -20. For information on setting the advertised precision, see *Configuring an NTP Client* on page 12-6 and *Set the Server's Advertised Precision* on page 12-14.

Leap. The status of leap second insertion for this association. Leap seconds are seconds that are added to the timestamp of an NTP entity to correct accumulated time errors. The possible values are:

00	No warning.
01	Last minute has 61 seconds.
10	Last minute has 59 seconds.
11	Alarm condition (clock not synchronized).

Refid. This is a 32-bit code identifying the particular reference clock. In the case of stratum 0 (unspecified) or stratum 1 (primary reference source), this is a four-octet, left-justified, zero-padded ASCII string. In the case of stratum 2 and greater (secondary reference) this is the four-octet Internet address of the peer selected for synchronization.

Rootdistance. This is a signed fixed-point number indicating the total roundtrip delay to the primary reference source at the root of the synchronization subnet, in seconds. Note that this variable can take on both positive and negative values, depending on clock precision and skew.

Rootdispersion. This is a signed fixed-point number indicating the maximum error relative to the primary reference source at the root of the synchronization subnet, in seconds. Only positive values are possible.

Ppoll. The poll time for this association when it is a peer. This number is the minimum interval between transmitted messages, in seconds as a power of two. For instance, a value of six indicates a minimum interval of 64 seconds.

Hpoll. The poll time for this association when it is a host. This number is the minimum interval between transmitted messages, in seconds as a power of two. For instance, a value of six indicates a minimum interval of 64 seconds.

KeyID. This is an integer identifying the cryptographic key used to generate the message authentication code.

Version. The version of NTP this association is using; the options are **1**, **2**, or **3**.

Association. The number of seconds since this NTP entity was associated with the switch.

Valid. This is an integer counter indicating the valid samples remaining in the filter register. It is used to determine the reachability state of an association, and when the poll interval should be increased or decreased.

Reach. This is a shift register used to determine the reachability status of this peer. The NTP algorithm uses this when determining timestamp information.

Unreach. The number of times this NTP entity was unreachable.

Flash. This field displays the number of error bits from the packet procedure.

Boffset. This field displays the default broadcast delay in seconds.

TTL/mode. This fields displays the Time-to-Live (TTL) time in seconds and the mode (unicast, multicast, or broadcast) of NTP messages sent to a broadcast address. For information on configuring an NTP broadcast address, see *Configuring a Broadcast Time Service* on page 12-13.

Timer. Shows the number of seconds until the next NTP message is sent to an association.

Flags Config. This counter lists what flags have been configured for this NTP entity. For more information about setting flags, see *Set a System Flag (Auth, Bclient, Monitor, Stats)* on page 12-35.

Reference Time. This is the local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.

Originate Timestamp. This is the local time, in timestamp format, of the peer when its last NTP message was sent. If the peer becomes unreachable the value is set to zero.

Receive Timestamp. This is the local time, in timestamp format, when the latest NTP message from the peer arrived. If the peer becomes unreachable the value is set to zero.

Transmit Timestamp. This is the local time, in timestamp format, when the last NTP message was sent from this association.

Filter delay. NTP comes with various filter routines as part of the algorithm that determines timestamp information. This field shows the delay in seconds the NTP algorithm uses to correct for delays caused by messages traversing through the NTP filters.

Filter offset. NTP comes with various filter routines as part of the algorithm that determines timestamp information. This counter indicates the offset of the peer clock relative to the local clock due to filters.

Filter order. The order in which NTP messages pass through filters.

Delay. The currently estimated delay of this remote association, in seconds. This number indicates the roundtrip delay of the peer clock relative to the local clock over the network path between them, in seconds. Note that this variable can take on both positive and negative values, depending on clock precision and skew-error accumulation. This time is determined by the NTP algorithm.

Offset. The currently estimated offset of this remote association, in seconds. This counter indicates the offset of the peer clock relative to the local clock. This time is determined by the NTP algorithm.

Disp. The currently estimated dispersion of this remote association, in seconds. This counter indicates the maximum error of the peer clock relative to the local clock over the network path between them, in seconds. Only positive values greater than zero are possible. This time is determined by the NTP algorithm.

Print Version Number

The **ntpvers** is used to show the version number of the xntp file. To display the version number, enter the **ntpvers** command at the system prompt. A message similar to the following is shown:

```
xntp Fri Apr 9 22:52:46 PDT 1999 (1)
```

Display Local Server Information

The `ntpinfo` command is used to display information about the local switch's implementation of NTP. To view local switch NTP information, enter the `ntpinfo` command at the system prompt. A screen similar to the following is shown:

```

system peer:          0.0.0.0
system peer mode:    unspec
leap indicator:      11
stratum:             16
precision:           -7
root distance:       0.00000 s
root dispersion:     0.00000 s
reference ID:        [0.0.0.0]
reference time:      00000000.00000000 Thu, Feb 7 1936 6:28:16.000
system flags:        monitor stats
frequency:           0.000 ppm
stability:           0.000 ppm
broadcastdelay:     0.003906 s
authdelay:           0.000122 s

```

Field Descriptions

The following section explains the fields shown using the `ntpinfo` command.

System peer. The IP address of the switch.

System peer mode. The peer mode of this remote association. There are five possible modes: symmetric active, symmetric passive, client, server, and broadcast. The displayed mode is assumed if this association becomes the switch's host NTP server. For a description of the modes, see *Display List of Peers the Server Knows About* on page 12-15. For a description of how to configure a peer, see *Configuring a New Peer Association* on page 12-12.

Leap indicator. The status of leap second insertion for this association. Leap seconds are seconds that are added to the timestamp of an NTP entity to correct accumulated time errors. The possible values are:

00	No warning.
01	Last minute has 61 seconds.
10	Last minute has 59 seconds.
11	Alarm condition (clock not synchronized)

Stratum. The stratum level of the remote peer. If this number is **16**, the remote peer has not been synchronized.

Precision. The advertised precision of the switch. It will be a number between -4 and -20.

Root distance. This is a signed fixed-point number indicating the total roundtrip delay to the primary reference source at the root of the synchronization subnet, in seconds. Note that this variable can take on both positive and negative values, depending on clock precision and skew.

Rootdispersion. This is a signed fixed-point number indicating the maximum error relative to the primary reference source at the root of the synchronization subnet, in seconds. Only positive values are possible.

Reference ID. This is a 32-bit code identifying the particular reference clock. In the case of stratum 0 (unspecified) or stratum 1 (primary reference source), this is a four-octet, left-justified, zero-padded ASCII string. In the case of stratum 2 and greater (secondary reference) this is the four-octet Internet address of the peer selected for synchronization.

Reference time. This is the local time at which the local clock was last set or corrected.

System Flags. This counter lists what flags have been configured for this NTP entity. For more information about setting flags, see *Set a System Flag (Auth, Bclient, Monitor, Stats)* on page 12-35.

Frequency. A number indicating the local clock's frequency in relation to a reference clock's Pulse per Second (PPS). If the clock is running in perfect synchronization, this number should be 1. Otherwise, it will be slightly lower or higher in order to compensate for the time difference.

Stability. The residual frequency error (in seconds) remaining after the system frequency correction is applied.

Broadcastdelay. The broadcast delay, in seconds, of this association. For information on how to set the broadcast delay, see *Configuring a Broadcast Time Service* on page 12-13.

Authdelay. The authentication delay, in seconds, of this association. For information on how to set the authentication delay, see *Set the Delay Added to Encryption Time Stamps* on page 12-33.

NTP Statistics Menu

To view the NTP Statistics Menu, enter the **ntstats** command at the system prompt. If you are in verbose mode the NTP configuration menu is displayed. Otherwise, enter a question mark (?) at the prompt to display this menu:

Command	NTP Statistics Menu
ntpstat	display local server statistics
ntppstat	display server statistics associated with particular peer(s)
ntploopinfo	display loop filter information
ntpmem	display peer memory usage statistics
ntpio	display I/O subsystem statistics
ntptimer	display event timer subsystem statistics
ntppreset	reset various subsystem statistics counters
ntppreset	reset stat counters associated with particular peer(s)
ntpctlstat	display packet count statistics from the control module
ntpleap	display the current leap second state
ntpmmon	turn the server's monitoring facility on or off
ntpmplist	display data the server's monitor routines have collected

Related Menus:

Ntconfig Ntinfo Ntstats Ntadmin Ntaccess

The main menu options are shown in the **Related Menus** list for quick access if you need to change menus.

Display Local Server Statistics

The **ntpstat** command allow you to view statistics for the local NTP entity (switch). To view statistics, enter the **ntpstat** command at the system prompt. A display similar to the following is displayed:

```

system uptime:           0
time since reset:       0
bad stratum in packet:  0
old version packets:    0
new version packets:    16
unknown version number: 0
bad packet length:      0
packets processed:      0
bad authentication:     0
limitation rejects:     0

```

Field Descriptions

The following section describes the fields displayed using the **ntpstat** command.

system uptime. The number of seconds the local NTP server has been associated with the switch.

time since reset. The number of seconds since the last time the local NTP server was restarted.

bad stratum in packet. The number of NTP packets received that had a corrupted stratum bit in the data of the packet.

old version packets. The number of NTP packets received that were of an older version of NTP (either version 1 or 2).

new version packets. The number of NTP packets received that were version 3 of NTP.

unknown version number. The number of NTP packets received for which the version was unknown (most likely due to packet corruption).

bad packet length. The number of NTP packets received that did not fit the NTP packet structure (most likely due to packet corruption).

packets processed. The total number of NTP packets processed.

bad authentication. The number of NTP packets rejected because they did not meet authentication standards.

limitation rejects. The number of NTP packets rejected because there were restrictions set on their point of origin. For information on setting restrictions, see *Create Restrict Entry/Add Flags to Entry* on page 12-39.

Display Server Statistics Associated with Particular Peer(s)

The **ntppstat** command allows you to view statistics for a specific NTP peer. To view statistics for a peer, enter the **ntppstat** command as shown:

```
ntppstat <ipAddress>
```

where **<ipAddress>** is the address of the peer for which you want to view statistics. For example, to view statistics for a peer with IP address 131.218.18.4, enter the following:

```
ntppstat 131.216.18.4
```

A screen similar to the following displays:

```
remote host           : 131.216.18.4
local interface       : 0.0.0.0
time last received    : 9s
time until next send  : 6s
reachability change   : 2973s
packets sent          : 184
packets received      : 181
bad authentication    : 2
bogus origin          : 2
duplicate             : 6
bad dispersion        : 69
bad reference time    : 1
candidate order       : 1
```

Field Descriptions

The following section describes the fields displayed using the **ntpstat** command.

remote host. The IP address of the host whose statistics you are viewing.

local interface. The local interface address assigned by NTP to the remote association. If this address is **0.0.0.0**, then the local address has yet to be determined.

time last received. The number of seconds since the last NTP message packet was received from another NTP entity in the network.

time until next send. The number of seconds until this NTP peer sends out an NTP message packet.

reachability change. This field displays the number of times this client/server's reachability has changed.

packets sent. The number of NTP message packets this peer has sent out.

packets received. The number of NTP message packets this peer has received.

bad authentication. The number NTP message packets this peer has rejected due to failed authentication.

bogus origin. The number of times a response packet from another NTP entity doesn't match the request packet sent out by this client/server.

duplicate. The number of identical NTP message packets this peer has received.

bad dispersion. The number of packets that were discarded due to overly large error dispersions.

bad reference time. The number of packets that were discarded because the contained reference time didn't match the local peer expectation.

candidate order. A number that represents this client/server's synchronization order. A lower number represents a reliable synchronization source.

Display Loop Filter Information

The loop filter is used to control and correct the phase of timestamps as processed by the local clock. The loop filter examines timestamps sent to and from the local clock and can adjust them to account for natural wander and jitter.

To view the statistics of the loop filter, enter the **ntploop** command at the system prompt. A screen similar to the following is shown:

```

offset:          0.000000 s
frequency:      0.000 ppm
poll adjust:    0
watchdog timer: 0 s
    
```

All of these field variables are determined by the NTP algorithm

Field Descriptions

The following section describes the fields displayed using the **ntploop** command.

offset. The currently estimated offset of this remote association, in seconds. This counter indicates the offset of the peer clock relative to the local clock.

frequency. A number indicating the local clock's frequency in relation to a reference clock's Pulse per Second (PPS). If the clock is running in perfect synchronization, this number should be 1. Otherwise, it will be slightly lower or higher in order to compensate for the time discrepancy between the reference clock and the local clock.

poll adjust. The number of times the poll time has been adjusted to conform to the network.

watchdog timer. The number of seconds since the local clock for this client/server was last adjusted.

Display Peer Memory Usage Statistics

The memory usage for the NTP information on the switch can be displayed using the **ntpmem** command. To view memory information, enter the **ntpmem** command at the system prompt. A screen similar to the following is shown:

```

time since reset: 0
total peer memory: 15
free peer memory: 11
calls to findpeer: 0
new peer allocations: 0
peer demobilizations: 0
hash table counts: 1 0 1 0 0 1 0 0
                   0 0 0 0 0 0 0 0
                   0 0 0 0 0 0 0 0
                   0 0 0 0 0 0 1 0
    
```

Field Descriptions

The following section describes the fields displayed using the **ntpmem** command.

time since reset. The number of seconds since the last reset of NTP (usually a reboot of the switch).

total peer memory. The total number of NTP associations possible for this switch.

free peer memory. The number of available spots on this switch for NTP associations.

calls to findpeer. The number of times the switch sent an NTP packet of any kind to a configured NTP association.

new peer allocations. The number of new NTP associations created since the last restart.

peer demobilizations. The number NTP associations lost since the last restart.

hash table counts. The number of peer tables hashed to the index.

Display I/O Subsystem Statistics

The **ntpio** command allows you to view general statistics on received and transmitted NTP packets for this switch. To view the I/O statistics, enter the **ntpio** command at the system prompt. A screen similar to the following is displayed:

```

time since reset:      0
receive buffers:     10
free receive buffers:  9
used receive buffers: 0
low water refills:    0
dropped packets:      0
ignored packets:      0
received packets:     18
packets sent:         17
packets not sent:     0
interrupts handled:   18
received by int:      18

```

Field Descriptions

The following section describes the fields displayed using the **ntpio** command.

time since reset. The number of seconds since the last restart of NTP.

receive buffers. The number of switch receive buffers currently allocated by this NTP entity.

free receive buffers. The number of free receive buffers.

used receive buffers. The number of receive buffers being used.

low water refills. The number of times memory has been added.

dropped packets. The number of packets discarded due to lack of resources (i.e., memory).

ignored packets. The number of packets ignored by this client/server.

received packets. The total number of NTP packets received by the switch.

packets sent. The total number of NTP packets sent by the switch.

packets not sent. The number of NTP packets generated but not sent due to restrictions. For information on NTP restrictions, see *Create Restrict Entry/Add Flags to Entry* on page 12-39.

interrupts handled. The number of times NTP information was interrupted in the process of transmitting or receiving.

received by int. The number of packets received by interrupts.

Display Event Timer Subsystem Statistics

The **ntptimer** command allows you to view significant NTP events that have occurred on this switch. To view significant NTP events, enter the **ntptimer** command at the system prompt. A screen similar to the following is displayed:

```
time since reset:      0
alarms handled:       0
alarm overruns:       0
calls to transmit:    0
```

Field Descriptions

The following section describes the fields displayed using the **ntptimer** command.

time since reset. The number of seconds since the last reset of NTP.

alarms handled. The number of NTP alarms generated by this switch. NTP alarms occur when the NTP algorithm determines that an NTP entity is out of synchronization.

alarm overruns. The number of times the NTP alarm routine was backed up.

calls to transmit. The number of requests from other NTP entities for information, either configuration, statistical, or timestamp.

Reset Various Subsystem Statistics Counters

To reset the counters displayed for the commands used in the NTP Statistics Menu (**ntpstat**, **ntploopinfo**, **ntpio**, and **ntptimer**), use the **ntppreset** command. To reset the statistics, enter the **ntppreset** command at the system prompt followed by one or more of the following flags:

- **io**
- **sys**
- **mem**
- **timer**
- **auth**
- **allpeers**

A brief message is displayed confirming the command.

Reset Stat Counters Associated With Particular Peer(s)

It is possible to remotely reset statistics for other NTP associations from the switch. To reset statistics for an NTP association, enter the **ntppreset** command as follows:

```
ntppreset <address>
```

where **<address>** is either the domain name or IP address of the remote association. For example, to reset statistics for a peer with IP address 1.1.1.4, enter:

```
ntppreset 1.1.1.4
```

It is possible to reset the statistics for more than one NTP association at a time by adding more than one address to the command. For example, to reset statistics for a peer with IP address 1.1.1.4 and a peer with IP address 1.1.1.5, you would enter:

```
ntppreset 1.1.1.4 1.1.1.5
```

A brief message is displayed confirming the command.

Display Packet Count Statistics from the Control Module

In a comprehensive network-management environment, facilities should exist to perform routine NTP control and monitoring functions. The control module of NTP is responsible for sending and receiving control messages. To display the statistics for the control module, enter the **ntpctlstat** command at the system prompt. A screen similar to the following is shown:

```
time since reset:      0
requests received:    0
responses sent:       0
fragments sent:       0
async messages sent:  0
error msgs sent:      0
total bad pkts:       0
packet too short:     0
response on input:    0
fragment on input:    0
error set on input:   0
bad offset on input:  0
bad version packets:  0
data in pkt too short: 0
unknown op codes:    0
```

Field Descriptions

The following section describes the fields displayed using the **ntpctlstat** command.

time since reset. The number of seconds since the last reset of NTP (usually a switch reboot).

requests received. The number of NTP requests received from any NTP association.

responses sent. The number of NTP messages sent from this switch in response to NTP association requests.

fragments sent. The number of NTP messages sent from this switch that did not contain all appropriate NTP data. This can occur if timestamp information from other NTP entities is judged by this switch to be incorrect.

async messages sent. The number of async trap packets sent.

error msgs sent. The number of error messages sent from the switch to other NTP entities because the switch was not able to respond to the NTP entity's request.

total bad pkts. The total number of packets received that NTP was not able to read.

packet too short. The number of packets received that NTP rejected because the packet was the incorrect length.

response on input. The number of packets received that required the switch to respond to the sender with an NTP message.

fragment on input. The number of packets received that the switch that did not contain complete NTP data.

error set on input. The number of input control packets received with the error bit set.

bad offset on input. The number of NTP timestamps received that the switch disallowed because the added time offset parameter appeared to be incorrect. This can occur if an NTP entity becomes unsynchronized and generates false timestamp information.

bad version packets. The number of packets received where the version number of NTP was undefinable. This is usually caused by packet corruption.

data in pkt too short. The number of packets received that NTP rejected because the packet information was incomplete.

unknown op codes. The number of NTP packets received that contained an unreadable request or information. This is usually caused by packet corruption.

Display the Current Leap Second State

If necessary, NTP adds or subtracts a second from the timestamps sent out on the network to correct for errors in time information. These modifications are called leap seconds. To display leap second information for the switch, enter the **ntpleap** command at the system prompt. A screen similar to the following is displayed:

```
sys.leap:                11 (clock out of sync)
leap.indicator:          00 (leap controlled by lower stratum)
leap.warning:            00 (leap controlled by lower stratum)
leap.bits:               00 (no leap second scheduled)
time to next leap interrupt: 1 s
date of next leap interrupt: Tue, Jul 6 1999 12:38:45
calls to leap process:   0
leap more than month away: 0
leap less than month away: 0
leap less than day away: 0
leap in less than 2 hours: 0
leap happened:           0
```

Field Descriptions

The following section describes the fields displayed using the **ntpleap** command.

sys.leap. The current status of the leap second monitor. There are four possible codes:

00	No warning.
01	Last minute has 61 seconds.
10	Last minute has 59 seconds.
11	Alarm condition (clock not synchronized)

leap.indicator. The number of leap seconds that occurred during the current day.

leap.warning. The number of leap seconds that will occur in the current month.

leap.bits. The number of leap bits set within the last hour.

time to next leap interrupt. A leap interrupt occurs when the NTP algorithm examines the topology of the network and determines if a leap second is needed (it may or may not be necessary at the time of the interrupt). This counter displays seconds until the next interrupt.

date of next leap interrupt. The time, in standard date notation, of the next leap interrupt after the most current leap interrupt is finished.

calls to leap process. The number of times a leap second has been added or subtracted.

leap more than month away. A scheduled leap second insertion more than a month away.

leap less than month away. A scheduled leap second insertion less than a month away.

leap less than day away. A scheduled leap second insertion less than a day away.

leap in less than 2 hours. A scheduled leap second insertion less than two hours away.

leap happened. The date of the last leap second insertion.

Turn the Server's Monitoring Facility On or Off

The Server Monitoring Facility keeps track of all NTP association for this switch. When it is On, it is possible to display a list of all NTP associations. For more information on displaying the Monitoring Facility list of NTP associations, see *Display Data The Server's Monitor Routines Have Collected* on page 12-31.

To turn the Monitoring Facility on or off, enter the **ntpmon** command as shown:

```
ntpmon <on:off>
```

where **<on:off>** is the status of the monitoring facility. For example, to turn the facility on, enter:

```
ntpmon on
```

Display Data The Server's Monitor Routines Have Collected

If the NTP monitoring facility is turned on, you can display a list of all known NTP associations with general information using the **ntpmlist** command.

To display a list of collected monitoring statistics, enter the **ntpmlist** command at the system prompt. A screen similar to the following is displayed:

remote address	port	local address	count	m	ver	drop	last	first
127.0.0.1	1025	127.0.0.1	1	7	3	0	0	0

This table is useful in establishing which entity is associated with the switch, and if entities have formed associations independent of administrator configuration (for example, if a user sets up an association with NTP without notifying the network administrator).

Field Descriptions

The following section describes the fields displayed using the **ntpmlist** command.

remote address. The IP address of the remote association.

port. The port the association was learned on and on which the association communicates with the switch.

◆ **Note** ◆

This is the TCP and UDP definition of a port, not a switch interface port.

local address. The local interface address for this association as created by the NTP configuration on the switch.

count. The number of NTP packets received from this association.

m. The mode the NTP associations uses in relation to the switch.

ver. The version of NTP the association is using (1,2, or 3)

drop. The number of NTP packets received from this association that were dropped (due to restrictions, bad packet data, etc.).

last. The number of seconds since the last NTP message was received from this association.

first. The number of seconds since the first NTP message was received from this association.

NTP Administration Menu

To view the NTP Administration Menu, enter the **ntadmin** command at the system prompt. If you are using verbose mode the NTP configuration menu is displayed. Otherwise, enter a question mark (?) at the prompt to display this menu:

Command	NTP Administration Menu
ntptimeo	set the primary receive time out
ntpdelay	set the delay added to encryption time stamps
ntphost	specify the host whose NTP server we talk to
ntpasswd	specify a password to use for authenticated requests
ntpkeyid	set keyid to use for authenticated requests
ntpkeytype	set key type to use for authenticated requests (des md5)
ntpdisable	clear a system flag (auth, bclient, monitor, stats)
ntpenable	set a system flag (auth, bclient, monitor, stats)

Related Menus:

Ntconfig Ntinfo Ntstats Ntadmin Ntaccess

The main menu options are shown in the **Related Menus** list for quick access if you need to change menus.

Set the Primary Receive Timeout

The **ntptimeo** command allows you to specify the number of milliseconds the server waits for a response to queries before the operation times out. The default is 8000 milliseconds. To change the timeout, enter the **ntptimeo** command as shown:

```
ntptimeo <value>
```

where **<value>** is the number of milliseconds of the new timeout length. For example, to set the timeout value to 3000 milliseconds, enter the following:

```
ntptimeo 3000
```

To view the current timeout setting with out changing it, enter the **ntptimeo** command with no value. A message similar to the following is shown:

```
primary timeout is 6000 ms
```

Set the Delay Added to Encryption Time Stamps

The **ntpdelay** command specifies a set time interval to add to timestamps included in server requests that require authentication. This can be used to enable server configuration over long delay network paths or between machines whose clocks are not synchronized.

To set the delay time, enter the **ntpdelay** command as shown:

```
ntpdelay <value>
```

where **<value>** is the number of milliseconds of the new delay time length. For example, to set the delay value to 30 milliseconds, enter the following:

```
ntpdelay 30
```

To view the current delay setting with out changing it, enter the **ntpdelay** command with no value. A message similar to the following is shown:

```
delay 30 ms
```

Specify the Host Whose NTP Server We Talk To

The **ntpghost** command specifies the name of the NTP server to which server queries are sent. This can be a domain name or an IP address. The default is localhost (the local server).

To change the NTP server for the switch, enter the **ntpghost** command as shown:

```
ntpghost <address>
```

where **<address>** is either the domain name or IP address of the NTP server. For example, to configure the switch to use an NTP server with an IP address of 1.1.1.4, enter:

```
ntpghost 1.1.1.4
```

To view the current NTP server used by the switch, enter the **ntpghost** command at the prompt with no address. A message similar to the following is shown:

```
current host is 1.1.1.4
```

Specify a Password to Use for Authenticated Requests

The **ntpasswd** command allows you to specify a password that must be entered when making configuration requests. The password must correspond to the key configured for use by the NTP server.

To specify a password:

1. Enter the **ntpasswd** command at the system prompt. A prompt displays asking for the Key ID number for the server, as shown:

```
Keyid:
```

Enter the key ID number for the server (as specified in the key file) and press **<return>**.

2. The following prompt appears requesting a password, as shown:

```
Password:
```

Enter the new password. This password is now required before making a configuration request of the server.

Set Key ID to Use for Authenticated Requests

The **ntpkeyid** command allows you to specify a key number to be used to authenticate configuration requests. This must correspond to the key number the server has been configured to use in the key file.

To set a new key ID, enter the **ntpkeyid** command as shown:

```
ntpkeyid <value>
```

where **<value>** is the new key ID number. For example, to set the key ID to 2, you would enter the following:

```
ntpkeyid 2
```

To view the currently configured key ID, enter the **ntpkeyid** command at the prompt and press **<return>**. A message similar to the following is shown:

```
keyid is 2
```


Set Key Type to Use for Authenticated Requests (DES|MD5)

NTP supports two types of encryption: DES or MD5. If you decide to use encryption to authenticate NTP information and configuration requests, you must specify which type of encryption to use.

To specify an encryption type enter the **ntpkeytype** command as shown:

```
ntpkeytype <value>
```

where **<value>** is either DES or MD5. For example, to set the key type to MD5, you would enter:

```
ntpkeytype MD5
```

To view the currently specified key type, enter the **ntpkeytype** command at the system prompt, and press **<return>**. A message similar to the following is displayed:

```
keytype is MD5
```

Set a System Flag (Auth, Bclient, Monitor, Stats)

The **ntpenable** command provides a way to enable various server options by creating flags added to NTP messages sent to the server.

To set a system flag, enter the **ntpenable** command as shown:

```
ntpenable <flag>
```

where **<flag>** is the type of flag the server will receive. There are six flag types that can be set:

auth	This flag causes the server to synchronize with unconfigured peers only if the peer has been correctly authenticated using a trusted key and key identifier. The default for this flag is disabled (off).
bclient	This flag causes the server to listen for a message from a broadcast or multicast server, following which an association is automatically instantiated for that server. The default for this flag is disabled (off).
monitor	This flag enables the monitoring facility. The default for this flag is disabled (off).
stats	This flag enables the statistics facility file generator. The default for this flag is enable (on).

When you have finished specifying a flag, press **<enter>**. A brief message appears to confirm the operation.

Clear a System Flag (Auth, Bclient, Monitor, Stats)

The **ntpdisable** command allows you to remove previously set flags from NTP messages sent to the server.

To disable a flag, enter the **ntpdisable** command as follows:

```
ntpdisable <flag>
```

where **<flag>** is the type of flag the server will receive. There are six flag types that can be set and removed. The flags are described in the section *Set a System Flag (Auth, Bclient, Monitor, Stats)* on page 12-35.

NTP Access Control Menu

To view the NTP Access Control Menu, enter the **ntaccess** command at the system prompt. If you are using verbose mode the NTP configuration menu is displayed. Otherwise, enter a question mark (?) at the prompt to display this menu:

Command	NTP Access Control Menu
ntpreqk	change the request message authentication keyid
ntpctlk	change the control message authentication keyid
ntpckey	add one or more key ID's to the trusted list
ntpvkey	display the trusted key ID list
ntpdkey	remove one or more key ID's from the trusted list
ntpauth	display the state of the authentication code
ntpcres	create restrict entry/add flags to entry
ntpvres	view the server's restrict list
ntpmres	remove flags from a restrict entry
ntpdres	delete a restrict entry
ntpctrap	configure a trap in the server
ntpvtrap	display the traps set in the server
ntpdtrap	remove a trap (configured or otherwise) from the server

Related Menus:
Ntconfig Ntinfo Ntstats Ntadmin Ntaccess

The main menu options are shown in the **Related Menus** list for quick access if you need to change menus.

Change the Request Message Authentication Key ID

There are two types of messages an NTP entity can send to another NTP entity: request and control. Request messages ask for information from the NTP entity such as timestamp information, statistics, etc. It is possible to change the authentication key identifier for request messages sent from the switch to another NTP entity.

To change the authentication key ID, enter the **ntpreqk** command as shown:

ntpreqk <value>

where **<value>** is the new key ID. Press **<return>**, and a brief message is displayed confirming the operation.

◆ **Note** ◆

The authentication key ID must match in both the switch sending the message and the switch receiving the message.

Change the Control Message Authentication Key ID

There are two types of messages an NTP entity can send to another NTP entity: request and control. Control messages attempt to change the configuration of the NTP entity in some fashion. It is possible to change the authentication key identifier for control messages sent from the switch to another NTP entity.

To change the authentication key ID, enter the **ntpctlk** command as shown:

```
ntpctlk <value>
```

where **<value>** is the new key ID. Press **<return>**, and a brief message is displayed confirming the operation.

◆ Note ◆

The authentication key ID must match in both the switch sending the message, and the switch receiving the message.

Add One or More Key ID's to the Trusted List

The trusted list in the key file is a list of all keys that are considered authentic and uncompromised. Messages from an NTP entity using one of these keys are accepted and acted upon. It is possible to add a key to the trusted list.

To add a key ID to the trust list in the key file, enter the **ntpckey** command as shown:

```
ntpckey <value>
```

where **<value>** is the new key ID to be added to the trusted list. For example, to add key ID 5 to the trusted list, enter the following:

```
ntpckey 5
```

A brief message is displayed confirming the operation.

◆ Note ◆

Adding a key ID using the **ntpckey** command adds the key to the working version of the key file in the switch's RAM. If you reset the switch or re-initialize NTP, the added key is lost.

Display the Trusted Key ID List

The trusted list in the key file is a list of all keys that are considered authentic and uncompromised. Messages from an NTP entity using one of these keys are accepted and acted upon.

To display a list of the trusted keys for this NTP client or server, enter the **ntpkey** command at the system prompt. A list of the key numbers accepted by this client or server is displayed. For more information on authentication, see *NTP and Authentication* on page 12-4.

Remove One or More Key ID's from the Trusted List

The trusted list in the key file is a list of all keys that are considered authentic and uncompromised. Messages from an NTP entity using one of these keys are accepted and acted upon. It is possible to remove a key from the trusted list.

To remove a key ID from the trusted list, enter the **ntpdkey** command as shown:

```
ntpdkey <value>
```

where **<value>** is the new key ID to be remove from the trusted list. For example, to remove key ID 5 from the trusted list, enter the following:

```
ntpdkey 5
```

A brief message is displayed confirming the operation.

◆ Note ◆

Removing a key ID using the **ntpdkey** command removes the key from the working version of the key file in the switch's RAM. If you reset the switch or re-initialize NTP, the removed key is reinstated.

Display the State of the Authentication Code

The **ntpauth** command allows you to look at the statistics of the authentication routine. These statistics consist of counters for various functions of the authentication code.

To view the statistics of the authentication code, enter the **ntpauth** command at the system prompt. A screen similar to the following is shown:

```
time since reset:      0
key lookups:          0
keys not found:       0
uncached keys:        0
encryptions:          0
decryptions:          0
```

Field Descriptions

The following sections explains the fields displayed using the **ntpauth** command.

time since reset. The number of seconds since the last restart of the switch.

key lookups. The number of times the switch has examined the key file to find a key.

keys not found. The number of times the switch failed to find a key in its key file.

uncached keys. The number of keys added to the key file using the **ntpdkey** command.

encryptions. The number of times the switch sent NTP messages or information out in encrypted form.

decryptions. The number of times the switch received NTP messages of information that was encrypted, and successfully decrypted the information.

Create Restrict Entry/Add Flags to Entry

It is possible to place restriction flags on specific NTP entities in relation to the switch. Restriction flags prevent messages or information coming from the NTP entity from affecting the switch.

To create a restriction flag, enter the **ntpcres** command as shown:

```
ntpcres <address> <mask> <restriction>
```

where **<address>** is the IP address of the NTP entity, **<mask>** is the entity's subnet mask, and **<restriction>** is the specific flag you want to place on the entity. For example to put an **ignore** restriction on an entity with address 1.1.1.1 and a subnet mask of 255.255.0.0, enter the following:

```
ntpcres 1.1.1.1 255.255.0.0 ignore
```

The following is a list of possible restriction flags that can be used:

ignore	Ignore all packets from hosts which match this entry. If this flag is specified neither queries nor time server polls will be responded to.
noquery	Ignore all NTP information queries and configuration requests from the source. Time service is not affected.
nomodify	Ignore all NTP information queries and configuration requests that attempt to modify the state of the server (i.e., run time reconfiguration). Queries which return information are permitted.
notrap	Decline to provide control message trap service to matching hosts. The trap service is a subsystem of the control message protocol which is intended for use by remote event logging programs.
lowpriotrap	Declare traps set by matching hosts to be low priority. The number of traps a server can maintain is limited (the current limit is 3). Traps are usually assigned on a first come, first serve basis, with later trap requestors being denied service. This flag modifies the assignment algorithm by allowing low priority traps to be overridden by later requests for normal priority traps. For more information on setting traps see <i>Configure a Trap in the Server</i> on page 12-41
noserve	Ignore NTP packets other than information queries and configuration requests. In effect, time service is denied, though queries may still be permitted.
nopeer	Provide stateless time service to polling hosts, but do not allocate peer memory resources to these hosts even if they otherwise might be considered useful as future synchronization partners.
notrust	Treat these hosts normally in other respects, but never use them as synchronization sources.

limited	These hosts are subject to a limitation of the number of clients from the same net. Net in this context refers to the IP notion of net (class A, class B, class C, etc.). Only the first client limit hosts that have shown up at the server and that have been active during the last client limit period (in seconds) are accepted. Requests from other clients from the same net are rejected. Only time request packets are taken into account. Query packets sent by the ntpq and xntpd programs are not subject to these limits. A history of clients is kept using the monitoring capability of xntpd. Thus, monitoring is always active as long as there is a restriction entry with the limited flag. For more information on enabling monitoring, see <i>Turn the Server's Monitoring Facility On or Off</i> on page 12-31.
ntpport	This is actually a match algorithm modifier, rather than a restriction flag. Its presence causes the restriction entry to be matched only if the source port in the packet is the standard NTP UDP port (123). Both ntpport and non-ntpport may be specified. The ntpport is considered more specific and is sorted later in the list.

View the Server's Restrict List

The **ntpvres** command allows you to view a list of all the configured restrictions for the switch. To view a list of configured restriction, enter the **ntpvres** command at the system prompt. A screen similar to the following appears:

address	mask	count	flags
0.0.0.0	0.0.0.0	12	none
127.0.0.1	255.255.255.255	0	ntpport, ignore

Field Descriptions

The following section describes the fields displayed with the **ntpvres** command.

address. The IP address of the NTP entity for which flags have been configured.

mask. The subnet mask of the NTP entity for which flags have been configured.

count. The number of NTP messages from the NTP entity that have been affected by the configured flags.

flags. The flags configured for this NTP entity. For a description of all possible flags, see *Create Restrict Entry/Add Flags to Entry* on page 12-39.

Remove Flags from a Restrict Entry

It is possible to place restriction flags on specific NTP entities in relation to the switch. Restriction flags prevent messages or information coming from the NTP entity from affecting the switch.

To remove a restriction flag from an NTP entity, enter the **ntpmres** command as shown:

```
ntpmres <address> <mask> <restriction>
```

where **<address>** is the IP address of the NTP entity, **<mask>** is the entity's subnet mask, and **<restriction>** is the specific flag you want to remove from the entity. For example, to remove an **ignore** restriction from an entity with address 1.1.1.1 and a subnet mask of 255.255.0.0, enter the following:

```
ntpmres 1.1.1.1 255.255.0.0 ignore
```

Delete a Restrict Entry

To remove an entry completely from the restriction list, enter the **ntpdres** command in the following manner:

```
ntpdres <address> <mask>
```

where **<address>** is the IP address of the NTP entity, and **<mask>** is the entity's subnet mask. For example to remove an entity with address 1.1.1.1 and a subnet mask of 255.255.0.0, enter the following:

```
ntpmres 1.1.1.1 255.255.0.0
```

This entity will no longer be listed in the restriction list and has no restriction flags placed on messages it sends to the switch.

Configure a Trap in the Server

The **ntpctrap** command allows you to set a trap receiver for the given address and port number. The trap receiver will log event messages and other information for the server in a log file.

To create a trap receiver, enter the **ntpctrap** command in the following manner:

```
ntpctrap <address> [<port>] [<interface>]
```

where address is the IP address of the switch. There are two optional items you can specify:

port The port on the switch used for sending NTP messages. If no port is specified, a default port of 18447 is used.

◆ Note ◆

This is the TCP and UDP definition of a port, not a switch interface port.

interface The local interface address for this NTP entity. If no interface is specified, the interface for the local NTP entity is used. For more information on interface addresses, see *Display Peer Summary Information* on page 12-16.

Display the Traps Set in the Server

The **ntpvttrap** command allows you to view a list of trap receivers set for the server. To view the trap list, enter the **ntpvttrap** command at the system prompt. A display similar to the following is shown:

```
address 127.0.0.1, port 18447
interface: 0.0.0.5, configured
set for 0 seconds, last set 0 seconds ago
sequence 1, number of resets 1
```

Field Descriptions

The following section describes the fields shown with the **ntpvttrap** command.

address. The address of the server where the trap was set.

port. The port on which the server is listening for NTP messages.

◆ Note ◆

This is the TCP and UDP definition of a port, not a switch interface port.

interface. The local interface address of the NTP server.

set for n seconds. The time the trap was initially set.

last set. The time in seconds from when the last trap was set for this server.

sequence. The number of times the trap was set.

number of resets. The number of times the trap has been reset.

Remove a Trap (Configured or Otherwise) from the Server

The **ntpdtrap** command allows you to remove a trap receiver for the given address. The trap receiver will log event messages and other information for the server in a log file.

To delete a trap receiver, enter the **ntpdtrap** command in the following manner:

```
ntpdtrap <address> [<port>] [<interface>]
```

where address is the IP address of the switch. There are two optional items you can specify:

port. The port on the switch used for sending NTP messages.

◆ Note ◆

This is the TCP/IP and UDP definition of a port, not a switch interface port.

interface. The local interface address for this NTP entity. For more information on interface addresses, see *Display Peer Summary Information* on page 12-16.

13 SNMP (Simple Network Management Protocol)

Introduction

Simple Network Management Protocol (SNMP) is an application layer protocol that allows network devices to exchange management information. SNMP works by sending messages, called protocol data units (PDUs), to network devices. Network administrators use SNMP to monitor network performance and to solve network problems.

An SNMP-managed network is comprised of three fundamental parts: agents, managed devices, and network management systems (NMSs). An agent, which resides within a managed device (i.e., a switch), is responsible for translating its local knowledge of management information into a form compatible with SNMP. When certain defined asynchronous events occur within a switch, the managed device sends traps, using the SNMP protocol, to a designated NMS. The NMS then views and monitors the switch's information through management software applications such as HP Open View or X-Vision.

SNMP parameters and traps are configurable through the **snmpc** command. For more information on this command, refer to *Configuring SNMP Parameters and Traps* on page 13-2. You can view SNMP statistics through the **snmps** command. For more information on this command, refer to *Viewing SNMP Statistics* on page 13-8. Both of these commands are also listed on the **Networking** menu.

Configuring SNMP Parameters and Traps

The **snmpc** command allows you to configure SNMP parameters and set traps that will be sent to network management stations. The **snmpc** command also enables you to add, modify, or delete SNMP parameters. The **snmpc** command is listed under the **Networking** menu. For more information about the networking menu, see Chapter 25, “IP Routing.” To configure SNMP parameters, enter the following command:

```
snmpc
```

A screen similar to the following displays:

```
SNMP current configuration:
```

```
1) Process SNMP Packets - enabled
2) Utilization Threshold - 60%
3) Set Community Name   - public
4) Get Community Name   - public
5) Trap Community Name  - public
6) Broadcast Traps      - disabled
7) 0 Unicast Traps      - disabled
```

```
(save/quit/cancel)
```

```
:
```

- To change a value, enter the number corresponding to that value, an equal sign (=), and the new value. For example, to enable broadcast traps, enter **5=enabled**.
- To clear an entry, specify the value as a period (.), as in **2=.** Note that true/false values and enabled/disabled values cannot be cleared.
- To save all your modifications, enter **save**.
- To cancel all your modifications, enter **Cancel** or **Ctrl-C**.
- To view the parameters currently configured, enter a question mark (?).

1) Process SNMP Packets

To enable or disable SNMP, enter 1, an equal sign (=), and the enable or disable command. The following is an example:

```
1=enable
```

2) Utilization Threshold

Utilization is the percentage of time that a resource is in use over a given period of time. Setting the Utilization Threshold places an upper limit on system utilization. To set this value, enter 2, an equal sign (=), and an integer between 1 and 99 to represent percentage of time in use. The default Utilization Threshold is 60%.

```
2=60%
```

3) Set Community Name

The Set Community Name variable is a password (up to 16 characters) that enables NMS stations to read and write objects through SNMP. The default Set Community Name is “public,” which allows all NMS stations read access to readable objects. If you want to specify a Set Community Name password, enter a **2**, an equal sign (=), and the new Set Community Name. The following is an example:

```
2=alpha
```

◆ **Note** ◆

Set Community Names with spaces must be enclosed in quotations (e.g., “test lab”).

4) Get Community Name

The Get Community Name variable is a password (up to 16 characters) that enables NMS stations to read objects defined in the MIBs. The default Get Community Name is “public,” which allows all NMS station read access to readable objects. If you want to specify a Get Community Name password, enter a **2**, an equal sign (=), and the new Get Community Name. The following is an example display:

```
2=beta
```

◆ **Note** ◆

Get Community Names with spaces must be enclosed in quotations (e.g., “data center”).

5) Trap Community Name

The Trap Community Name (up to 16 characters) is a password that enables NMS stations to collect traps (provided the NMS stations are configured with the same corresponding Trap Community Name). The default Trap Community Name is “public,” which allows the switch to send traps to all NMS stations configured with the Trap Community Name, “public.” If you want to specify a Trap Community Name password, enter a **4**, an equal sign (=), and the new Trap Community Name. The following is an example display.

```
4=trap1
```

◆ **Note** ◆

Trap Community Names with spaces must be enclosed in quotations (e.g., “trap 1”).

6) Broadcast Traps

When broadcast traps are enabled, the switch transmits traps to all NMS stations in the default group. If you enable this parameter, unicast traps (see option 6 below) will automatically be disabled. The default for broadcast traps is **disabled**. To enable broadcast traps, enter the following command:

```
5=enabled
```

The following prompt displays:

```
UDP destination port (162):
```

Enter the UDP destination port for the traps. UDP port 162 is the default port and is commonly used for traps; however, the destination port can be re-defined to accommodate a network management station using a nonstandard port.

◆ Note ◆

The destination port configured here must correspond to the UDP destination port configured at the receiving network management station(s).

7) Unicast Traps

When unicast traps are enabled, the switch transmits traps only to the IP address(es) defined in the **snmpc** list below this field.

◆ Note ◆

If both broadcast and unicast traps are disabled, then the switch does not transmit any traps.

If you enable this parameter, broadcast traps (see option 5 above) will automatically be disabled. The default for unicast traps is disabled. To enable unicast traps, enter the following command:

```
6=enabled
```

Configuring a New Network Management Station

- a. To define a new network management station, enter 8, followed by an equal sign (=), and the IP address of the network management station to receive traps. You can define a maximum of ten network management stations. They must be numbered sequentially from 8 through 17. If network management stations are already shown on the display for this menu, use the next highest number to add another station. The following is an example of how to define the first network management station:

```
8=123.12.1.1
```


The following prompt displays:

```
Enter trap mask words 0:1 (ffffff:ffffff):
```

Each trap in the switch is assigned a mask that consists of “words”. The mask value **ffffff:ffffff** indicates that *all* traps are enabled for words 0 and 1. If you want to accept this default (all traps enabled for words 0 and 1), press <Enter>. If you want to enable one or more specific traps for words 0 and 1, you must calculate their bit configurations and enter the new mask value at the prompt. Trap types and their bit positions are listed in the tables beginning on page 13-11.

Here is a sample configuration for setting a combination of traps.

Bit Configurations for Setting Traps



Example: To set a combination of trap types, add the hex values of the bits as follows:

Trap Type	Bit Settings	
	Word 0	Word 1
tempAlarm	00 00 00 00	: 00 00 00 01
risingAlarm	00 00 40 00	: 00 00 00 00
fallingAlarm	00 00 80 00	: 00 00 00 00
portPartitioned	00 00 00 00	: 00 00 02 00
Total =	00 00 C0 00	: 00 00 02 01

You would then enter the total mask value of the traps, as follows:

Enter trap mask words 0:1 (ffffff:ffffff): 0000C000:00000201

This setting would enable only these four traps for words 0 and 1.

- b. The following prompt displays:

Enter trap mask words 2:3 (ffffff:ffffff):

Enter the trap type(s) for words 2 and 3. If you want to accept the default (all traps enabled for words 2 and 3), press **<Enter>**. To set one or more specific traps, again calculate the bit configurations and enter the new mask value at the prompt.

- c. The following prompt displays:

Enter destination port (162):

Enter the UDP destination port for the traps configured above. If you choose the default in field four, port 162, press **<Enter>** at the prompt.

- d. The following prompt displays:

NMS state (on):

Indicate whether or not traps will be sent to this Network Management Station (the NMS defined in step a). If the NMS state is enabled (**on**), the NMS will be notified of traps. Press **<Enter>** to accept the default (**on**). If the NMS state is disabled (**off**), the NMS will not be notified of traps.

- e. The following prompt displays:

Special Access? (no): yes

Select whether or not this Network Management Station has special access. If you enter **yes**, this NMS will have administrative privileges such as modifying, deleting, or adding to other trap entries as well as its own. Without special access, an NMS can only update its own entry. If you choose the default, **no**, simply press **<Enter>** at the prompt.

Save your configuration by typing **save** and then **<Enter>**.

- f. After you have saved your configuration, the prompt re-displays. The above entries will create an NMS number 8 in the list. Traps will be sent to the IP address specified for that NMS station (provided the NMS state is **on** and unicast traps are **enabled**).

To view your new SNMP configuration, enter the **snmpc** command. The following is a sample display of the output from the **snmpc** command after the above sample configuration:

SNMP current configuration:

```
1) Process SNMP Packets - enabled
2) Utilization Threshold - 60%
3) Set Community Name - admin
4) Get Community Name - public
5) Trap Community Name - trap1
6) Broadcast Traps - disabled
7) 1 Unicast Traps - enabled
8) NMS IP address - 123.12.1.1 /162 --bfffffff:ffffffff (on) (SA)
-- ffffffff:ffffffff
```

(save/quit/cancel)

:

The values that appear to the immediate right of the NMS IP address are: the UDP destination port number (**162**), the trap bit masks (**fffffff:ffffffff**), the functional state of the NMS (**on**), and the special access (**SA**) status (this does not appear if you selected **no** for special access in step above).

To add network management stations to this current SNMP configuration, enter the next highest entry number from the last defined NMS. For example, if you wanted to add another NMS to the above sample configuration, you would enter the following:

9=123.22.2.2

Please note that any additional NMS entries must have a unique IP address. Repeat steps **b** through **f** to continue configuring additional NMS entries. Once you save your configuration and re-enter the **snmpc** command at the prompt, the screen refreshes to include the new NMS entry. The following is a sample display:

```
SNMP current configuration:
1) Process SNMP Packets - enabled
2) Utilization Threshold - 60%
3) Set Community Name - public
4) Get Community Name - public
5) Trap Community Name - public
6) Broadcast Traps - disabled
7) 1 Unicast Traps - enabled
8) NMS IP address - 123.12.1.1 /162 -- ffffffff:bfffffff (on) (SA
-- ffffffff:ffffffff)
9) NMS IP address - 123.22.2.2 /162 -- ffffffff:ffffffff (on)
-- ffffffff:ffffffff
(save/quit/cancel)
:
```

- g.** To delete an IP address added to this list, enter the NMS index number of the entry followed by the decimal (.) character. The following example would delete the NMS IP address listed at number **9**.

```
9=.
```

Viewing SNMP Statistics

The **snmps** command is used to display SNMP statistics. The command displays the SNMP activities since the last time the switch was powered on, or since the last Reset was executed. It also displays a list of the current traps.

The **snmps** command is listed on the **Networking** menu. For more information about the networking menu, see Chapter 25, "IP Routing." To display SNMP statistics, enter the following command:

```
snmps
```

A screen similar to the following displays:

SNMP Statistics		
	In	Out
Total Packets	67	67
Bad Versions	0	
Bad Community Names	0	
Bad Community Use	0	
Bad Type Discards	0	
ASN Parse Errors	0	
Too Big Errors	0	0
No Such Name Errors	0	1
Bad Value Errors	0	0
Read Only Errors	0	0
General Errors	0	0
Total Variable Requests	186	
Total Set Variable Requests	0	
Get Requests	17	0
Get Next Requests	50	0
Set Requests	0	0
Get Responses	0	67
Authentication Trap Enables:	0	
Traps	0	0

Trap generation is ENABLED to these management stations:

```
198.206.1.1 /162 -- ffffffff:bfffffff (on)
198.2.1.1   /162 -- ffffffff:7fffffff (off) (SA)
```

Total Packets

The total number of packets received and sent.

Bad Versions

The total number of SNMP messages delivered to the switch SNMP protocol entity that were for an unsupported SNMP version.

Bad Community Names

The total number of SNMP message names delivered to the switch SNMP protocol entity that used an unknown SNMP community name.

Bad Community Use

The total number of SNMP messages delivered to the SNMP protocol entity which represented an SNMP operation that was not allowed by the SNMP community named in the message.

Bad Type Discards

The total number of SNMP entries discarded because the request type was not recognized.

ASN Parse Errors

The total number of ASN.1 or BER errors encountered by the SNMP protocols entity when decoding received SNMP Messages.

Too Big Errors

The total number of SNMP PDUs delivered to the SNMP protocol entity with a value in the error-status field of 'tooBig'.

No Such Name Error

The total number of SNMP PDUs delivered to the SNMP protocol entity with value in the error-status field of 'noSuchName'.

Bad Value Errors

The total number of valid SNMP PDUs delivered to the SNMP protocol entity with a value in the error-status field of 'readOnly.' It is a protocol error to generate an SNMP PDU that contains the value 'readOnly' in the error-status field; as such this object is provided as a means of detecting incorrect implementations of the SNMP.

Read Only Errors

The total number of valid SNMP PDUs delivered to the SNMP protocol entity for with an error-status field value of 'Read Only'.

General Errors

The total number of SNMP PDUs delivered to the switch SNMP protocol entity with an error-status field value of 'GenError'.

Total Variable Requests

The total number of MIB objects from which Requests have been retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.

Total Set Variable Requests

The total number of MIB objects from which Requests have been retrieved successfully by the SNMP entity as the result of receiving valid SNMP Set-Request PDUs.

Get Requests

The total number of SNMP Get-Request PDUs accepted and processed by the switch SNMP protocol entity.

Get Next Requests

The total number of SNMP Get-Next PDUs accepted and processed by the switch SNMP protocol entity.

Set Requests

The total number of SNMP Set-Request PDUs which have been accepted and processed by the switch SNMP protocol entity.

Get Responses

The total number of SNMP Response PDUs accepted and processed by the switch SNMP protocol entity.

Authentication Trap Enables

Indicates whether the SNMP agent Enable process is permitted to generate authentication-failure traps. The value of this object overrides any configuration information, providing a means to enable all authentication-failure traps.

Traps

The number of SNMP Trap PDUs generated by the SNMP protocol entity. Traps are broadcast only.

Traps are broadcast only

This appears if traps are set to broadcast. The address is the broadcast address of the default VLAN of AutoTracker group 1.

Trap generation is ENABLED to these management stations

This appears if you have used the **snmpc** command to set up one or more management stations to receive traps. The trap tables on the following pages list the traps that are currently supported.

Trap Tables

The following table is a summary list of the supported SNMP traps and their values.

<i>Trap or Mask Name</i>	<i>Object ID</i>	<i>Bit Position</i>	<i>Hex Value</i>	<i>Page</i>
coldStart	1.3.6.1.2.1.11.0	(word 0) 0	(word 0) 1	13-15
warmStart	1.3.6.1.2.1.11.1	(word 0) 1	(word 0) 2	13-16
linkDown	1.3.6.1.2.1.11.2	(word 0) 2	(word 0) 4	13-16
linkUp	1.3.6.1.2.1.11.3	(word 0) 3	(word 0) 8	13-17
authentication failure	1.3.6.1.2.1.11.4	(word 0) 4	(word 0) 10	13-17
egpNeighborLoss	1.3.6.1.2.1.11.5	(word 0) 5	(word 0) 20	13-18
frDLCIStatusChange	1.3.6.1.2.1.11.7	(word 0) 7	(word 0) 80	13-18
ipxTrapCircuitDown	1.3.6.1.4.1.23.2.5.5.1	(word 0) 8	(word 0) 100	13-19
ipxTrapCircuitUp	1.3.6.1.4.1.23.2.5.5.2	(word 0) 9	(word 0) 200	13-19
newRoot	1.3.6.1.2.17.0.1	(word 0) 10	(word 0) 400	13-19
topologyChange	1.3.6.1.2.17.0.2	(word 0) 11	(word 0) 800	13-20
atmfVpcChange	1.3.6.1.4.1.353.0.1	(word 0) 12	(word 0) 1000	13-21
atmfVccChange	1.3.6.1.4.1.353.0.2	(word 0) 13	(word 0) 2000	13-22
rising Alarm	1.3.6.1.2.16.0.1	(word 0) 14	(word 0) 4000	13-23
falling Alarm	1.3.6.1.2.16.0.2	(word 0) 15	(word 0) 8000	13-24
dsx3LineStatusChange	1.3.6.1.2.1.10.20.15.0.1	(word 0) 16	(word 1) 1 0000	13-25
dsx1LineStatusChange	1.3.6.1.2.1.10.18.15.0.1	(word 0) 17	(word 1) 2 0000	13-26
MPLS_LDP_THRESHOLD_MASK *		(word 0) 18	(word 0) 4 0000	
POS3_STAT_CHANGE_MASK *		(word 0) 19	(word 0) 8 0000	
IMA_FAILURE_ALARM_MASK *		(word 0) 20	(word 0) 10 0000	
SYSLOG_TRAP_MASK *		(word 0) 29	(word 0) 2000 0000	
NMS_MASTER_MASK *		(word 0) 30	(word 0) 4000 0000	
NMS_TRAP_DISABLE_MASK *		(word 0) 31	(word 0) 8000 0000	
* This mask name does not necessarily match the trap name.				

Trap Tables

<i>Trap or Mask Name</i>	<i>Object ID</i>	<i>Bit Position</i>	<i>Hex Value</i>	<i>Page</i>
tempAlarm	1.3.6.1.4.1.800.3.1.1.4.0.1	(word 1) 0	(word 1) 1	13-27
moduleChange	1.3.6.1.4.1.800.3.1.1.4.0.2	(word 1) 1	(word 1) 2	13-28
powerEvent	1.3.6.1.4.1.800.3.1.1.4.0.3	(word 1) 2	(word 1) 4	13-29
controllerEvent	1.3.6.1.4.1.800.3.1.1.4.0.4	(word 1) 3	(word 1) 8	13-30
loginViolation	1.3.6.1.4.1.800.3.1.1.4.0.5	(word 1) 4	(word 1) 10	13-31
macVlanViolation	1.3.6.1.4.1.800.3.1.1.4.0.6	(word 1) 5	(word 1) 20	13-31
macDuplicatePort	1.3.6.1.4.1.800.3.1.1.4.0.7	(word 1) 6	(word 1) 40	13-32
portLinkUpEvent	1.3.6.1.4.1.800.3.1.1.4.0.8	(word 1) 7	(word 1) 80	13-33
portLinkDownEvent	1.3.6.1.4.1.800.3.1.1.4.0.9	(word 1) 8	(word 1) 100	13-34
portPartitioned	1.3.6.1.4.1.800.3.1.1.4.0.10	(word 1) 9	(word 1) 200	13-35
portRecordMismatch	1.3.6.1.4.1.800.3.1.1.4.0.11	(word 1) 10	(word 1) 400	13-36
groupChange	1.3.6.1.4.1.800.3.1.1.4.0.14	(word 1) 13	(word 1) 2000	13-37
vlanChange	1.3.6.1.4.1.800.3.1.1.4.0.15	(word 1) 14	(word 1) 4000	13-38
portMove	1.3.6.1.4.1.800.3.1.1.4.0.16	(word 1) 15	(word 1) 8000	13-39
moduleResetReload	1.3.6.1.4.1.800.3.1.1.4.0.17	(word 1) 16	(word 1) 1 0000	13-40
systemEvent	1.3.6.1.4.1.800.3.1.1.4.0.18	(word 1) 17	(word 1) 2 0000	13-41
vlanRouteTableFull	1.3.6.1.4.1.800.3.1.1.4.0.19	(word 1) 18	(word 1) 4 0000	13-42
sapTableFull	1.3.6.1.4.1.800.3.1.1.4.0.20	(word 1) 19	(word 1) 8 0000	13-42
atmSSCOPstate	1.3.6.1.4.1.800.3.1.1.4.0.21	(word 1) 20	(word 1) 10 0000	13-43
ilmiState	1.3.6.1.4.1.800.3.1.1.4.0.22	(word 1) 21	(word 1) 20 0000	13-43
atmConnection	1.3.6.1.4.1.800.3.1.1.4.0.23	(word 1) 22	(word 1) 40 0000	13-44
atmService	1.3.6.1.4.1.800.3.1.1.4.0.24	(word 1) 23	(word 1) 80 0000	13-45
dldciNew	1.3.6.1.4.1.800.3.1.1.4.0.27	(word 1) 26	(word 1) 400 0000	13-46
dldciDel	1.3.6.1.4.1.800.3.1.1.4.0.28	(word 1) 27	(word 1) 800 0000	13-47
dldciUp	1.3.6.1.4.1.800.3.1.1.4.0.29	(word 1) 28	(word 1) 1000 0000	13-48
dldciDn	1.3.6.1.4.1.800.3.1.1.4.0.30	(word 1) 29	(word 1) 2000 0000	13-49
portManualForwarding Mode	1.3.6.1.4.1.800.3.1.1.4.0.31	(word 1) 30	(word 1) 4000 0000	13-50
fdciCFStateChange	1.3.6.1.4.1.800.3.1.1.4.0.32	(word 1) 31	(word 1) 8000 0000	13-51
duplicateIPAddress	1.3.6.1.4.1.800.3.1.1.4.0.35	(word 2) 2	(word 2) 4	13-52
duplicateMACaddress	1.3.6.1.4.1.800.3.1.1.4.0.36	(word 2) 3	(word 2) 8	13-53

<i>Trap or Mask Name</i>	<i>Object ID</i>	<i>Bit Position</i>	<i>Hex Value</i>	<i>Page</i>
healthThresholdRising	1.3.6.1.4.1.800.3.1.1.4.0.37	(word 2) 4	(word 2) 10	13-54
healthThresholdFalling	1.3.6.1.4.1.800.3.1.1.4.0.38	(word 2) 5	(word 2) 20	13-54
healthThresholdDevice	1.3.6.1.4.1.800.3.1.1.4.0.39	(word 2) 6	(word 2) 40	13-55
healthThresholdModule	1.3.6.1.4.1.800.3.1.1.4.0.40	(word 2) 7	(word 2) 80	13-55
xylanXIPXMAPPort StatusChange	1.3.6.1.4.1.800.3.1.1.4.0.41	(word 2) 8	(word 2) 100	13-56
xylanSIPXMAPPortState Change	1.3.6.1.4.1.800.3.1.1.4.0.42	(word 2) 9	(word 2) 200	13-57
clkBusLineStateChange	1.3.6.1.4.1.800.3.1.1.4.0.45	(word 2) 10	(word 2) 400	13-60
xylanXIPGMAPFailed Update	1.3.6.1.4.1.800.3.1.1.4.0.44	(word 2) 11	(word 2) 800	13-59
avlAuthAttempt	1.3.6.1.4.1.800.3.1.1.4.0.43	(word 2) 16	(word 2) 1 0000	13-58
mcpStatisticsOverflow	1.3.6.1.4.1.800.3.1.1.4.0.67	(word 2) 18	(word 2) 4 0000	13-62
mcpShortCut	1.3.6.1.4.1.800.3.1.1.4.0.68	(word 2) 19	(word 2) 8 0000	13-66
mcpIngressRetryTime	1.3.6.1.4.1.800.3.1.1.4.0.69	(word 2) 20	(word 2) 10 0000	13-67
vrrpTrapNewMasterOut	1.3.6.1.2.1.46.1.3.1.0.3	(word 2) 21	(word 2) 20 0000	13-68
vrrpAuthFailure	1.3.6.1.2.1.46.1.3.1.0.4	(word 2) 22	(word 2) 40 0000	13-69
blind-violation	1.3.6.1.4.1.800.3.1.1.1.0.46	(word 2) 23	(word 2) 80 0000	13-61
mpcStatisticsOverflow	1.3.6.1.4.1.800.3.1.1.1.0.47	(word 2) 18	(word 2) 4 0000	13-62
fddiLerFlagChange	1.3.6.1.4.1.800.3.1.1.4.0.65	(word 3) 0	(word 3) 1	13-63
fddiCLTFailCntIncr	1.3.6.1.4.1.800.3.1.1.4.0.66	(word 3) 1	(word 3) 2	13-64
oamVCAIS	1.3.6.1.4.1.800.3.1.1.4.0.71	(word 3) 10	(word 3) 400	13-70
oamVCRDI	1.3.6.1.4.1.800.3.1.1.4.0.72	(word 3) 11	(word 3) 800	13-71
oamVCLOC	1.3.6.1.4.1.800.3.1.1.4.0.73	(word 3) 12	(word 3) 1000	13-72
oamVCUnsuccessLoop	1.3.6.1.4.1.800.3.1.1.4.0.74	(word 3) 13	(word 3) 2000	13-73
oamVPAIS	1.3.6.1.4.1.800.3.1.1.4.0.75	(word 3) 14	(word 3) 4000	13-74
oamVPRDI	1.3.6.1.4.1.800.3.1.1.4.0.76	(word 3) 15	(word 3) 8000	13-75
oamVPLOC	1.3.6.1.4.1.800.3.1.1.4.0.77	(word 3) 16	(word 3) 1 0000	13-76
oamVPUnsuccessLoop	1.3.6.1.4.1.800.3.1.1.4.0.78	(word 3) 17	(word 3) 2 0000	13-77
accountEvent	1.3.6.1.4.1.800.3.1.1.4.0.86	(word 3) 21	(word 3) 20 0000	13-78
Over1Alarm	1.3.6.1.4.1.800.3.1.1.4.0.87	(word 3) 22	(word 3) 40 0000	13-78

Trap Tables

<i>Trap or Mask Name</i>	<i>Object ID</i>	<i>Bit Position</i>	<i>Hex Value</i>	<i>Page</i>
Under1Event	1.3.6.1.4.1.800.3.1.1.4.0.88	(word 3) 23	(word 3) 80 0000	13-79
Over2Alarm	1.3.6.1.4.1.800.3.1.1.4.0.89	(word 3) 24	(word 3) 100 0000	13-79
Under2Event	1.3.6.1.4.1.800.3.1.1.4.0.90	(word 3) 25	(word 3) 200 0000	13-80
Over3Alarm	1.3.6.1.4.1.800.3.1.1.4.0.91	(word 3) 26	(word 3) 400 0000	13-80
Under3Event	1.3.6.1.4.1.800.3.1.1.4.0.92	(word 3) 27	(word 3) 8000 0000	13-81
NoDeviceAlarm	1.3.6.1.4.1.800.3.1.1.4.0.93	(word 3) 28	(word 3) 1000 0000	13-81
FileAlarm	1.3.6.1.4.1.800.3.1.1.4.0.94	(word 3) 29	(word 3) 2000 0000	13-82
ldpPeerCreate	1.3.6.1.4.1.800.3.1.1.4.0.80	(word 3) 5	(word 3) 20	13-83
ldpPeerDelete	1.3.6.1.4.1.800.3.1.1.4.0.81	(word 3) 6	(word 3) 40	13-84
ldpSessionCreate	1.3.6.1.4.1.800.3.1.1.4.0.82	(word 3) 17	(word 3) 80	13-85
ldpSessionDelete	1.3.6.1.4.1.800.3.1.1.4.0.83	(word 3) 8	(word 3) 100	13-86
lecStateChangeEvent	1.3.6.1.4.1.800.3.1.1.4.0.96	(word 2) 26	(word 2) 40 0000	13-87

SNMP Standard Traps

This section lists the standard traps that are defined within RFC (MIB) documents. These traps signify events as they occur on common network devices. The following information on traps is provided in the tables.

Trap. The object name of the trap as it is defined in the corresponding MIB (Management Information Base). Alcatel supports standardized and proprietary MIBS.

Object ID. The SNMP object identifier (OID) for this trap.

Description. A brief explanation describing the circumstances under which a specific trap is generated.

Bit Position. The trap's specific position in a bit mask (a bit mask is a binary notation which represents a combination of all four trap words). By mapping a specific trap to its binary position, you can determine whether or not a trap is enabled. For example, a trap is enabled if its corresponding bit is set to 1 and disabled if its corresponding bit is set to 0.

Word. A word is a set of four consecutive bytes within a system's memory. Alcatel allocates a total of four words for trap representation. Each of the 32 bit positions within a word corresponds to a specific trap. The first word, Word 0, contains only standard traps as they are defined within RFC (MIB) documents. Words 1, 2, and 3 contain Alcatel-specific traps.

Hex Value. The resulting hexadecimal value of the bit mask.

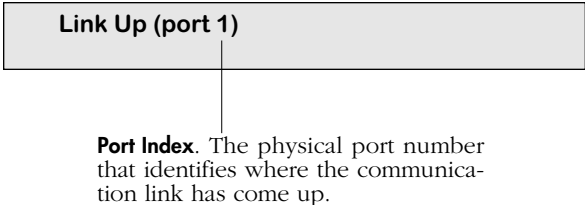
Trap Text and Variable Description. Trap text is a brief statement containing additional information that can help you narrow down the source of the trap, such as slot/port numbers, module types, and MAC addresses (variable descriptions have been added for your convenience). When a specific trap is triggered, it may display in various text formats, depending on the software application through which it is viewed. The trap text in the following tables are examples of trap text displayed through the HP OpenView Alarm Log and the Traps window in X-Vision Discovery. For more information on X-Vision, see the on-line documentation included with the application.

Trap	coldStart
Object ID	1.3.6.1.2.1.11.0
Description	The sending protocol entity is re-initializing itself such that the agent's configuration or the protocol entity implementation may be altered.
Bit Position (Word 0)	0
Hex Value (Word 0)	1
Trap Text and Variable Descriptions	Cold Start

Trap Tables

Trap	warmStart
Object ID	1.3.6.1.2.1.11.1
Description	The sending protocol entity is re-initializing itself such that neither the agent's configuration nor the protocol entity implementation may be altered.
Bit Position (Word 0)	1
Hex Value (Word 0)	2
Trap Text and Variable Descriptions	Warm Start

Trap	linkDown
Object ID	1.3.6.1.2.1.11.2
Description	The sending protocol entity recognizes a failure in one of the communication links represented in the agent's configuration.
Bit Position (Word 0)	2
Hex Value (Word 0)	4
Trap Text and Variable Descriptions	<div data-bbox="391 1354 972 1423" style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; display: inline-block; margin-bottom: 10px;">Link Down (port 1)</div> <p>Port Index. The physical port number that identifies the failed communication link.</p>

Trap	linkUp
Object ID	1.3.6.1.2.1.11.3
Description	The sending protocol entity recognizes that one of the communication links represented in the agent's configuration has come up.
Bit Position (Word 0)	3
Hex Value (Word 0)	8
Trap Text and Variable Descriptions	 <p>Port Index. The physical port number that identifies where the communication link has come up.</p>

Trap	authenticationFailure
Object ID	1.3.6.1.2.1.11.4
Description	The sending protocol entity is the addressee of a protocol message that is not properly authenticated.
Bit Position (Word 0)	4
Hex Value (Word 0)	10
Trap Text and Variable Descriptions	Authentication Failure

Trap	egpNeighborLoss
Object ID	1.3.6.1.2.1.11.5
Description	An EGP neighbor for whom the sending protocol entity was an EGP peer has been marked down and the peer relationship no longer exists.
Bit Position (Word 0)	5
Hex Value (Word 0)	20
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;">Neighbor Loss (neigh addr 192.168.10.1)</div> <p style="text-align: center; margin-top: 10px;">Neighbor IP Address. The IP address of this entry's EGP neighbor.</p>

Trap	frDLCIStatusChange
Object ID	1.3.6.1.2.1.11.6
Description	This trap is sent when the indicated virtual circuit has changed state. It has either been created or invalidated, or has toggled between the active and inactive states. However, if the reason for the state change is due to the DLCMI going down, traps should not be generated for each DLCI.
Bit Position (Word 0)	7
Hex Value (Word 0)	80
Variable Description	<p>frCircuitIfIndex - The ifIndex value of the ifEntry this virtual circuit is layered into.</p> <p>frCircuitDlci - The DLCI for this virtual circuit.</p> <p>frCircuitState - Indicates whether this virtual circuit is active or inactive.</p>

Trap	ipxTrapCircuitDown
Object ID	1.3.6.1.4.1.23.2.5.5.1
Description	This trap indicates that the specified circuit has gone down.
Bit Position (Word 0)	8
Hex Value (Word 0)	100
Variable Description	ipxCircSysInstance - The identifier of this instance of IPX. ipxCircIndex - The identifier of this circuit, for this instance of IPX.

Trap	ipxTrapCircuitUp
Object ID	1.3.6.1.4.1.23.2.5.5.2
Description	This trap indicates that the specified circuit has come up.
Bit Position (Word 0)	9
Hex Value (Word 0)	200
Variable Description	ipxCircSysInstance - The identifier of this instance of IPX. ipxCircIndex - The identifier of this circuit, for this instance of IPX.

Trap Type	newRoot
Object ID	1.3.6.1.2.1.17.0.1
Description	Sent by a bridge that became the new root of the Spanning Tree.
Bit Position (Word 0)	10
Hex Value (Word 0)	400
Trap Text and Variable Descriptions	Spanning Tree: A new agent has become the root of the Spanning Tree.

Trap Tables

Trap	topologyChange
Object ID	1.3.6.1.2.1.17.0.2
Description	A bridge's configured ports either transitioned from Learning state to Forwarding state or from Forwarding state to Blocking state. This trap will not be sent if a newRoot trap was sent for the same transition.
Bit Position (Word 0)	11
Hex Value (Word 0)	800
Trap Text and Variable Descriptions	Spanning Tree: A configured port's state has transitioned.

Trap	atmfVpcChange
Object ID	1.3.6.1.4.1.353.0.1
Description	Either a permanent VPC was added or deleted at this ATM interface, or an existing VPC was modified.
Bit Position (Word 0)	12
Hex Value (Word 0)	1000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>A permanent VPC has been added or deleted at this ATM Interface, or the attributes of an existing VPC have been modified (index 0, Vpi 2, Status 3)</p> </div> <p>Port Index. The port number of this ATM interface. Valid values range from 0 to 2147483647.</p> <p>VPI. The Virtual Path Identifier at this ATM interface. Valid values range from 0 to 4095.</p> <p>Operational Status. The present operating status of the VPC. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 unknown 2 end2endUp 3 end2endDown 4 localUpEnd2endUnknown 5 localDown

Trap	atmfVccChange
Object ID	1.3.6.1.4.1.353.0.2
Description	Either a permanent VCC was added or deleted at this ATM interface, or an existing VCC was modified.
Bit Position (Word 0)	13
Hex Value (Word 0)	2000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>A permanent VCC has been added or deleted at this ATM Interface, or the attributes of an existing VPC have been modified (index 0, Vpi 2, Vci 6, status 3)</p> </div> <p>Operational Status. The present operational status of the VCC. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 unknown 2 end2endUp 3 end2endDown 4 localUpEnd2endUnknown 5 localDown <p>Port Index. The port number which identifies this ATM interface. Valid values range from 0 to 2147483647.</p> <p>VPI. The Virtual Channel Identifier at this ATM interface. Valid values range from 0 to 4095. For virtual interfaces, this value has no meaning and is set to zero.</p> <p>VCI. The Virtual Channel Identifier at this ATM interface. Valid values range from 0 to 65535. For virtual interfaces, this value has no meaning and is set to zero.</p>

Trap	risingAlarm
Object ID	1.3.6.1.2.1.16.0.1
Description	The value of an Ethernet statistical variable (i.e., a member of the Ethernet statistics group as defined by RFC 1757) has exceeded its rising threshold. The variable's rising threshold and whether it will generate an SNMP trap for this condition are configured by a network management station running RMON.
Bit Position (Word 0)	14
Hex Value (Word 0)	4000
Trap Text and Variable Descriptions	<p>Variable. The MIB object identifier for the variable being sampled.</p> <p>Alarm Index. An index value for this entry in the alarm table. Each entry defines a diagnostic sample at a particular interval for an object on the device.</p> <p>An RMON alarm entry crossed its rising threshold (index 25 var 2 type 1 value 201 rising threshold 200)</p> <p>Value. The value of the statistic during the last sampling period. For example, if the sample method is Delta Value, this value will be the difference between the samples at the beginning and end of the period. If the sample method is Absolute Value, this value will be the sampled value at the end of the period. This is the value that is compared with the rising threshold.</p> <p>Rising Threshold. A threshold for the sampled statistic. This trap is generated when the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold.</p> <p>After a rising event is generated, another such event will not be generated until the sampled value falls below this threshold and reaches the Falling Threshold value.</p> <p>Sampling Method. The method of sampling the selected variable and calculating the value for comparison with the thresholds. Possible values are integers 1 and 2:</p> <ol style="list-style-type: none"> 1 Absolute Value. The value of the selected variable will be compared directly with the thresholds at the end of the sampling interval. 2 Delta Value. The value of the selected variable at the last sample will be subtracted from the current value, and the difference compared with the thresholds.

Trap	fallingAlarm
Object ID	1.3.6.1.2.1.16.0.2
Description	The value of an Ethernet statistical variable (i.e., a member of the Ethernet statistics group as defined by RFC 1757) has dipped below its falling threshold. The variable's falling threshold and whether it will generate an SNMP trap for this condition are configured by a network management station running RMON.
Bit Position (Word 0)	15
Hex Value (Word 0)	8000
Trap Text and Variable Descriptions	<p>Variable. The MIB object identifier for the variable being sampled.</p> <p>Alarm Index. An index value for this entry in the alarm table. Each entry defines a diagnostic sample at a particular interval for an object on the device.</p> <div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; text-align: center;"> <p>An RMON alarm entry crossed its falling threshold (index 25 var 2 type 1 value 100 falling threshold 9)</p> </div> <p>Value. The value of the statistic during the last sampling period. For example, if the sample method is Delta Value, this value will be the difference between the samples at the beginning and end of the period. If the sample method is Absolute Value, this value will be the sampled value at the end of the period. This is the value that is compared with the falling threshold.</p> <p>Sampling Method. The method of sampling the selected variable and calculating the value for comparison with the thresholds. Possible values are:</p> <ol style="list-style-type: none"> 1 Absolute Value. The value of the selected variable will be compared directly with the thresholds at the end of the sampling interval. 2 Delta Value. The value of the selected variable at the last sample will be subtracted from the current value, and the difference compared with the thresholds. <p>Falling Threshold. A threshold for the sampled statistic. This trap is generated when the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was more than this threshold.</p> <p>After a falling event is generated, another such event will not be generated until the sampled value rises above this threshold and reaches the Rising Threshold value.</p>

Trap Type	dsx3LineStatusChange
Object ID	1.3.6.1.2.1.10.30.15.0.1
Description	The value of an instance dsx3LineStatus changed.
Bit Position (Word 0)	16
Hex Value (Word 1)	1 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px; text-align: center;"> Line Status Change (line status 1, last change 4) </div> <p>DSX3 Line Status. The line status of the interface. It contains loopback, failure, received alarm, and transmitted alarm information. Valid values range from 1 to 8191.</p> <p>Last Change. The last value of MIB II's sysUpTime object at the time this DS3 entered its current line status state. If the current state was entered prior to the last re-initialization of the proxy-agent, this value is zero.</p>

Trap	dsx1LineStatusChange
Object ID	1.3.6.1.2.1.10.18.15.0.1
Description	The value of an instance dsx1LineStatus changed.
Bit Position (Word 0)	17
Hex Value (Word 1)	2 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px; text-align: center;"> Line Status Change (line status 1, last change 2) </div> <p>DSX1 Line Status. The line status of the interface. It contains loopback, failure, received alarm, and transmitted alarm information. Valid values range from 1 to 8191.</p> <p>Last Change. The last value of MIB II's sysUpTime object at the time this DS1 entered its current line status state. If the current state was entered prior to the last re-initialization of the proxy-agent, this value is zero.</p>

Extended Traps

This section lists Alcatel-specific traps. These extended traps are generated specifically by Alcatel switch devices.

Trap Type	tempAlarm
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.1
Description	The temperature sensor(s) have detected a temperature in the chassis that exceeds the threshold. These sensors are physically located on the MPX module, but can detect temperature changes throughout the chassis.
Bit Position (Word 1)	0
Hex Value (Word 1)	1
Trap Text and Variable Descriptions	Temperature Sensor has changed state to Over Threshold

Trap Type	moduleChange		
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.2		
Description	A module was either inserted or removed from the chassis. In some cases, this trap may also be generated when a module is reset.		
Bit Position (Word 1)	1		
Hex Value (Word 1)	2		
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p align="center">Module was inserted or removed from chassis (slot 4, subunit 1, type 10)</p> </div> <p>Slot Number. The slot number on the front of the chassis where this module was inserted or removed.</p> <p>Submodule Type. Indicates the submodule that was inserted or removed. Typically this value will be 1, meaning the base module was inserted or removed. If this value is 2, then HSX module 1 was moved. If this value is 3, then HSX module 2 was moved.</p> <p>Module Type. Indicates the module type that was inserted or removed. The following integers are valid values:</p> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; vertical-align: top;"> <ul style="list-style-type: none"> 4 HSM 5 MPM 6 ESM 8-port 10BASE-T 7 ESM 16-port 8 TSM 6-port UTP/STP 9 FSM FDDI module 10 FSM CDDI module 11 ESM 4-port 12 ASM .5 MB multi-mode 13 ESM 12-port 10BASE-T 14 ESM 6-port universal module 15 MPM version II 16 ATM DS-3 17 FSM FDDI single mode 18 ASM .5 MB single mode 19 ASM UTP 20 ESM 8-port fiber </td> <td style="width: 50%; vertical-align: top;"> <ul style="list-style-type: none"> 21 ESM 12-port Telco 22 TSM fiber 23 ASM 2 MB multi-mode 24 ASM 2 MB single mode 25 WSM 26 WSM BRI 27 HSM2 base slot type 28 PizzaSwitch reserved 29 TSM CD-6 30 ASM 2 MB single mode 33 10Meg Ether Universal 34 ATM E3 (European) 35 Ether 100 FX Sngl Full Dup 36 Ether 100 FX Multi Full Dup 37 Ether 100 TX CU Full Dup 39 PizzaPort (repeater) </td> </tr> </table>	<ul style="list-style-type: none"> 4 HSM 5 MPM 6 ESM 8-port 10BASE-T 7 ESM 16-port 8 TSM 6-port UTP/STP 9 FSM FDDI module 10 FSM CDDI module 11 ESM 4-port 12 ASM .5 MB multi-mode 13 ESM 12-port 10BASE-T 14 ESM 6-port universal module 15 MPM version II 16 ATM DS-3 17 FSM FDDI single mode 18 ASM .5 MB single mode 19 ASM UTP 20 ESM 8-port fiber 	<ul style="list-style-type: none"> 21 ESM 12-port Telco 22 TSM fiber 23 ASM 2 MB multi-mode 24 ASM 2 MB single mode 25 WSM 26 WSM BRI 27 HSM2 base slot type 28 PizzaSwitch reserved 29 TSM CD-6 30 ASM 2 MB single mode 33 10Meg Ether Universal 34 ATM E3 (European) 35 Ether 100 FX Sngl Full Dup 36 Ether 100 FX Multi Full Dup 37 Ether 100 TX CU Full Dup 39 PizzaPort (repeater)
<ul style="list-style-type: none"> 4 HSM 5 MPM 6 ESM 8-port 10BASE-T 7 ESM 16-port 8 TSM 6-port UTP/STP 9 FSM FDDI module 10 FSM CDDI module 11 ESM 4-port 12 ASM .5 MB multi-mode 13 ESM 12-port 10BASE-T 14 ESM 6-port universal module 15 MPM version II 16 ATM DS-3 17 FSM FDDI single mode 18 ASM .5 MB single mode 19 ASM UTP 20 ESM 8-port fiber 	<ul style="list-style-type: none"> 21 ESM 12-port Telco 22 TSM fiber 23 ASM 2 MB multi-mode 24 ASM 2 MB single mode 25 WSM 26 WSM BRI 27 HSM2 base slot type 28 PizzaSwitch reserved 29 TSM CD-6 30 ASM 2 MB single mode 33 10Meg Ether Universal 34 ATM E3 (European) 35 Ether 100 FX Sngl Full Dup 36 Ether 100 FX Multi Full Dup 37 Ether 100 TX CU Full Dup 39 PizzaPort (repeater) 		

Trap Type	powerEvent
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.3
Description	A power supply was either inserted or removed from the chassis, or there is a problem with the power supply. This trap is also generated when a power supply is switched on or off.
Bit Position (Word 1)	2
Hex Value (Word 1)	4
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p style="text-align: center;">Power Supply was inserted or removed from chassis or has a problem (ps1 3, ps2 2)</p> </div> <p>Power Supply Status. The current state of power supply 1 (ps1) and power supply 2 (ps2). The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Unknown. 2 No power supply present. 3 Power supply okay. 4 Power supply bad.

Trap Type	controllerEvent
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.4
Description	A chassis controller (MPX) lost or gained the state of the master.
Bit Position (Word 1)	3
Hex Value (Word 1)	8
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px; display: inline-block;"> Chassis controller (MPX) lost or gained master control (slot 1, state 3) </div> <p>Slot. The slot number of the MPX that has lost or gained master control. Valid values are:</p> <ul style="list-style-type: none"> 1 Slot Number 1 2 Slot Number 2 <p>State. The current state of the MPX in the slot. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Unknown 2 Invalid 3 Master 4 Slave

Trap Type	loginViolation
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.5
Description	A login attempt for the User Interface (UI) failed due to an incorrect login ID or an invalid password. Three (3) consecutive unsuccessful attempts will trigger this alarm.
Bit Position (Word 1)	4
Hex Value (Word 1)	10
Trap Text and Variable Descriptions	Login Attempt failed due to invalid ID or password.

Trap Type	macVlanViolation
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.6
Description	Data from a MAC address that previously came from one a port with a VLAN-ID different from the VLAN where the frame had been previously received.
Bit Position (Word 1)	5
Hex Value (Word 1)	20
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; display: inline-block; margin-bottom: 10px;"> Receiving Port VLAN ID has changed (bridge address 0036589adf01) </div> <p>MAC Address. The MAC address from which data has come from two different ports in two different groups.</p>

Trap Tables

Trap Type	macDuplicatePort
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.7
Description	Data from a MAC address that previously came from a source port different from the port where the frame previously was received although they both ports belong to the same VLAN.
Bit Position (Word 1)	6
Hex Value (Word 1)	40
Trap Text and Variable Descriptions	<div data-bbox="425 709 1258 783" style="border: 1px solid black; padding: 5px; text-align: center;">VLAN Receiving Port has changed (bridge address 00145221cd02)</div> <p>MAC Address. The MAC address from which data has come from two different ports in the same group.</p>

Trap Type	portLinkUpEvent																																												
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.8																																												
Description	A physical, logical, or virtual port was enabled. These ports may be enabled through the UI or Switch Manager. Note that if you enable a physical port, any associated logical and virtual ports will also be enabled. And if you enable a logical port, such as an ATM service, associated virtual ports will be enabled.																																												
Bit Position (Word 1)	7																																												
Hex Value (Word 1)	80																																												
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; text-align: center;">Physical, logical or virtual port was enabled (slot 2 IF 2 type 203 instance 1)</div> <p>Slot Number. The slot number for the module that contains this port.</p> <p>Port Number. The port number on this module that was enabled.</p> <p>Port Type. The physical type of this port. The following integers are valid values:</p> <table style="margin-left: 40px;"> <tr><td>1</td><td>Unknown</td></tr> <tr><td>2</td><td>Other</td></tr> <tr><td>3</td><td>Router</td></tr> <tr><td>4</td><td>Bridge</td></tr> <tr><td>5</td><td>Trunk</td></tr> <tr><td>6</td><td>ATM trunk port</td></tr> <tr><td>7</td><td>ATM LAN Emulation port</td></tr> <tr><td>8</td><td>Classical IP</td></tr> <tr><td>9</td><td>ATM MUX</td></tr> <tr><td>203</td><td>Ethernet 10BASE-T</td></tr> <tr><td>204</td><td>Ethernet 100BASE-T</td></tr> <tr><td>205</td><td>Token Ring 4 mbs</td></tr> <tr><td>206</td><td>Token Ring 16 mbs</td></tr> <tr><td>207</td><td>FDDI</td></tr> <tr><td>208</td><td>CDDI</td></tr> <tr><td>209</td><td>ATM 25 mbs</td></tr> <tr><td>210</td><td>ATM 50 mbs</td></tr> <tr><td>211</td><td>DS-1</td></tr> <tr><td>212</td><td>DS-3</td></tr> <tr><td>213</td><td>OC-3</td></tr> <tr><td>214</td><td>OC-12</td></tr> <tr><td>215</td><td>OC-48</td></tr> </table> <p>Physical Instance. The specific instance of this slot/port/type. In most cases this value will be 1 (only one instance of the port), but an ATM port may have multiple instances. Possible values range from 1 to 254.</p>	1	Unknown	2	Other	3	Router	4	Bridge	5	Trunk	6	ATM trunk port	7	ATM LAN Emulation port	8	Classical IP	9	ATM MUX	203	Ethernet 10BASE-T	204	Ethernet 100BASE-T	205	Token Ring 4 mbs	206	Token Ring 16 mbs	207	FDDI	208	CDDI	209	ATM 25 mbs	210	ATM 50 mbs	211	DS-1	212	DS-3	213	OC-3	214	OC-12	215	OC-48
1	Unknown																																												
2	Other																																												
3	Router																																												
4	Bridge																																												
5	Trunk																																												
6	ATM trunk port																																												
7	ATM LAN Emulation port																																												
8	Classical IP																																												
9	ATM MUX																																												
203	Ethernet 10BASE-T																																												
204	Ethernet 100BASE-T																																												
205	Token Ring 4 mbs																																												
206	Token Ring 16 mbs																																												
207	FDDI																																												
208	CDDI																																												
209	ATM 25 mbs																																												
210	ATM 50 mbs																																												
211	DS-1																																												
212	DS-3																																												
213	OC-3																																												
214	OC-12																																												
215	OC-48																																												

Trap Type	portLinkDownEvent																																												
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.9																																												
Description	A physical, logical, or virtual port was disabled. These ports may be disabled through the UI or Switch Manager. Note that if you disable a physical port, any associated logical and virtual ports will also be disabled. And if you disable a logical port, such as an ATM service, associated virtual ports will also be disabled.																																												
Bit Position (Word 1)	8																																												
Hex Value (Word 1)	100																																												
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> Physical, logical or virtual port was disabled (slot 2 IF 2 type 203 instance 1) </div> <p>Slot Number. The slot number for the module that contains this port.</p> <p>Port Number. The port number on this module that was disabled.</p> <p>Port Type. The physical type of this port. The following integers are valid values:</p> <table style="margin-left: 40px;"> <tr><td>1</td><td>Unknown</td></tr> <tr><td>2</td><td>Other</td></tr> <tr><td>3</td><td>Router</td></tr> <tr><td>4</td><td>Bridge</td></tr> <tr><td>5</td><td>Trunk</td></tr> <tr><td>6</td><td>ATM trunk port</td></tr> <tr><td>7</td><td>ATM LAN Emulation port</td></tr> <tr><td>8</td><td>Classical IP</td></tr> <tr><td>9</td><td>ATM MUX</td></tr> <tr><td>203</td><td>Ethernet 10BASE-T</td></tr> <tr><td>204</td><td>Ethernet 100BASE-T</td></tr> <tr><td>205</td><td>Token Ring 4 mbs</td></tr> <tr><td>206</td><td>Token Ring 16 mbs</td></tr> <tr><td>207</td><td>FDDI</td></tr> <tr><td>208</td><td>CDDI</td></tr> <tr><td>209</td><td>ATM 25 mbs</td></tr> <tr><td>210</td><td>ATM 50 mbs</td></tr> <tr><td>211</td><td>DS-1</td></tr> <tr><td>212</td><td>DS-3</td></tr> <tr><td>213</td><td>OC-3</td></tr> <tr><td>214</td><td>OC-12</td></tr> <tr><td>215</td><td>OC-48</td></tr> </table> <p>Physical Instance. The specific instance of this slot/port/type. In most cases this value will be 1 (only one instance of the port), but an ATM port may have multiple instances. Possible values range from 1 to 254.</p>	1	Unknown	2	Other	3	Router	4	Bridge	5	Trunk	6	ATM trunk port	7	ATM LAN Emulation port	8	Classical IP	9	ATM MUX	203	Ethernet 10BASE-T	204	Ethernet 100BASE-T	205	Token Ring 4 mbs	206	Token Ring 16 mbs	207	FDDI	208	CDDI	209	ATM 25 mbs	210	ATM 50 mbs	211	DS-1	212	DS-3	213	OC-3	214	OC-12	215	OC-48
1	Unknown																																												
2	Other																																												
3	Router																																												
4	Bridge																																												
5	Trunk																																												
6	ATM trunk port																																												
7	ATM LAN Emulation port																																												
8	Classical IP																																												
9	ATM MUX																																												
203	Ethernet 10BASE-T																																												
204	Ethernet 100BASE-T																																												
205	Token Ring 4 mbs																																												
206	Token Ring 16 mbs																																												
207	FDDI																																												
208	CDDI																																												
209	ATM 25 mbs																																												
210	ATM 50 mbs																																												
211	DS-1																																												
212	DS-3																																												
213	OC-3																																												
214	OC-12																																												
215	OC-48																																												

Trap Type	portPartitioned
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.10
Description	The physical port detected jabber (i.e., the port has transitioned through enable/disable states more than 50 times in the past 200 ms). Jabber may be produced by a bad port connection, such as a faulty cable.
Bit Position (Word 1)	9
Hex Value (Word 1)	200
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p style="text-align: center;">Port jabber detected (enabled/disabled faster than 50 times in 200 ms) (slot 2, IF 2, type 203, instance 1)</p> </div> <p>Slot Number. The slot number for the module that contains this port.</p> <p>Port Number. The port number on this module that detected jabber.</p> <p>Port Type. The physical type of this port. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Unknown 2 Other 3 Router 4 Bridge 5 Trunk 6 ATM trunk port 7 ATM LAN Emulation port 8 Classical IP 9 ATM MUX 203 Ethernet 10BASE-T 204 Ethernet 100BASE-T 205 Token Ring 4 mbs 206 Token Ring 16 mbs 207 FDDI 208 CDDI 209 ATM 25 mbs 210 ATM 50 mbs 211 DS-1 212 DS-3 213 OC-3 214 OC-12 215 OC-48 <p>Physical Instance. The specific instance of this slot/port/type. In most cases this value will be 1 (only one instance of the port), but an ATM port may have multiple instances. Possible values range from 1 to 254.</p>

Trap Type	portRecordMismatch																																												
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.11																																												
Description	The port configuration is different from the previous configuration. Typically this trap is generated when a NIC of one type is swapped for a different type (i.e., Ethernet for FDDI, ATM for Token Ring, etc.).																																												
Bit Position (Word 1)	10																																												
Hex Value (Word 1)	400																																												
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Port configuration different than previously detected (slot 2, IF 2, type 203, instance 1)</p> </div> <p>Slot number. The slot number for the module that contains this port.</p> <p>Port number. The port number on this module that has a different configuration.</p> <p>Port Type. The physical type of this port. The following integers are valid values:</p> <table style="margin-left: 20px;"> <tr><td>1</td><td>Unknown</td></tr> <tr><td>2</td><td>Other</td></tr> <tr><td>3</td><td>Router</td></tr> <tr><td>4</td><td>Bridge</td></tr> <tr><td>5</td><td>Trunk</td></tr> <tr><td>6</td><td>ATM trunk port</td></tr> <tr><td>7</td><td>ATM LAN Emulation port</td></tr> <tr><td>8</td><td>Classical IP</td></tr> <tr><td>9</td><td>ATM MUX</td></tr> <tr><td>203</td><td>Ethernet 10BASE-T</td></tr> <tr><td>204</td><td>Ethernet 100BASE-T</td></tr> <tr><td>205</td><td>Token Ring 4 mbs</td></tr> <tr><td>206</td><td>Token Ring 16 mbs</td></tr> <tr><td>207</td><td>FDDI</td></tr> <tr><td>208</td><td>CDDI</td></tr> <tr><td>209</td><td>ATM 25 mbs</td></tr> <tr><td>210</td><td>ATM 50 mbs</td></tr> <tr><td>211</td><td>DS-1</td></tr> <tr><td>212</td><td>DS-3</td></tr> <tr><td>213</td><td>OC-3</td></tr> <tr><td>214</td><td>OC-12</td></tr> <tr><td>215</td><td>OC-48</td></tr> </table> <p>Physical Instance. The specific instance of this slot/port/type. In most cases this value will be 1 (only one instance of the port), but an ATM port may have multiple instances. Possible values range from 1 to 254.</p>	1	Unknown	2	Other	3	Router	4	Bridge	5	Trunk	6	ATM trunk port	7	ATM LAN Emulation port	8	Classical IP	9	ATM MUX	203	Ethernet 10BASE-T	204	Ethernet 100BASE-T	205	Token Ring 4 mbs	206	Token Ring 16 mbs	207	FDDI	208	CDDI	209	ATM 25 mbs	210	ATM 50 mbs	211	DS-1	212	DS-3	213	OC-3	214	OC-12	215	OC-48
1	Unknown																																												
2	Other																																												
3	Router																																												
4	Bridge																																												
5	Trunk																																												
6	ATM trunk port																																												
7	ATM LAN Emulation port																																												
8	Classical IP																																												
9	ATM MUX																																												
203	Ethernet 10BASE-T																																												
204	Ethernet 100BASE-T																																												
205	Token Ring 4 mbs																																												
206	Token Ring 16 mbs																																												
207	FDDI																																												
208	CDDI																																												
209	ATM 25 mbs																																												
210	ATM 50 mbs																																												
211	DS-1																																												
212	DS-3																																												
213	OC-3																																												
214	OC-12																																												
215	OC-48																																												

Trap Type	groupChange
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.14
Description	A Group was either created or deleted through the UI or Switch Manager.
Bit Position (Word 1)	13
Hex Value (Word 1)	2000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px; text-align: center;"> Group created or deleted (vlan 2 admin status 4) </div> <p>Group number. The Group number that has been created or deleted.</p> <p>Administrative Status. The administrative status for this group. Possible options are:</p> <ol style="list-style-type: none"> 1 Disabled. All ports in this Group are disabled. 2 Enabled. All ports in this Group are enabled. 3 Deleted. This Group was deleted, and all attached virtual ports and routers are detached and deleted. 4 Created. This Group has been created. 5 Modify. This Group has been modified.

Trap Type	vlanChange
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.15
Description	A VLAN was either created or deleted through the UI or Switch Manager.
Bit Position (Word 1)	14
Hex Value (Word 1)	4000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px; text-align: center;"> VLAN Change created or deleted (group 2, admin status 4) </div> <p>Group number. The Group number to which this VLAN belongs.</p> <p>Administrative status. The administrative status for this VLAN. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Enabled. 2 Disabled. 3 Deleted. This VLAN was deleted. 4 Created. This Group has been created. 5 Modify. This Group has been modified.

Trap Type	portMove																																												
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.16																																												
Description	The specified port has moved from a Group or has had its configuration changed.																																												
Bit Position (Word 1)	15																																												
Hex Value (Word 1)	8000																																												
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px;"> Port VLAN, group or configuration change (slot 2, IF 8, type 4, instance 1) </div> <p>Slot number. The slot number for the module that contains this port.</p> <p>Port number. The port number on this module that was changed.</p> <p>Port Type. The physical type of this port. The following integers are valid values:</p> <table style="margin-left: 40px;"> <tr><td>1</td><td>Unknown</td></tr> <tr><td>2</td><td>Other</td></tr> <tr><td>3</td><td>Router</td></tr> <tr><td>4</td><td>Bridge</td></tr> <tr><td>5</td><td>Trunk</td></tr> <tr><td>6</td><td>ATM trunk port</td></tr> <tr><td>7</td><td>ATM LAN Emulation port</td></tr> <tr><td>8</td><td>Classical IP</td></tr> <tr><td>9</td><td>ATM MUX</td></tr> <tr><td>203</td><td>Ethernet 10BASE-T</td></tr> <tr><td>204</td><td>Ethernet 100BASE-T</td></tr> <tr><td>205</td><td>Token Ring 4 mbs</td></tr> <tr><td>206</td><td>Token Ring 16 mbs</td></tr> <tr><td>207</td><td>FDDI</td></tr> <tr><td>208</td><td>CDDI</td></tr> <tr><td>209</td><td>ATM 25 mbs</td></tr> <tr><td>210</td><td>ATM 50 mbs</td></tr> <tr><td>211</td><td>DS-1</td></tr> <tr><td>212</td><td>DS-3</td></tr> <tr><td>213</td><td>OC-3</td></tr> <tr><td>214</td><td>OC-12</td></tr> <tr><td>215</td><td>OC-48</td></tr> </table> <p>Physical Instance. The specific instance of this slot/port/type. In most cases this value will be 1 (only one instance of the port), but an ATM port may have multiple instances. Possible values range from 1 to 254.</p>	1	Unknown	2	Other	3	Router	4	Bridge	5	Trunk	6	ATM trunk port	7	ATM LAN Emulation port	8	Classical IP	9	ATM MUX	203	Ethernet 10BASE-T	204	Ethernet 100BASE-T	205	Token Ring 4 mbs	206	Token Ring 16 mbs	207	FDDI	208	CDDI	209	ATM 25 mbs	210	ATM 50 mbs	211	DS-1	212	DS-3	213	OC-3	214	OC-12	215	OC-48
1	Unknown																																												
2	Other																																												
3	Router																																												
4	Bridge																																												
5	Trunk																																												
6	ATM trunk port																																												
7	ATM LAN Emulation port																																												
8	Classical IP																																												
9	ATM MUX																																												
203	Ethernet 10BASE-T																																												
204	Ethernet 100BASE-T																																												
205	Token Ring 4 mbs																																												
206	Token Ring 16 mbs																																												
207	FDDI																																												
208	CDDI																																												
209	ATM 25 mbs																																												
210	ATM 50 mbs																																												
211	DS-1																																												
212	DS-3																																												
213	OC-3																																												
214	OC-12																																												
215	OC-48																																												

Trap	moduleResetReload																								
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.17																								
Description	The specified module has been either reset or reloaded. A reload may occur during a firmware download.																								
Bit Position (Word 1)	16																								
Hex Value (Word 1)	1 0000																								
Trap Text and Variable Descriptions	<p>Submodule Type. Indicates the submodule that was reset or reloaded. Typically this value will be 1, meaning the base module was reset or reloaded. If this value is 2, then HSX module 1 was affected. If this value is 3, then HSX module 2 was affected.</p> <p>.Slot number. The slot number of the module that was reset or reloaded.</p> <div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin: 10px 0;"> <p>Module reset or reloaded by chassis manager (slot 4 subunit 1 type 6 status 3)</p> </div> <p>Module Type. Indicates the module type that was reset or reloaded. The following integers are valid values:</p> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%;">4 HSM</td> <td style="width: 50%;">13 ESM 12-port 10BASE-T</td> </tr> <tr> <td>5 MPM</td> <td>14 ESM 6-port universal module</td> </tr> <tr> <td>6 ESM 8-port 10BASE-T</td> <td>15 MPM version II</td> </tr> <tr> <td>7 ESM 16-port</td> <td>16 ATM DS-3</td> </tr> <tr> <td>8 TSM 6-port UTP/STP</td> <td>17 FSM FDDI single mode</td> </tr> <tr> <td>9 FSM FDDI module</td> <td>18 ASM .5 MB single mode</td> </tr> <tr> <td>10 FSM CDDI module</td> <td>19 ASM UTP</td> </tr> <tr> <td>11 ESM 4-port</td> <td>20 ESM 8-port fiber</td> </tr> <tr> <td>12 ASM .5 MB multi-mode</td> <td>21 ESM 12-port Telco</td> </tr> <tr> <td></td> <td>22 TSM fiber</td> </tr> <tr> <td></td> <td>23 ASM 2 MB multi-mode</td> </tr> <tr> <td></td> <td>24 ASM 2 MB single mode</td> </tr> </table> <p>Operational State. Indicates the current state of the module that was reset or reloaded. The following integers are valid values:</p> <ol style="list-style-type: none"> 1 Unknown state. The module may have failed low-level self-test. 2 Invalid. The module may exist, by the chassis does not have control of it. 3 Operational. The module is running fine with no errors. 4 Disabled. The module has been set to disable through the UI or SNMP. 5 Reset. The module has been reset. 6 Loading. The module is in the middle of loading. 7 Testing. The module is in self-test. 8 Warning. A warning was detected during operation. 9 Non-fatal error. A non-fatal error was detected during operation. 10 Fatal error. A fatal error occurred during operation. The module may or may not be functional. 	4 HSM	13 ESM 12-port 10BASE-T	5 MPM	14 ESM 6-port universal module	6 ESM 8-port 10BASE-T	15 MPM version II	7 ESM 16-port	16 ATM DS-3	8 TSM 6-port UTP/STP	17 FSM FDDI single mode	9 FSM FDDI module	18 ASM .5 MB single mode	10 FSM CDDI module	19 ASM UTP	11 ESM 4-port	20 ESM 8-port fiber	12 ASM .5 MB multi-mode	21 ESM 12-port Telco		22 TSM fiber		23 ASM 2 MB multi-mode		24 ASM 2 MB single mode
4 HSM	13 ESM 12-port 10BASE-T																								
5 MPM	14 ESM 6-port universal module																								
6 ESM 8-port 10BASE-T	15 MPM version II																								
7 ESM 16-port	16 ATM DS-3																								
8 TSM 6-port UTP/STP	17 FSM FDDI single mode																								
9 FSM FDDI module	18 ASM .5 MB single mode																								
10 FSM CDDI module	19 ASM UTP																								
11 ESM 4-port	20 ESM 8-port fiber																								
12 ASM .5 MB multi-mode	21 ESM 12-port Telco																								
	22 TSM fiber																								
	23 ASM 2 MB multi-mode																								
	24 ASM 2 MB single mode																								

Trap Type	systemEvent
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.18
Description	A potentially fatal error occurred in the system.
Bit Position (Word 1)	17
Hex Value (Word 1)	2 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px; text-align: center;"> <p>Potentially fatal error occurred (trap 10)</p> </div> <p>Event Trap Type. A number that identifies the specific error that occurred in the system. The following integers are valid values:</p> <ul style="list-style-type: none"> 10 Unspecified Log Event 11 Log file full 12 Log file erased 20 Unspecified memory event 21 Memory shortage 30 Unspecified CPU event 31 Long term CPU overload 32 Short term CPU overload 40 Unspecified ffs event 41 Attempt to write to full ffs 42 System/user directed purge 43 Removed imgs/cfgs 44 Exec file removed 45 Config file removed 46 Exec file updated 47 Config file updated 50 Unspecified chassis event 51 Module failed to init 52 Module failed to load 53 Module startup failed 54 Module failed 55 Driver failed

Trap Tables

Trap Type	vlanRouteTableFull
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.19
Description	The IP or IPX route table is full.
Bit Position (Word 1)	18
Hex Value (Word 1)	4 0000
Trap Text and Variable Descriptions	IP or IPX route table is full on insertion.

Trap Type	sapTableFull
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.20
Description	The SAP table is full upon insertion.
Bit Position (Word 1)	19
Hex Value (Word 1)	8 0000
Trap Text and Variable Descriptions	SAP table full on insertion.

Trap Type	atmSSCOPstate
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.21
Description	A specified port changed.
Bit Position (Word 1)	20
Hex Value (Word 1)	10 0000
Trap Text and Variable Descriptions	<div style="text-align: center; border: 1px solid black; padding: 5px; width: fit-content; margin: 0 auto;"> Signalling state changed (slot 3 port 1) </div> <p>Slot number. The slot number where this ASM module is located.</p> <p>Port number. The port number on this ASM module where the signalling state has changed.</p>

Trap Type	ilmiState
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.22
Description	The ILMI state for the specified port changed. This change of state indicates whether address registration was successful, and whether the switch knows the network prefix provided by the external ATM switch.
Bit Position (Word 1)	21
Hex Value (Word 1)	20 0000
Trap Text and Variable Descriptions	<div style="text-align: center; border: 1px solid black; padding: 5px; width: fit-content; margin: 0 auto;"> ILMI state changed (slot 3 port 1) </div> <p>Slot number. The slot number where this ASM module is located.</p> <p>Port number. The port number on this ASM module where the ILMI state has changed.</p>

Trap Type	atmConnection
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.23
Description	The specified ATM VCC was created or deleted.
Bit Position (Word 1)	22
Hex Value (Word 1)	40 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> ATM VCC created or deleted (slot 3, port 1, Vpi 0, Vci 100, admin status 2) </div> <p>Slot Number. The slot number where this ASM module is located.</p> <p>Port Number. The port number on the ASM module where this VCC was created or deleted.</p> <p>VPI Number. The virtual path identifier for this virtual connec-</p> <p>VCI Number. The virtual channel identifier for this virtual connection.</p> <p>Admin Status. Indicates the current status of this ATM VCC. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Disabled. This VCC was disabled. 2 Enabled. This VCC was enabled. 3 Deleted. This VCC was deleted.

Trap Type	atmService
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.24
Description	The specified ATM service (Port-to-Port Bridging, Trunking, LAN Emulation, etc.) was created or deleted.
Bit Position (Word 1)	23
Hex Value (Word 1)	80 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> ATM service created or deleted (slot 3, port 1, service 2, admin status 2) </div> <p>Slot Number. The slot number where this ASM module is located.</p> <p>Port Number. The port number on the ASM module where the service was created or deleted.</p> <p>Service Number. The ATM service number assigned to this service when it was set up.</p> <p>Admin Status. The current status of this ATM VCC. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Disabled. This VCC has disabled. 2 Enabled. This VCC was enabled. 3 Deleted. This VCC was deleted.

Trap Type	dlciNew
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.27
Description	Frame Relay DLCI was created.
Bit Position (Word 1)	26
Hex Value (Word 1)	400 0000
Trap Text and Variable Descriptions	<div style="text-align: center; border: 1px solid black; padding: 5px; margin: 10px auto; width: fit-content;"> Frame Relay DLCI created (slot 3 port 1 DLCI Number 100) </div> <p>Slot number. The slot number where this Frame Relay module is located.</p> <p>Port number. The port number on this Frame Relay module where the DLCI was created.</p> <p>DLCI Number. The number of the DLCI that was created.</p>

Trap Type	dlciDel
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.28
Description	Frame Relay DLCI was deleted.
Bit Position (Word 1)	27
Hex Value (Word 1)	800 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; text-align: center; margin-bottom: 10px;"> Frame Relay DLCI deleted (slot 3 port 1 DLCI Number 100) </div> <p>Slot number. The slot number where this Frame Relay module is located.</p> <p>Port number. The port number on this Frame Relay module where the DLCI was deleted.</p> <p>DLCI number. The number of the DLCI that was just deleted.</p>

Trap Tables

Trap Type	dlciUp
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.29
Description	Frame Relay DLCI changed to active state.
Bit Position (Word 1)	28
Hex Value (Word 1)	1000 0000
Trap Text and Variable Descriptions	<div data-bbox="396 646 1284 705" style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; text-align: center;">Frame Relay DLCI Changed to Active (slot 3 port 1 DLCI Number 100)</div> <p style="text-align: center;">Slot Number. The slot number where this Frame Relay module is located.</p> <p style="text-align: center;">Port Number. The port number on this Frame Relay module where the DLCI was activated.</p> <p style="text-align: center;">DLCI Number. The number of the DLCI that was just activated.</p>

Trap Type	dlciDn
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.30
Description	Frame Relay DLCI changed to inactive state.
Bit Position (Word 1)	29
Hex Value (Word 1)	2000 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p style="text-align: center;">Frame Relay DLCI Changed to Inactive (slot 3 port 1 DLCI Number 100)</p> </div> <p style="text-align: center;">Slot Number. The slot number where this Frame Relay module is located.</p> <p style="text-align: center;">Port Number. The port number on this Frame Relay module where the DLCI was de-activated.</p> <p style="text-align: center;">DLCI Number. The number of the DLCI that was just de-activated.</p>

Trap Type	portManualForwardingMode
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.31
Description	The specified port was placed into manual mode forwarding as its default setting.
Bit Position (Word 1)	30
Hex Value (Word 1)	4000 0000
Trap Text and Variable Descriptions	<p style="text-align: center;"> Slot Number. The slot number where this port is located. </p> <p style="text-align: center;"> Port number. The port number on the module. </p> <div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; text-align: center; margin: 10px auto; width: fit-content;"> Port placed into manual mode forwarding (slot 3, port 1, type 1, instance 1) </div> <p style="text-align: center;"> Port Type. The physical type of this port. The following integers are valid values: </p> <ul style="list-style-type: none"> 1 Unknown 2 Other 3 Router 4 Bridge 5 Trunk 6 ATM trunk port 7 ATM LAN Emulation port 8 Classical IP 9 ATM MUX 203 Ethernet 10BASE-T 204 Ethernet 100BASE-T 205 Token Ring 4 mbs 206 Token Ring 16 mbs 207 FDDI 208 CDDI 209 ATM 25 mbs 210 ATM 50 mbs 211 DS-1 212 DS-3 213 OC-3 214 OC-12 215 OC-48 <p style="text-align: center;"> Physical Instance. The specific instance of this slot/port/type. In most cases this value will be 1 (only one instance of the port), but an ATM port may have multiple instances. Possible values range from 1 to 254. </p>

Trap Type	fddiCFStateChange																										
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.32																										
Description	The specified FDDI physical port changed from wrap configuration state.																										
Bit Position (Word 1)	31																										
Hex Value (Word 1)	8000 0000																										
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px;"> FDDI physical port changes from wrap configuration state (index 1, state 2) </div> <p>SMT Index. A unique value for each SMT (Station Management Station). The value for each SMT must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.</p> <p>SMT State. The attachment configuration for the station or concentrator. The following integers are valid values:</p> <table style="margin-left: 40px;"> <tr><td>1</td><td>isolated</td></tr> <tr><td>2</td><td>local_a</td></tr> <tr><td>3</td><td>local_b</td></tr> <tr><td>4</td><td>local_ab</td></tr> <tr><td>5</td><td>local_s</td></tr> <tr><td>6</td><td>wrap_a</td></tr> <tr><td>7</td><td>wrap_b</td></tr> <tr><td>8</td><td>wrap_ab</td></tr> <tr><td>9</td><td>wrap_s</td></tr> <tr><td>10</td><td>c_wrap_a</td></tr> <tr><td>11</td><td>c_wrap_b</td></tr> <tr><td>12</td><td>c_wrap_s</td></tr> <tr><td>13</td><td>thru</td></tr> </table>	1	isolated	2	local_a	3	local_b	4	local_ab	5	local_s	6	wrap_a	7	wrap_b	8	wrap_ab	9	wrap_s	10	c_wrap_a	11	c_wrap_b	12	c_wrap_s	13	thru
1	isolated																										
2	local_a																										
3	local_b																										
4	local_ab																										
5	local_s																										
6	wrap_a																										
7	wrap_b																										
8	wrap_ab																										
9	wrap_s																										
10	c_wrap_a																										
11	c_wrap_b																										
12	c_wrap_s																										
13	thru																										

Trap Tables

Trap Type	duplicateIPaddress
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.35
Description	The switch detected a duplicate IP address.
Bit Position (Word 2)	2
Hex Value (Word 2)	4
Trap Text and Variable Descriptions	<p>IP Address. The IP address of the station that reported the duplicate IP address.</p> <p>MAC Address. The MAC address of the station that reported the duplicate IP address.</p> <p style="text-align: center;">Duplicate IP address detected (IP addr 192.168.10.1, Mac 0036589adf01, slot 3, IF 4, dup Mac 00145221cd02, dup slot 1, dup IF 3)</p> <p>Port Number. The port on the module of the reporting station from which the trap was sent.</p> <p>Duplicate Slot. The slot number on the reporting station where the duplicate address was discovered.</p> <p>Duplicate Port. The port on the module of the reporting station where the duplicate address was discovered.</p> <p>Slot Number. The slot number of the reporting station from which the trap was sent.</p> <p>Duplicate MAC. The MAC address associated with the duplicated IP address.</p>

Trap Type	duplicateMACAddress
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.36
Description	The switch detected a duplicate MAC address of one of its own router ports.
Bit Position (Word 2)	3
Hex Value (Word 2)	8
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Duplicate MAC address detected (Mac 00145221cd02, slot 2, IF 3, time 4</p> </div> <p>MAC Address. The router port's MAC address for which the last duplicate MAC address was detected.</p> <p>Slot. The slot number where the duplicate MAC address was last received.</p> <p>Interface. The interface number where the duplicate MAC address was last received.</p> <p>Time. The time, in seconds, when the duplicate MAC was detected.</p>

Trap Tables

Trap Type	healthThresholdRising
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.37
Description	At least one of the user-specified thresholds was exceeded.
Bit Position (Word 2)	4
Hex Value (Word 2)	10
Trap Text and Variable Descriptions	Thresh-hold rising trap

Trap Type	healthThresholdFalling
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.38
Description	At least one of the user-specified thresholds was exceeded during the previous cycle and none of them are exceeded in the current cycle.
Bit Position (Word 2)	5
Hex Value (Word 2)	20
Trap Text and Variable Descriptions	Thresh-hold falling trap

Trap Type	healthThresholdDevice
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.39
Description	At least one of the device-level threshold crossing was detected.
Bit Position (Word 2)	6
Hex Value (Word 2)	40
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #f0f0f0; padding: 5px; margin-bottom: 10px;"> Device-level threshold crossing is detected (Data 0a 09 0d 53 00 00 00 00 00 00 00 00 00 00) </div> <p>Data. An octet string that represents the contents of device-level rising/falling threshold trap.</p>

Trap Type	healthThresholdModule
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.40
Description	At least one module-level threshold crossing was detected.
Bit Position (Word 2)	7
Hex Value (Word 2)	80
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #f0f0f0; padding: 5px; margin-bottom: 10px;"> Module-level threshold crossing is detected (count 2, data 0a 09 0d 53 00 00 00 00 00 00 00 00 00 00) </div> <p>Count. The number of modules with threshold crossing data in module-level rising/falling threshold traps.</p> <p>Data. An octet string that represents the contents of device-level rising/falling threshold trap.</p>

Trap Type	xylanXIPXMAPPortStatusChange
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.41
Description	An XMAP turned on or off.
Bit Position (Word 2)	8
Hex Value (Word 2)	100
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px;"> <p>The status of an XMAP-tracked virtual port has changed (port 1, reason 2)</p> </div> <p>Port Number. The virtual port number of the port that most recently changed.</p> <p>Reason. The reason for the last port status change. The following integers are valid values:</p> <ul style="list-style-type: none"> 0 No trap was sent. 1 A port was added. 2 A change of information on an existing port. 3 A port was deleted.

Trap Type	xylanXIPXMAPPortStateChange
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.42
Description	An XMAP turned on or off.
Bit Position (Word 2)	9
Hex Value (Word 2)	200
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px;"> <p style="text-align: center;">The state of the XMAP agent has changed to (state 1)</p> </div> <p>Operating State. The XMAP's operating state. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 inactive 2 active

Trap Type	aviAuthAttempt
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.43
Description	Indicates the last authenticated VLAN attempt.
Bit Position (Word 2)	16
Hex Value (Word 2)	1 0000
Trap Text and Variable Descriptions	<p style="text-align: center;">User. The last user who made an authentication attempt.</p> <div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; text-align: center; margin: 10px auto; width: fit-content;"> <p>The last VLAN authentication attempt was: (user 1, event 2, MAC 0036589adf01, port 4, slot 5)</p> </div> <p>MAC Address. The last MAC address to make an authentication attempt.</p> <p>Port. The last port number from which the authentication attempt originated.</p> <p>Event Type. The last authorization attempt type. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Successful login 2 Failed Login Attempt 3 Logout/Drop <p>Slot. The last slot number from which the authentication attempt originated.</p>

Trap Type	xylanXIPGMAPFailedUpdate
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.44
Description	GMAP is unable to update the forwarding database to reflect information in its internal database.
Bit Position (Word 2)	11
Hex Value (Word 2)	800
Trap Text and Variable Descriptions	<p>Reason. The reason for the last GMAP update was not applied. The following integers are valid values:</p> <ol style="list-style-type: none"> 1 The target group is an authenticated group. 2 The update would conflict with a binding rule. 3 The update would create two different group entries for the same protocol. 4 The update would create two different protocol entries for the same group. 5) The target group is not mobile. <div style="border: 1px solid black; padding: 5px; margin: 10px 0; text-align: center;"> GMAP is unable to update the forwarding database (reason 1, port 2, Mac 0036589adf01, protocol 4, group 5) </div> <p>MAC Address. The last MAC address for which a GMAP change was not applied.</p> <p>Group. The group identifier of the last GMAP change that was not applied.</p> <p>Port. The virtual port number of the last port on which the GMAP change was not applied.</p> <p>Protocol. The protocol identifier of the last GMAP change that was not applied.</p>

Trap Tables

Trap Type	clkBusLineStateChange								
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.45								
Description	Either the bus line's status changed (active or inactive) or clock switching occurred.								
Bit Position (Word 2)	10								
Hex Value (Word 2)	400								
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Bus Line's status changed (bus line 1, operating state 1) or clock switching has occurred.</p> </div> <p>Bus Line. The specific bus line where the status change occurred. The following integers are valid values:</p> <table style="margin-left: 20px;"> <tr> <td>1</td> <td>8 khz</td> </tr> <tr> <td>2</td> <td>19 mhz</td> </tr> </table> <p>Operating State. The bus line's operating state. The following integers are valid values:</p> <table style="margin-left: 20px;"> <tr> <td>1</td> <td>inactive</td> </tr> <tr> <td>2</td> <td>active</td> </tr> </table>	1	8 khz	2	19 mhz	1	inactive	2	active
1	8 khz								
2	19 mhz								
1	inactive								
2	active								

Trap Type	bind-violation
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.46
Description	A configured binding rule was violated.
Bit Position (Word 2)	23
Hex Value (Word 2)	80 0000
Trap Text and Variable Descriptions	<p style="text-align: right;">IP Address. The IP address for which this binding is configured.</p> <p style="text-align: center;">VLAN ID. The VLAN ID for which this rule is configured.</p> <p style="text-align: center;">Group ID. The group ID of the VLAN for which this rule is configured.</p> <div style="border: 1px solid black; padding: 5px; text-align: center; margin: 10px auto; width: fit-content;"> <p>A binding rule has been violated (groupId 1, vlanId 2, IP 192.168.10.1 3, Mac 0036589adf01, protocol 5, port 6, rule 4, index 8)</p> </div> <p style="text-align: center;">Protocol. The protocol for which this binding is configured.</p> <p style="text-align: center;">Port. The port for which this binding is configured.</p> <p style="text-align: center;">MAC Address. The MAC address for which this binding is configured.</p> <p style="text-align: center;">Rule. The rule for which this binding is configured.</p> <p style="text-align: right;">Rule Index. The index which uniquely defines the rule for this VLAN.</p>

Trap Type	mpcStatisticsOverflow
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.47
Description	An entry in the mpcStatisticsTable reached the threshold value.
Bit Position (Word 2)	18
Hex Value (Word 2)	4 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>MPC: Statistics threshold value reached (MpcIndex, Insufficient resources replies.)</p> </div> <p>MPC Index. A unique number that identifies a conceptual row in the mpcConfig-Table.</p> <p>Insufficient resources replies. The reply from the MPC Statistics Table came back as insufficient resources.</p>

Trap Type	fddiLerFlagChange				
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.65				
Description	The LER (Link Error Rate) flag on a port changed from CLEAR to SET.				
Bit Position (Word 3)	0				
Hex Value (Word 3)	1				
Trap Text and Variable Descriptions	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; margin-bottom: 10px;"> <p style="text-align: center;">FDDI: Link Error Rate on a port is set (SMTIndex 1, port 2, LerFlag 3)</p> </div> <p>SMT Index. A unique value for each SMT (Station Management). The value for each SMT must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.</p> <p>Port index. A unique value for each port with in a given SMT, which is the same as the corresponding resource index in SMT. The value for each port must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.</p> <p>LER Flag. The condition becomes active when the value of the fddiPRTLerEstimate is less than or equal to fddimibPORTLerEstimate. The following integers are valid values:</p> <table style="margin-left: 40px;"> <tr> <td>1</td> <td>True</td> </tr> <tr> <td>2</td> <td>False</td> </tr> </table>	1	True	2	False
1	True				
2	False				

Trap Type	fddiLCTFailCntIncr
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.66
Description	The LCT (Link Confidence Test) flag on a port incremented.
Bit Position (Word 3)	1
Hex Value (Word 3)	2
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Fddi: Link Confidence Test flag on a port incremented (SMTIndex 1, port index 2, failure counts 3)</p> </div> <div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>Port Index . A unique value for each port within a given SMT, which is the same as the corresponding resource index in SMT. The value for each port must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.</p> </div> <div style="width: 45%;"> <p>SMT Index. A unique value for each SMT. The value for each SMT must remain constant at least from one re-unitization of the entity's network management system to the next re-initialization.</p> </div> </div> <div style="margin-top: 20px; text-align: center;"> <p>Failure Counts. The count of the consecutive times the link confidence test (LCT) failed during connection management.</p> </div>

Trap Type	mpcStatisticsOverflow
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.67
Description	The statisticsNum value of the mpcStatisticsTable reached the threshold value.
Bit Position (Word 2)	18
Hex Value (Word 2)	4 0000
Variables	mpcIndex mpcStatRxMpoaResolveReplyInsufECResources
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>GMAP is unable to update the forwarding database (index 1, MPOA replies 3)</p> </div> <p>MPC Index. A unique number that identifies a conceptual row in the mpcConfigTable.</p> <p>MPOA Resolution Replies. The number of MPOA Resolution Replies received with an MPOA CIE Code of 0x81.</p>

Trap Type	mpcShortCut
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.68
Description	The established shortcut path either closed or failed to complete the path.
Bit Position (Word 2)	19
Hex Value (Word 2)	8 0000
Variables	mpcRowStatus lecControlDirectVci mpcFlowDetectProtocol mpcIngressCacheDestAddr, mpcIngressCacheDestAtmAddr mpcIndex mpcMpsIndex
Trap Text and Variable Descriptions	<p>Row Status. This object allows creation and deletion of MPOA clients.</p> <div style="border: 1px solid black; padding: 5px; margin: 10px 0;"> <p>GMAP is unable to update the forwarding database (rowStatus 1, control direct Vci 2, protocol 4, dest addr 192.168.40.12, dest ATM addr 3903488001bc900001020000090020da00000900, index 1, mps index 2)</p> </div> <p>Control Direct VCI. The VCI that identifies the VCC at the point where it connects to a LANE client. If the Control Direct VCC does not exist, this value is zero.</p> <p>Protocol. The protocol on which flow detection is performed.</p> <p>Destination ATM Address. The destination ATM address received in the MPOA Resolution Reply.</p> <p>Destination Address. The destination internet-work layer address.</p> <p>MPC Index. A unique number that identifies a conceptual row in the mpcConfig-Table.</p> <p>MPC MPS Index. The MPS's index that is used to identify a row in the mpcConig Table.</p>

Trap Type	mpcIngressRetryTimeOut
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.69
Description	The retry time exceeded the MPC-p5 time.
Bit Position (Word 2)	20
Hex Value (Word 2)	10 0000
Variables	mpcIndex mpcRetryTimeMaximum mpcIngressCacheDestAddr mpcIngressCacheDestAtmAddr mpcFlowDetectProtocol mpcMpsIndex
Trap Text and Variable Descriptions	<p>MPC Index. A unique number that identifies a conceptual row in the mpcConfig-Table.</p> <p>Maximum Retry Time. The MPC-p5 cumulative maximum value for retry time.</p> <p>Destination Address. The destination internet-network layer address.</p> <p>Destination ATM Address. The destination ATM address received in the MPOA Resolution Reply.</p> <p>Detect Protocol. The protocol on which flow detection is performed.</p> <p>GMAP is unable to update the forwarding database (index 1, max time 5, dest addr 192.168.40.12, ATM addr 3903488001bc900001020000090020da00000900, protocol 1)</p>

Trap Tables

Trap Type	vrrpTrapNewMaster
Object ID	1.3.6.1.2.1.46.1.3.1.0.3
Description	The sending agent has transitioned from “Backup” state to “Master” state.
Bit Position (Word 2)	21
Hex Value (Word 2)	20 0000
Trap Text and Variable Descriptions	<div data-bbox="404 638 1263 720" style="border: 1px solid black; background-color: #e0e0e0; padding: 5px;">Agent has transitioned from Backup to Master state (If index 1, vrid 2)</div> <p style="text-align: center;">Interface Index Number. A unique value that identifies the sending agent.</p> <p style="text-align: center;">Virtual Router ID. The number that identifies the virtual router on this VRRP. Possible values range from 1 to 255.</p>

Trap Type	vrrpAuthFailure
Object ID	1.3.6.1.2.1.46.1.3.1.0.4
Description	A packet was received from a router whose authentication key or authentication type conflicts with this router's authentication key or type.
Bit Position (Word 2)	22
Hex Value (Word 2)	40 0000
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>A packet with a wrong authentication key or type is received (If index 1, vrid 2, source 192.168.10.1, error type 3)</p> </div> <div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>Packet Source IP. The IP address of an inbound VRRP packet.</p> </div> <div style="text-align: center;"> <p>Interface Index Number. A unique value that identifies the sending agent.</p> </div> </div> <div style="display: flex; justify-content: space-around; margin-top: 20px;"> <div style="text-align: center;"> <p>Virtual Router ID. The number that identifies the virtual router on this VRRP. Possible values range from 1 to 255.</p> </div> <div style="text-align: center;"> <p>Error Type. The type of configuration conflict. The following integers are valid values:</p> <ul style="list-style-type: none"> 1 Invalid authentication type 2 Mismatched authentication 3 Authentication Failure </div> </div>

Trap Tables

Trap Type	oamVCAIS
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.71
Description	The specified connection is in the VC-AIS state.
Bit Position (Word 3)	10
Hex Value (Word 3)	400
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex xylanOamF5VCVciIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 10px; margin-bottom: 10px;"> <p>The specified connection is in VC-AIS state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>Slot Number. The slot number for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>VPI. The virtual path identifier for the specified connection.</p> <p>VCI. The virtual circuit identifier for the specified connection.</p>

Trap Type	oamVCRDI
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.72
Description	The specified connection is in the VC-RDI state.
Bit Position (Word 3)	11
Hex Value (Word 3)	800
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex xylanOamF5VCVciIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>The specified connection is in VC-RDI state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>Slot Number. The slot number for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>VPI. The virtual path identifier for the specified connection.</p> <p>VCI. The virtual circuit identifier for the specified connection.</p>

Trap Type	oamVCLOC
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.73
Description	The specified connection is in the VC-LOC state.
Bit Position (Word 3)	12
Hex Value (Word 3)	1000
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex xylanOamF5VCVciIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 10px; margin-bottom: 10px;"> <p>The specified connection is in VC-LOC state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>Slot Number. The slot number for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>VPI. The virtual path identifier for the specified connection.</p> <p>VCI. The virtual circuit identifier for the specified connection.</p>

Trap Type	oamVCUnsuccessLoop
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.74
Description	The specified connection is in the Unsuccessful Loopback state.
Bit Position (Word 3)	13
Hex Value (Word 3)	2000
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex xylanOamF5VCVciIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>The specified connection is in VC-UnsuccessLoop state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>VPI. The virtual path identifier for the specified connection.</p> <p>VCI. The virtual circuit identifier for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>Slot Number. The slot number for the specified connection.</p>

Trap Type	oamVPAIS
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.75
Description	The specified VP connection is in the VP-AIS state.
Bit Position (Word 3)	14
Hex Value (Word 3)	4000
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 10px; margin-bottom: 10px;"> <p>The specified connection is in VP-AIS state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>Slot Number. The slot number for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>VPI. The virtual path identifier for the specified connection.</p> <p>VCI. The virtual circuit identifier for the specified connection.</p>

Trap Type	oamVPRDI
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.76
Description	The specified VP connection is in the VP-RDI state.
Bit Position (Word 3)	15
Hex Value (Word 3)	8000
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 10px; margin-bottom: 10px;"> <p>The specified connection is in VP-LOC state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>Slot Number. The slot number for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>VPI. The virtual path identifier for the specified connection.</p> <p>VCI. The virtual circuit identifier for the specified connection.</p>

Trap Type	oamVPLOC
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.77
Description	The specified VP connection is in the VP-LOC state.
Bit Position (Word 3)	16
Hex Value (Word 3)	1 0000
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>The specified connection is in VPUnsuccessLoop state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>VCI. The virtual circuit identifier for the specified connection.</p> <p>Slot Number. The slot number for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>VPI. The virtual path identifier for the specified connection.</p>

Trap Type	oamVPUnsuccessLoop
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.78
Description	The specified VP connection is in the unsuccessful loopback state.
Bit Position (Word 3)	17
Hex Value (Word 3)	2 0000
Variables	xylanOamF5VCSlotIndex xylanOamF5VCPortIndex xylanOamF5VCVpiIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 10px; margin-bottom: 10px;"> <p>The specified connection is in VP-RDI state. (Slot 1, Port 2, VPI 2, VCI 1)</p> </div> <p>Slot Number. The slot number for the specified connection.</p> <p>Port Number. The port number for the specified connection.</p> <p>VPI. The virtual path identifier for the specified connection.</p> <p>VCI. The virtual circuit identifier for the specified connection.</p>

Trap Tables

Trap	accountEvent
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.86
Description	An account event is generated to signal that a new accounting file is available on the switch
Bit Position (Word 3)	21
Hex Value (Word 3)	20 0000
Variable Description	chasAccountName - Path name of the most recently terminated accounting file. chasAccountFileCount - The number of terminated accounting files awaiting collection and removal by an external accounting collection agent.

Trap	Over1Alarm
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.87
Description	This alarm is generated when the filling level exceeds the first threshold. It signals that the switch changes over to the alternate collection device.
Bit Position (Word 3)	22
Hex Value (Word 3)	40 0000
Variable Description	chasAccountFilingLevel - The amount of buffer taken up by accounting data. Value shown as a percentage of the buffer size. chasAccountThreshold1 - The first filling level of the intermediate storage area for accounting data. Crossing this threshold generates a warning. Value shown as a percentage of the buffer size. chasAccountDeviceInUse - The IP address of the collection device with which a TCP connection was most recently established.

Trap Type	Under1Event
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.88
Description	This event is generated when the filling level goes below the first threshold. This event is for information only.
Bit Position (Word 3)	23
Hex Value (Word 3)	80 0000
Variable Description	chasAccountFilingLevel - The amount of buffer taken up by accounting data. Value shown as a percentage of the buffer size. chasAccountThreshold1 - The first filling level of the intermediate storage area for accounting data. Crossing this threshold generates a warning. Value shown as a percentage of the buffer size.

Trap	Over2Alarm
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.89
Description	This alarm is generated when the filling level exceeds the second threshold. It signals that the switch changes over to the alternate collection device.
Bit Position (Word 3)	24
Hex Value (Word 3)	100 0000
Variable Description	chasAccountFilingLevel - The amount of buffer taken up by accounting data. Value shown as a percentage of the buffer size. chasAccountThreshold2 - The second filling level of the intermediate storage area for accounting data. Crossing this threshold generates a warning. Value shown as a percentage of the buffer size. chasAccountDeviceInUse - The IP address of the collection device with which a TCP connection was most recently established.

Trap Tables

Trap	Under2Event
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.90
Description	This event is generated when the filling level is lowered below the second threshold.
Bit Position (Word 3)	25
Hex Value (Word 3)	200 0000
Variable Description	chasAccountFilingLevel - The amount of buffer taken up by accounting data. Value shown as a percentage of the buffer size. chasAccountThreshold2 - The second filling level of the intermediate storage area for accounting data. Crossing this threshold generates a warning. Value shown as a percentage of the buffer size.

Trap	Over3Alarm
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.91
Description	This event is generated when the filling level exceeds the third threshold. It signals that the switch is now in congestion.
Bit Position (Word 3)	26
Hex Value (Word 3)	400 0000
Variable Description	chasAccountFilingLevel - The amount of buffer taken up by accounting data. Value shown as a percentage of the buffer size. chasAccountThreshold3 - The third filling level of the intermediate storage area for accounting data. Crossing this threshold generates a warning. Value shown as a percentage of the buffer size. chasAccountDeviceInUse - The IP address of the collection device with which a TCP connection was most recently established.

Trap	Under3Event
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.92
Description	This event is generated when the filling level goes below the third threshold.
Bit Position (Word 3)	27
Hex Value (Word 3)	8000 0000
Variable Description	chasAccountFilingLevel - The amount of buffer taken up by accounting data. Value shown as a percentage of the buffer size. chasAccountThreshold3 - The third filling level of the intermediate storage area for accounting data. Crossing this threshold generates a warning. Value shown as a percentage of the buffer size.

Trap Type	NoDeviceAlarm
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.93
Description	This alarm is generated when the TCP connection establishment fails with both the primary and the secondary collection devices.
Bit Position (Word 3)	28
Hex Value (Word 3)	1000 0000
Variable Description	chasAccountDevicePrimary - The IP address of the primary collection device. chasAccountDeviceSecondary - The IP address of the secondary collection device.

Trap Tables

Trap	FileAlarm
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.94
Description	This alarm is generated when too many files are awaiting collection.
Bit Position (Word 3)	29
Hex Value (Word 3)	2000 0000
Variable Description	chasAccountFileCount - The number of terminated accounting files awaiting collection and removal by an external accounting collection agent.

Trap Type	fantrayEvent
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.1
Description	A fantrayEvent trap occurs when a problem condition is recognized on a chassis fan tray.
Bit Position (Word 3)	30
Hex Value (Word 3)	4000 0000
Variable Description	fantray1State - Status of fan tray 1. chasAccountDeviceSecondary - Status of fan tray 2.

Trap Type	ldpPeerCreate
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.80
Description	A LDP peer is identified by the LDP hello mechanism and a peer entry is created.
Bit Position (Word 3)	5
Hex Value (Word 3)	20
Variables	mplsLdpEntityID mplsLpdPeerIndex mplsLdpPeerID
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Peer Entity is Created. (EntityId 1, PeerIndex 2, PeerId 3)</p> </div> <p>EntityId. The identification number assigned to the new entity.</p> <p>PeerIndex. The index number assigned to the peer.</p> <p>PeerId. The identification number assigned to the peer.</p>

Trap Tables

Trap Type	ldpPeerDelete
Object ID	1.3.6.1.4.1.800.3.1.1.1.0.81
Description	An LDP peer is lost and the peer entry is deleted.
Bit Position (Word 3)	6
Hex Value (Word 3)	40
Variables	mplsLdpEntityID mplsLpdPeerIndex mplsLdpPeerID
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>Peer Entity is Deleted. (EntityId 1, PeerIndex 2, PeerId 3)</p> </div> <p>EntityId. The identification number of the deleted entity.</p> <p>PeerIndex. The index number of the deleted peer.</p> <p>PeerId. The identification number of the deleted peer.</p>

Trap Type	ldpSessionCreate
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.82
Description	An LDP session with the peer is established and a session entry is created.
Bit Position (Word 3)	17
Hex Value (Word 3)	80
Variables	mplsLdpEntityID mplsLpdPeerIndex mplsLdpPeerID mplsLdpSessionIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>LDP Session Created. (EntityId 1, PeerIndex 2, PeerId 3, Session Id 4)</p> </div> <p>EntityId. The identification number assigned to the new-entity.</p> <p>PeerIndex. The index number of the peer with which the session is created.</p> <p>PeerId. The identification number of peer with which the session is created.</p> <p>SessionId. The identification number of the new session.</p>

Trap Tables

Trap Type	ldpSessionDelete
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.83
Description	An LDP session with the peer is lost and the session entry is deleted.
Bit Position (Word 3)	8
Hex Value (Word 3)	100
Variables	mplsLdpEntityID mplsLpdPeerIndex mplsLdpPeerID mplsLdpSessionIndex
Trap Text and Variable Descriptions	<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p>LDP Session Deleted. (EntityId 1, PeerIndex 2, PeerId 3, Session Id 4)</p> </div> <p>EntityId. The identification number of the deleted entity.</p> <p>PeerIndex. The index number of the peer with whom the session entry was lost.</p> <p>PeerId. The identification number of the peer with whom the session entry was lost.</p> <p>SessionId. The identification number of the deleted session.</p>

Trap Type	lecStateChangeEvent
Object ID	1.3.6.1.4.1.800.3.1.1.4.0.96
Description	A trap message is sent to a network manager when a LEC status changes.
Bit Position (Word 2)	26
Hex Value (Word 3)	40 00000
Variables	lecID lecActualLanName lecAtmAddress, xylanLecSlotNumber xylanLecPortNumber xylanLecServiceNumber lecInterfaceState xylanReasonOfChange

<p>Trap Text and Variable Descriptions</p>	<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>ELAN Name. The name of the ELAN whose status changed to generate this trap.</p> </div> <div style="width: 45%;"> <p>Service Instance. The specific instance of this service. In most cases this value will be 1 but an ATM port may have multiple instances</p> </div> </div> <div style="text-align: center; margin: 10px 0;"> <div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px; display: inline-block;"> LEC Status Change (ELAN Name, Service Instance, New state, previous state). </div> </div> <div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>New State. The new, current status of the LEC that changed to generate this trap. Displayed as an integer as shown below in the State List.</p> </div> <div style="width: 45%;"> <p>Previous State. The previous status of the LEC that changed to generate this trap. Displayed as an integer as shown below in the State List.</p> </div> </div> <div style="margin-top: 20px;"> <p>State List</p> <ol style="list-style-type: none"> 1. none 2. timeout 3. undefined error 4. version not supported 5. invalid request parameters 6. duplicate LAN destination 7. duplicate ATM address 8. insufficient resources 9. access denied 10. invalid requester id 11. invalid LAN destination 12. invalid ATM address 13. no configuration 14. leconfigureError 15. insufficient information </div>
---	--

14 DNS Resolver and RMON

Introduction

This chapter describes commands related to the Domain Name Server (DNS) resolver and remote network monitoring (RMON) feature in the switch. This chapter also describes how to configure router port MAC addresses with the **chngmac** command.

The commands for these features are available from the Networking submenu, which is described in Chapter 25, "IP Routing."

Configuring the DNS Resolver

The Names Submenu

The **Names** command takes you to the Names submenu. The one command in this menu, **res**, is used to view and to configure the Domain Name Server (DNS) resolver. You can configure up to three Domain Name Servers. The switch searches all three servers until it resolves the name to an IP address or until it fails to find the name.

To display the **Names** submenu, enter the following command:

```
names
```

A screen similar to the following displays:

```
Command      Resolver Configuration Menu
-----
res          View/Configure the DNS resolver

Main   File   Summary  VLAN   Networking
Interface Security System  Services Help
```

To configure one or more Domain Name Servers, enter the following command:

```
res
```

If the resolver function has not been enabled, a screen similar to the following displays:

```
DNS Resolver Configuration

1) Resolver Enabled : No

Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :
```

Configuring the DNS Resolver

To enable the resolver function, enter **1=y**. A screen similar to the following then displays:

DNS Resolver Configuration

```
1) Resolver Enabled      : Yes
2) Domain                : UNSET
3) Server Address 1     : UNSET
4) Server Address 2     : UNSET
5) Server Address 3     : UNSET
```

Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :

The prompts allow you to enter a Domain Name and up to three Domain Name Servers (identified by their IP addresses).

- To change a value, enter the number corresponding to that value, an (=), then the new value. For example, to set a Domain Name to Company.Com, enter **2=Company.Com**.
- To clear an entry, specify the value as (.) as in **2=.**
- To save all your modifications, enter **save**
- To cancel all your modifications, enter **quit**
- To view the parameters currently configured, enter **?**

Remote Network Monitoring (RMON)

Remote Network Monitoring (RMON) allows you to set up remote monitoring within your Omni Switch/Router. RMON consists of “probes” and “events.” There are two commands in the Networking menu, **probes** and **events**, which you can use to monitor, activate and inactivate probes and events. Be aware that you cannot create probes from within the switch’s User Interface; to do so requires a network application such as HP ProbeView.

Probes and Events

A **probe** is a task that runs in the switch. By using probes instead of sending repetitive inquiries to the switch, network traffic is significantly reduced.

There are three different kinds of probes: Ethernet, History, and Alarm.

A network management station (NMS) can configure either History or Alarm probes (a maximum of 16 is allowed). The status of a probe can be one of the following:

- Creating - The probe is under creation.
- Active - The probe is active.
- Inactive - The probe is inactive.

An **event** is an action that takes place based on an alarm condition detected by a probe. The event can take the form of an SNMP trap message and/or a log entry describing the alarm.

Ethernet Probes

An Ethernet probe monitors a selected Ethernet interface (port) and tracks Ethernet statistics. An Ethernet probe is automatically created on each Ethernet interface that is enabled. If the interface becomes disabled, that Ethernet probe is deleted.

History Probes

A History probe keeps a running history of all the statistics it has collected. When you set up a history probe you assign a sampling interval and a total number of samples to be taken. It keeps this information in a set of rotating buffers, so that it always retains the most recent samples.

The sampling rate is configurable from 1 second to 3600 seconds (1 hour). The total number of samples is configurable, however, it is limited by system resources (memory) available. The more samples you request, the more system resources needed. You may request as many samples as you want but the system will only grant as many as it has available.

Alarm Probes

An Alarm probe generates an alarm if the variable you are monitoring exceeds a set limit.

To set up an Alarm probe you need to select a variable (Ethernet statistic) that you want to monitor. You set an upper and lower threshold that you will allow this variable to reach. If it crosses the threshold, an event is triggered which results in the sending of an SNMP trap and/or the logging of the alarm.

There are two ways an Alarm probe monitors variables. One is by absolute value. For example, if you set an upper limit of 100, an alarm will be generated if the variable exceeds 100. The other is a delta value where you can set the amount of change allowable; for example, you could set the delta range to 10. If the current sample differs from the previous sample by more than 10, an alarm will be generated.

The Alarm probe attempts to prevent a flood of alarms from being generated by fluctuating values. It does so by continuously comparing the upper and lower limits. What this means is that the first time either an upper or lower limit is exceeded, an alarm will be generated. However, if the variable moves back inside the limit, then out again, another alarm will not be generated unless the opposite limit is exceeded. For example, consider a situation where an upper limit of 75 and a lower limit of 25 is set. The variable goes to 76. An alarm is generated. If it drops to 74 then goes back up to 76, no alarm will be generated. Only when the variable drops below 25 will another alarm be generated. If it goes back up to 76 then another alarm will be generated, etc. This procedure prevents a flood of alarms from being generated if the value fluctuates between 74 and 76.

Monitoring Probes

The **probes** command is used to monitor, activate, and inactivate existing probes (remember, you cannot create probes in the switch's UI). You can do three things with the command:

1. View all the current probes.
2. View a specific probe.
3. Activate or inactivate a History or Alarm Probe. (You can only do this with the "admin" login.)

The **probes** command has three optional parameters. The format is:

probes [active | inactive] [n]

where:

active - activates an existing probe

inactive - inactivates an existing probe

n - is the entry number of the probe to view

If you enter the **probes** command without parameters, it displays all the current probes.

RMON Probe Summary

Entry	Slot/Port	Flavor	Status	Time	System Resources
1	2/ 1	Ethernet	Active	0 hrs 39 mins	312 bytes
2	2/ 1	History	Active	0 hrs 4 mins	3656 bytes
3	2/ 1	Alarm	Active	0 hrs 0 mins	1336 bytes

Entry

The entry number in the list of probes (1-16).

Slot/Port

The slot port number (interface) that this probe is monitoring.

Flavor

Ethernet, History, or Alarm.

Status

Creating, Active, or Inactive.

Time

Time since the last change in status.

System Resources

Amount of memory that has been allocated to this probe.

To see the detail for each of the probes enter the **probes** command followed by the entry number as shown below.

```
/Networking % probes 1
```

RMON Probe Summary

Entry	Slot/Port	Flavor	Status	Time	System Resources
1	2/ 1	Ethernet	Active	0 hrs 39 mins	312 bytes

Probe's Owner: Omni Switch/Router Ethernet probe on slot 2 port 1

```
/Networking % probes 2
```

RMON Probe Summary

Entry	Slot/Port	Flavor	Status	Time	System Resources
2	2/ 1	History	Active	0 hrs 4 mins	3656 bytes

Probe's Owner: andy

- History Control Buckets Requested = 60
- History Control Buckets Granted = 60
- History Control Interval = 60 seconds
- History Sample Index = 6

```
/Networking % probes 3
```

RMON Probe Summary

Entry	Slot/Port	Flavor	Status	Time	System Resources
3	2/ 1	Alarm	Active	0 hrs 0 mins	1336 bytes

Probe's Owner: andy

- Alarm Rising Threshold = 3000
- Alarm Falling Threshold = 3000
- Alarm Rising Event Index = 1
- Alarm Falling Event Index = 3
- Alarm Interval = 30 seconds
- Alarm Sample Type = delta value
- Alarm Startup Alarm = rising or falling alarm
- Alarm Variable = ethernet octets received

Monitoring Events

The **events** command has one optional parameter. The format is:

```
events [clear]
```

where:

clear - clears the event log. (You can only do this with the "admin" login.)

RMON Logged Events Summary

Entry	Time	Description
1	0 hrs 26 mins	Rising threshold alarm for etherStatsOctets on slot 2 port 1
2	0 hrs 27 mins	Rising threshold alarm for etherStatsOctets on slot 2 port 1

Configuring Router Port MAC Addresses

You can use the **chnghmac** command if you want to configure a locally administered address (LAA) for a group that has an IP router port, IPX router port, or both. To use this command, enter **chnghmac** followed by the number of the group you want to modify (the default group number is **1**).

◆ Important Note ◆

You must add **chnghmacFlag=1** to the end of the **mpx.cmd** file and then reboot the switch to use the **chnghmac** command. See Chapter 7, “Managing Files,” for information on editing system files.

For example, if you want to modify a MAC address in Group 2, you would enter:

```
chnghmac 2
```

at the system prompt. Something similar to the following would then be displayed:

```
Current MAC address is factory default
Enter Router Port's MAC address ([XXYYZZ:AABBCC]) :
```

Enter the router port MAC address. (It cannot be a multicast address.) If you enter an incorrect address, the following will be displayed:

```
Invalid input format -- usage [XXYYZZ:AABBCC].
```

and the **chnghmac** command will terminate. If you enter a correct address, the following would then be displayed:

```
Is MAC address in Canonical or Non-Canonical (C or N) [C] :
```

Enter **C** if the address is canonical or **N** if it is non-canonical (the default is canonical). Note that if you execute the **chnghmac** command again it will display the user-defined instead of “factory default.”

Restoring Router Port Mac Addresses

If you want to restore the MAC address to the factory default, enter **chnghmac** followed by the group number. When the system asks for the MAC address, enter **000000:000000**. For example, to restore router port configured MAC address 003030:000001 in Group 2 to the factory default, enter

```
chnghmac 2
```

at the system prompt. The following would then be displayed:

```
Configured MAC Address: Canonical    Non-Canonical
                        003030:000001 000c0c:000080
{Address 000000:000000 requests use of factory default}
Enter Router Port's MAC address ([XXYYZZ:AABBCC]) :
```

Note that the **chnghmac** command displayed the user-defined instead of “factory default.” Enter **000000:000000** at the prompt.

15 Managing Ethernet Modules

Overview of Omni Switch/Router Ethernet Modules

This chapter describes User Interface commands for Ethernet, Fast Ethernet, and Gigabit Ethernet modules.

This chapter documents User Interface (UI) commands to manage Omni Switch/Router Ethernet modules. For documentation on Command Line Interface (CLI) commands to manage Ethernet modules, see the *Text-Based Configuration CLI Reference Guide*.

◆ Important Notes ◆

In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. See Chapter 4, “The User Interface,” for documentation on changing from CLI mode to UI mode.

In Release 4.5 and later, Mammoth-based Ethernet and early-generation Ethernet modules are no longer supported.

Port Mirroring and Port Monitoring

Port Mirroring and Port Monitoring can be used on all Ethernet modules. Both Port Mirroring and Port Monitoring are described at the end of Chapter 19, “Managing Groups and Ports.”

Fast Ethernet Backbones

Fast Ethernet ports can be used as backbone links. The switch has two features that can improve the performance and flexibility of Ethernet backbones. OmniChannel aggregates the bandwidth of up to four (4) Fast Ethernet ports. This feature allows you to scale Fast Ethernet links from 100 Mbps to 800 Mbps in 100 Mbps increments. OmniChannel is described in *OmniChannel* on page 15-9. Fast Ethernet ports also support the 802.1Q tagging mechanism, enhancing the compatibility of ports with other vendors’ equipment. 802.1Q is described in Chapter 16, “Managing 802.1Q Groups.”

Gigabit Ethernet Modules

Gigabit Ethernet modules can be used as backbone links and used to support high-speed servers. Kodiak Gigabit Ethernet modules support 802.1Q hardware tagging. See Chapter 16, “Managing 802.1Q Groups,” for more information on 802.1Q hardware tagging for Gigabit Ethernet Modules.

◆ Note ◆

For Kodiak-based 10/100 Ethernet modules, 802.1Q is supported over OmniChannel. See Chapter 16, “Managing 802.1Q Groups” for more information.

Variety of Connector Options

Ethernet and Fast Ethernet modules are available in a variety of connector types. On the OmniSwitch, Fast Ethernet modules use copper RJ-45 and fiber SC connectors. On the Omni Switch/Router, 10/100 Ethernet modules use copper RJ-45 connectors and the ESX-100FM/FS-12W Fast Ethernet module uses fiber MT-RJ connectors.

On the OmniSwitch, Ethernet 10 Mbps modules are available with copper RJ-45, fiber SC, Telco (RJ-21), BNC, and AUI connectors. On the Omni Switch/Router, the 10 Mbps ESX-FM-24W uses fiber VF-45 connectors.

Gigabit Ethernet modules on the OmniSwitch and Omni Switch/Router use fiber SC connectors. Refer to Chapter 3, “Omni Switch/Router Switching Modules,” for information on Omni Switch/Router Ethernet hardware.

Three Generations of Modules

Ethernet modules in Release 3.1 and later contained advanced chip technology referred to as “Mammoth.” This technology boosted the port density of modules, increasing the port count available in each chassis. The Mammoth technology also included ports with 10/100 autosensing capability. This generation of Ethernet modules also uses a different set of software commands to configure and monitor ports.

Ethernet modules in Release 4.3 and later contain another chip technology referred to as “Kodiak.” The new Kodiak-based modules combine several features of the Mammoth and early Ethernet modules. They support priority VLANs with 4 separate levels of priority; in addition, ESX-K Series Kodiak-based Ethernet modules support the addition of a server version of the OmniChannel. For information on priority VLANs, see Chapter 19, “Managing Groups and Ports.” For information on OmniChannel and Server Channel features, see *OmniChannel* on page 15-9.

The following table outlines the Kodiak Ethernet modules.

◆ Important Note ◆

In Release 4.5 and later, Mammoth-based Ethernet and early-generation Ethernet modules are no longer supported.

Kodiak Ethernet Modules

Ethernet Module (Chassis Type)	Speed Supported (per port)	Software Configurable?	Commands Available	OmniChannel Supported?
ESX-K-100C-32W (Omni Switch/Router)	10/100 Mbs	Yes	10/100cfg 10/100vc	Yes
ESX-K-100FM/FS-16W (Omni Switch/Router)	100 Mbs	Yes	10/100cfg 10/100vc	Yes
GSX-K-FM/FS-2W (Omni Switch/Router)	1000 Mbs	Yes	10/100cfg 10/100vc	No

ESX-K Series Modules and Optimized Ports

Kodiak-based modules will flood packets with unknown destination addresses on ports configured for optimized device mode. To prevent this condition, the following command can be entered into the `mpx.cmd` file:

```
MamOptSwitchPorts=1
```

If the port is set to optimized and has not learned a MAC address, it will flood these packets out regardless if the above condition is used. If the above flag is set, the port will not flood multicast packets.

◆ Note ◆

For information on editing the `mpx.cmd` text files, see Chapter 7, “Managing Files.”

Port Partitioning

Ethernet10BaseT, 10/100BaseT and 100BaseF boards can detect certain cabling errors and/or physical media misconfigurations which could lead to multiple retries or reception of multiple spurious frames, affecting performance of attached devices. In this event, the system will partition the affected port, which will be marked in the `vi` menu with Inactive (**Inactv**) operational status. (See Chapter 19, “Managing Groups and Ports,” for information about the `vi` command.) If a cable drop is detected, the system will remove the partitioned state, bringing the port back into a normal state once the link is detected.

If the original cabling problem has not been corrected, the link may become partitioned again. In this event, normal operation will be enabled when the problem has been corrected.

The Ethernet Management Menus

The **eth100** and **10/100** sub-menus are described in this chapter. These sub-menus are part of the physical interface sub-menu, which appears similar to the following display:

Command	Physical Interface Menu
slipc	Configure SLIP (Serial Line IP) on a TTY Port
atm	Enter the ATM Management sub-menu
eth100	Enter the 100BaseT sub-menu
10/100	Enter the 10/100BaseT sub-menu
tok	Enter the Token Ring Management sub-menu

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The **eth100** sub-menu contains commands for early generation Fast Ethernet modules. The **10/100** sub-menu has commands for Kodiak Ethernet modules.

When you enter **eth100** at a system prompt, you enter the early generation Fast Ethernet sub-menu. This sub-menu displays as follows:

Command	100BaseT Menu
eth100vc	View 100BaseT Port Configuration Table
eth100cfg	Configure 100BaseT Port Parameters

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

◆ Important Note ◆

In Release 4.5 and later, early-generation Ethernet modules are no longer supported.

When you enter **10/100** at a system prompt, you enter the Kodiak Ethernet configuration sub-menu. This sub-menu displays as follows:

Command	10/100 Menu
10/100vc	View 10/100 Port Configuration Table
10/100cfg	Configure 10/100 Port Parameters
crechnl	Create a Fast Ethernet Channel
delechnl	Delete a Fast Ethernet Channel
addprtchnl	Add port/s to a fast Ethernet Channel
delprtchnl	Delete port/s from a fast Ethernet Channel
chnlinfo	Display channel configuration parameters

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

Descriptions for these commands begin on page 15-5. The commands in this sub-menu below **crechnl** are used to configure OmniChannel; documentation for OmniChannel begins on page 15-9.

Configuring 10/100 Auto-Sensing Ports

The **10/100cfg** command allows you to enable auto-negotiation, as well as configure link speed (10 or 100 Mbps) and the link mode (full or half-duplex) on 10/100 Ethernet ports on the ESX-K-100C-32W modules on the Omni Switch/Router.

Follow these steps to configure a 10/100 port:

1. Enter **10/100cfg** at the system prompt and press **<Enter>**.
2. The system displays a prompt asking for the slot and port number:

Enter Slot/Interface :

Enter the slot number, a slash (/), and the port number of the Ethernet port that you want to configure. Press **<Enter>**.

3. The system prompts you to enable or disable auto-sensing:

Autonegotiate [y,n, or quit] (Currently enabled (y)) :

Enter **y** to enable auto-negotiation or **n** to disable auto-negotiation. Auto-negotiation can be used to determine the link speed *and* the link mode (full or half) of the connection.

If you choose **y** to enable auto-negotiation, the system will automatically detect whether the connection speed of the attached device is 10 Mbps or 100 Mbps. It can also determine whether the link mode of the connection is half- or full-duplex.

◆ Note ◆

Auto-negotiated ports on GSX modules display inactive ports as 1000 Mbps/full duplex.

If you enable auto-negotiation, continue with Step 6.

If you choose **n** to disable auto-negotiation, then you will be prompted for the Line Speed. Continue on with the next step.

4. If you chose to disable auto-sensing, then the following prompt displays showing the current line speed:

Line Speed [100 or 10] (Currently 100) :

Select whether you want the port to operate at 10 Mbps or 100 Mbps. The port will operate at this speed until you change it through the **10/100cfg** command later. Press **<Enter>** after you enter the Line Speed. The new line speed will take effect; no reboot is required. Continue with the next step.

5. The following prompt displays, showing the current link mode:

Link Mode [Full, Half] (Currently (H)alf Duplex) :

Enter **F** to set the port to full-duplex mode or **H** to set the port to half-duplex mode. In full-duplex mode, the full 100 or 10 Mbps of bandwidth is used for data traveling on each direction of the cable. Press **<Enter>** after you enter the Mode. The new mode will take effect; no reboot is required. You have completed the configuration of this port.

6. Since you have enabled auto-negotiation, the port will automatically sense the line speed of the connection. You can also further enable auto-negotiation for the link mode. When the following prompt displays:

Link Mode [Half or Auto] (Currently (H)alf Duplex) :

select whether you want the port to auto-sense the duplex mode (**Auto**) or whether you want the port to default to half-duplex mode (**Half**). Enter an **A** for auto-sensing or enter an **H** for half-duplex.

If you set the mode to half-duplex, then the port will always run in half-duplex. If you set the mode to **Auto**, then the port will automatically detect whether the connection is half- or full-duplex and then operate in that mode. You have completed the configuration of this port.

Connecting Kodiak Modules to Non-Auto-Negotiating Links

The ESX-K-100C-32W can auto-negotiate link speed. However, if you hard-configure (auto-negotiation disabled) a Kodiak 10/100 module port for 10 Mbps, then you should not connect that port to a non-auto-negotiating 100 Mbps port or device.

Configuring Kodiak Ethernet Ports

The **10/100cfg** command allows you to configure the link mode (full or half-duplex) for ports on newer Kodiak Ethernet modules.

This procedure describes how to configure Ethernet modules on the Omni Switch/Router.

Follow these steps to configure a Kodiak Ethernet port:

1. Enter **10/100cfg** at the system prompt and press **<Enter>**.
2. The system displays a prompt asking for the slot and port number:

Enter Slot/Interface :

Enter the slot number, a slash (/), and the port number of the Ethernet port that you want to configure. Press **<Enter>**.

3. The following prompt displays, showing the current link mode:

Link Mode [Full, Half] (Currently (H)alf Duplex) :

Enter **F** to set the port to full-duplex mode or **H** to set the port to half-duplex mode. In full-duplex mode, the full 100 or 10 Mbps of bandwidth is used for data traveling on each direction of the cable. Press **<Enter>** after you enter the Mode. The new mode will take effect; no reboot is required.

Viewing Configurations for 10/100 Ethernet Modules

The **10/100vc** command allows you to view the current status of newer Ethernet modules (see *Kodiak Ethernet Modules* on page 15-3). These modules support 100 Mbps, or 1000 Mbps Ethernet. Ethernet 10/100 ports (e.g., ESX-K-100C-32) can auto-sense the connection speed of the attached device.

Entering **10/100vc** displays information similar to the following:

10/100 Configure Values for all slots

Slot/ Intf	Auto- negotiate	DETECTED		SET	
		Line Speed	Duplex Mode	Line Speed	Duplex Mode
5/ 1	enabled	?	?	auto	half-d
5/ 2	enabled	10	HALF-D	auto	half-d
5/ 3	enabled	100	HALF-D	auto	half-d
5/ 4	enabled	100	HALF-D	auto	half-d
5/ 5	enabled	?	?	auto	half-d
5/ 6	enabled	10	HALF-D	auto	half-d
5/ 7	enabled	100	HALF-D	auto	half-d
5/ 8	enabled	?	?	auto	half-d

Slot/Intf. The slot and port number (Intf) where this Ethernet port is located.

Auto-negotiate. Indicates whether auto-negotiation is enabled on a 10/100 port. If enabled, the port will automatically sense whether the attached device operates at 10 Mbps or 100 Mbps and adjust accordingly. If disabled, the port does not automatically detect the connection speed and instead uses the line speed you configure through the **10/100cfg** command. You enable or disable auto-negotiation through **10/100vcfg**. A value of **n/a** in this column means the port does not support auto-sensing and the line speed defaults to either 10 or 100 Mbps.

The next set of columns are divided into DETECTED and SET. The columns under DETECTED are the current operational **Line Speed** or **Duplex Mode**. The columns under SET are the configured values; these configured values will either be defaults or the values configured through **10/100cfg**.

Line Speed. Indicates the speed (in Mbps) at which the port is currently operating (DETECTED) or configured to operate (SET).

DETECTED values will be **10** (Mbps), **100** (Mbps), or a question mark (?). A question mark (?) in this column indicates the port is not connected to a device.

SET values will be **auto**, **10** (Mbps,) or **100** (Mbps). The **auto** setting means auto-sensing is enabled and the Line Speed will equal the speed for which the attached device is configured.

Duplex Mode. Indicates whether the port is operating (DETECTED) or configured (SET) for half- or full-duplex mode.


DETECTED values will be half-duplex (**HALF-D**), full-duplex (**FULL-D**), or a question mark (?). A question mark (?) in this column indicates the port is not connected to a device.

SET values will be auto-sensing (**auto**), half-duplex (**half-d**), or full-duplex (**full-d**). If this value is **auto**, then the switch automatically sets the duplex mode to the network device's setting. If this value is **half-d**, then the port will always run in half-duplex mode. If this value is **full-d**, then the port will always run in full-duplex mode. You configure the duplex mode through the **10/100cfg** command. Note that you can only configure a 10/100 port for full-duplex if you disable auto-sensing.

OmniChannel

OmniChannel allows you to increase the bandwidth of Fast backbones by combining the capacity of up to four (4) Fast Ethernet ports into one channel. The combined channel operates within Spanning Tree as one virtual port, and can provide up to 800 Mbps (in full-duplex mode) of bandwidth. (In full-duplex mode, 400 Mbps is supported in each direction of the OmniChannel.) This feature is useful for Ethernet-intensive networks that need to increase bandwidth capacity without setting up ATM backbones using OC-3 or OC-12 connections.

The OmniChannel feature operates on 10/100 and 100 Mbps Ethernet ports employing Kodiak chip technology, such as those modules listed in the table, *Kodiak Ethernet Modules* on page 15-3. OmniChannel does not operate on 10 Mbps ports or on early-generation Fast Ethernet ports.



Up to Four 100 Mbps Links May Comprise an OmniChannel Backbone


◆ Note ◆

For Kodiak-based 10/100 Ethernet modules, 802.1Q is supported over OmniChannel. See Chapter 16, “Managing 802.1Q Groups” for more information.

OmniChannel balances the traffic load among links by MAC address. MAC addresses are assigned to physical links in the OmniChannel in a round-robin fashion. The first MAC address learned will transmit and receive data on the first link. The second MAC address learned will transmit and receive over the second link, and so on regardless of the bandwidth requirements of each MAC address.

The Server Channel Feature

For ESX-K Series Kodiak-based Ethernet boards, you can create an OmniChannel that connects to a server instead of another Omni Switch/Router. The intention of the Server Channel is to give the user the option to increase the bandwidth between a server and Omni Switch/Router for more client request support. This functionality is especially useful for internet servers such as B2C and B2B servers.



Up to Four 100 Mbps Links May Comprise a Server Channel backbone

Server Channel Limitations

The following are limitations to creating a server channel on the Omni Switch/Router:

- The maximum number of Server Channels in the whole box is not fixed; however, it is suggested that no more than 16 be created on the same switch.
- Each Server Channel can support up to 4 ports.
- Within one Server Channel, all of channel ports must be on the same slot.
- Within one Server Channel, all of channel ports must be in one VLAN.
- A port cannot be configured as Server Channel and Omni Channel port at the same time.
- Currently, Server Channel cannot be used with 802.1Q.

Creating an OmniChannel

You use the **crechnl** command to create an OmniChannel. Follow these steps:

1. Enter **crechnl**.
2. The following prompt displays:

Channel Number (2):

Enter the identification number you want to assign to this OmniChannel. By default, the software lists the next available channel number in parentheses. (In this example, the next available channel number is **2**.) If you want to select the default, simply press **<Enter>**. Otherwise, enter the desired channel number and press **<Enter>**.

3. The following prompt displays:

Channel type (1) omni_chnl (2) server_chnl

If the far end of the link is another Omni Switch/Router, you need to create an OmniChannel. Select **1** and proceed to the next step. If the far end of the link is a server, select **2** to create a Server Channel.

4. The following prompt displays:

To select a port, use the convention - Slot/Physical Port.

For eg. 2/1 is used to select Physical Port 1 on Slot 2

Primary Slot/Port:

Enter the slot and port that the switch will initially use as the Spanning Tree virtual port for this channel. Each OmniChannel is considered a single virtual port within the network, so only one physical port will participate in Spanning Tree.

◆ Note ◆

After a reboot or after a loss of a connection, the first port in an OmniChannel that the switch brings up will become the primary port. Therefore, one of the ports you choose as the secondary port (explained in Step 5 below) could become the primary port and thus participate in Spanning Tree.

If the port you enter is already part of another OmniChannel, then it cannot be used in a second OmniChannel. The following message displays for those ports that are already part of another OmniChannel:

Primary port in use

5. The following prompt displays:

To select a port, use the convention - Slot/Physical Port.

For eg. 2/1 is used to select Physical Port 1 on Slot 2

Secondary Slot/Port:

Enter the other ports that will be used in this OmniChannel. Up to four (4) Fast Ethernet Ports may participate in an OmniChannel. Therefore, you can specify up to three (3) additional ports which will initially become secondary ports. These secondary ports must be on the same module as the primary port. Secondary ports do not participate in the Spanning Tree algorithm; they are used for data transmission only.

◆ Note ◆

As explained in Step 4 above, a port that you initially configure as a secondary port can become a primary port.

Specifying a Range of Ports. To specify a range of ports, enter the slot number, a slash (/), the port number for the first secondary port, a dash (-), and the port number for the last secondary port. For example, to specify ports 3, 4, and 5 on the Fast Ethernet module in slot 2 as secondary ports in an OmniChannel, you would enter:

2/3-5

Specifying Multiple Ports. To specify multiple ports (on the same module) that are not physically contiguous, enter the slot number, a slash (/), the port number for the first secondary port, a comma (,), and then the slot and port for the next secondary port. For example, to specify ports 3 and 5 on the Fast Ethernet module in slot 2, you would enter:

2/3, 2/5

The order in which you specify secondary ports is important. In the event of a failure on the primary port, the first secondary port specified will become the primary port in the OmniChannel and participate in Spanning Tree.

Messages will display, informing you that secondary ports were saved in flash memory:

Successfully saved sec port in flash

Successfully saved sec port in flash

Adding Ports to an OmniChannel

After you create an OmniChannel with the **crechnl** command, you can add more secondary ports to the same channel as long as the channel contains less than 4 ports. You use the **addprtchnl** command to add ports to an OmniChannel. Follow these steps:

1. Enter **addprtchnl**.
2. The following prompt displays:

Channel Number :

Enter the channel number to which you want to add secondary ports. You can check the current port assignments for a given OmniChannel by using the **chnlinfo** command, which is described in *Viewing OmniChannel Parameters* on page 15-14.

3. The following prompt displays:

**To select a port, the convention - Slot/Physical Port or Slot/Phy.
Port Range. For eg. 2/1 is used to select Physical Port 1 on Slot
2 and 2/2-4 selects physical ports 2,3 and 4 on Slot 2
Slot/Port(s):**

Enter the additional ports that will be part of this OmniChannel. All the ports you enter will initially be secondary ports (i.e., they do not participate in the Spanning Tree algorithm and are used for data transmission only). You can specify up to 4 ports on an OmniChannel; only 3 of the ports can be secondary ports.

Specifying a Range of Ports. To specify a range of ports, enter the slot number, a slash (/), the port number for the first secondary port, a dash (-), and the port number for the last secondary port. For example, to specify ports 3, 4, and 5 on the Fast Ethernet module in slot 2 as secondary ports in an OmniChannel, you would enter:

2/3-5

Specifying Multiple Ports. To specify multiple ports (on the same module) that are not physically contiguous, enter the slot number, a slash (/), the port number for the first secondary port, a comma (,), and the slot and port for the next secondary port. For example, to specify ports 3 and 5 on the Fast Ethernet module in slot 2, you would enter:

2/3, 2/5

Messages will display, informing you that secondary ports were saved in flash memory:

**Successfully saved sec port in flash
Successfully saved sec port in flash**

Deleting an OmniChannel

You can delete any existing OmniChannel through the **delchnl** command. Follow these steps:

1. Enter **delechnl**.
2. The following prompt displays:

Channel to be deleted:

Enter the channel number that you want to delete. You can obtain information on a channel through the **chnlinfo** command, which is described in *Viewing OmniChannel Parameters* on page 15-14. Press **<Enter>** and the channel, along with all port assignments, will be deleted.

Deleting Ports from an OmniChannel

You can delete ports from an OmniChannel using the **delprtchnl** command. Follow these steps:

1. Enter **delprtchnl**.
2. The following prompt displays:

Channel Number :

Enter the channel number on which you want to delete ports. You can check the current port assignments for a given OmniChannel by using the **chnlinfo** command, which is described in *Viewing OmniChannel Parameters* on page 15-14.

3. The following prompt displays:

**To select a port, the convention - Slot/Physical Port or Slot/Phy.
Port Range. For eg. 2/1 is used to select Physical Port 1 on Slot
2 and 2/2-4 selects physical ports 2,3 and 4 on Slot 2
Slot/Port(s):**

Enter the port(s) that you want to delete from this OmniChannel.

Important Note

If you delete the primary port a secondary port will become the new primary port. The secondary port that will take over this role is the first secondary port specified through the **crechnl** command.

Deleting a Range of Ports. To delete a range of ports, enter the slot number, a slash (/), the port number for the first port, a dash (-), and the port number for the last port. For example, to delete ports 3, 4, and 5 on the Fast Ethernet module in slot 2, you would enter:

2/3-5

Deleting Multiple Ports. To delete multiple ports (on the same module) that are not physically contiguous, enter the slot number, a slash (/), the port number for the first port, a comma (,), and the slot and port for the next port. For example, to delete ports 3 and 5 on the Fast Ethernet module in slot 2, you would enter:

2/3, 2/5

Viewing OmniChannel Parameters

You can view the current configuration parameters and port assignments for an OmniChannel by using the **chnlinfo** command. Follow these steps:

1. Enter **chnlinfo**.
2. The following prompt displays:

Enter channel number for which information is required:

Enter the channel number for which you want to view information. If you want to view information on all OmniChannels in the switch, simply press **<Enter>**.

3. A screen similar to the following displays:

Displaying channel 2			
Channel Id	Phy. Port	Port Status	Mac Count
2	5/6	Inactive	0
	5/7	Inactive	0
3	5/3	Active	35
	5/4	Active	34
	5/5	Active	34

The following sections describe the variables in this table.

Channel Id. The identification number assigned to this OmniChannel during the **crechnl** configuration procedure.

Phy. Port. The physical slot and port number for all ports included in the OmniChannel. The slot number is listed first, then a slash (/), and the port number on the Ethernet module.

Port Status. The current operational status of this physical port. If the port is **Active**, then a cable is connected and data is capable of passing to and from the port. If the port is **Inactive**, then a cable may not be attached or the port is inoperational for hardware or software reasons.

Mac Count. The current number of MAC addresses that have been learned on this port. A separate MAC count is given for each physical port in the OmniChannel.

16 Managing 802.1Q Groups

This chapter documents User Interface (UI) commands to manage 802.1Q groups. For documentation on Command Line Interface (CLI) commands to manage 802.1Q groups, see the *Text-Based Configuration CLI Reference Guide*.

◆ Important Notes ◆

In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. See Chapter 4, “The User Interface,” for documentation on changing from CLI mode to UI mode.

In Release 4.5 and later, Mammoth-based Ethernet modules are no longer supported.

802.1Q is an IEEE standard for sending frames through the network tagged with VLAN identification. Alcatel has developed its own implementation of VLANs that closely follows the IEEE standard (and enhances it). However, Alcatel VLANs and 802.1Q VLANs cannot interoperate without special configuration.

If your network uses 802.1Q tagging, you will need to create 802.1Q groups and specify ports that will handle 802.1Q traffic. This can be done for 10/100, Fast Ethernet and Gigabit Ethernet Kodiak ASIC-based modules. Up to 64 groups can be supported using multiple spanning tree on an 802.1Q link for Kodiak ASIC-based Fast Ethernet and Gigabit Ethernet modules.

For Release 4.4 and later, Kodiak ASIC-based 10/100 Ethernet modules support 802.1Q traffic over OmniChannel in multiple spanning tree mode. However, you must first create an OmniChannel before creating 802.1Q groups. See Chapter 15, “Managing Ethernet Modules” for information about OmniChannel. See *Single vs. Multiple Spanning Tree* on page 16-4 for information on single and multiple spanning tree.

Support for 802.1Q in the Omni Switch/Router allows you to set up port-based groups that interoperate with 802.1Q-compliant equipment from other networking vendors.

Ports added to an 802.1Q group are done using Ethernet switch services. When using the service commands to add ports to an 802.1Q group, multiple spanning tree instances on a single port are supported. See *Single vs. Multiple Spanning Tree* on page 16-4 for additional information on the differences between single and multiple spanning tree.

The 802.1Q specification defines *trunk* and *access* ports (and links). Trunk links are LAN segments used for multiplexing VLANs between VLAN bridges. All devices that are directly connected to a trunk link must be VLAN-aware. Access links are LAN segments used to multiplex one or more VLAN-unaware devices into a port of a VLAN bridge. (This also includes a hybrid with some tagged and some untagged Groups.)

◆ Note ◆

The use of the word *trunk* in this document should not be confused with the IEEE use of *trunking* with link aggregation (such as OmniChannel and IEEE 802.3ad). The general meaning of a trunk is an inter-switch link over which different types of traffic are multiplexed.


IEEE 802.1Q Sections Not Implemented

Some portions of the 802.1Q specification have not yet been implemented in the Omni Switch/Router. These include the following:

- The tunneling of non-canonical 802.5 frames is not supported, since the Alcatel Omni S/R handles such traffic by frame translations. This tunneling mode of operation involves the Token Ring Encapsulation Flag in the 802.1Q header. It is not set or interpreted in the Alcatel Omni S/R implementation.
- The Alcatel Omni S/R implementation does not support the SNAP-encoded Tag Header (which is intended for Token Ring LANs). Only the Ethernet-encoded 4-byte Tag Header is supported (and only Ethernet LANs are supported).
- Alcatel Omni S/R does not support the Generic Attribute Registration Protocol (GARP) Multicast Registration Protocol (GMRP) and GARP VLAN Registration Protocol (GVRP) that are defined in 802.1Q.

Application Example

The following diagram illustrates a simple 802.1Q application:



Simple 802.1Q Application

In the above diagram, the PC devices (endstations) need to be segmented into different 802.1Q VLANs. The switch port to which each device attaches is assigned to an 802.1Q group (Group 2 for endstations A, B, E, and F, and Group 3 for endstations C, D, G, and H).

The ports connecting Switch X and Switch Y are also added to 802.1Q groups 2 and 3. All of the switch ports that handle 802.1Q traffic are now capable of passing 802.1Q information.

Prior to Release 4.4, only Mammoth ASIC-based Ethernet, Fast Ethernet and Gigabit Ethernet modules could be part of an 802.1Q group. For Release 4.4 and later, Kodiak ASIC-based 10/100, Fast Ethernet and Gigabit Ethernet modules also support 802.1Q groups. In either configuration, existing policies for a group will not be affected by the group's support for 802.1Q.


◆ Important Note ◆

Kodiak ASIC-based 10/100 Ethernet modules support 802.1Q traffic over OmniChannel in multiple spanning tree mode. However, for 802.1Q support over OmniChannel, you must first create an OmniChannel before creating 802.1Q groups. See Chapter 15 for information about OmniChannel. For information on the differences between single and multiple spanning tree, see *Single vs. Multiple Spanning Tree* on page 16-4.

By matching switch ports with 802.1Q groups, you are statically assigning the port to the group. Once assigned, an 802.1Q port cannot be dynamically assigned to another group. However, the same switch port can be statically assigned to more than one 802.1Q group.

Single vs. Multiple Spanning Tree

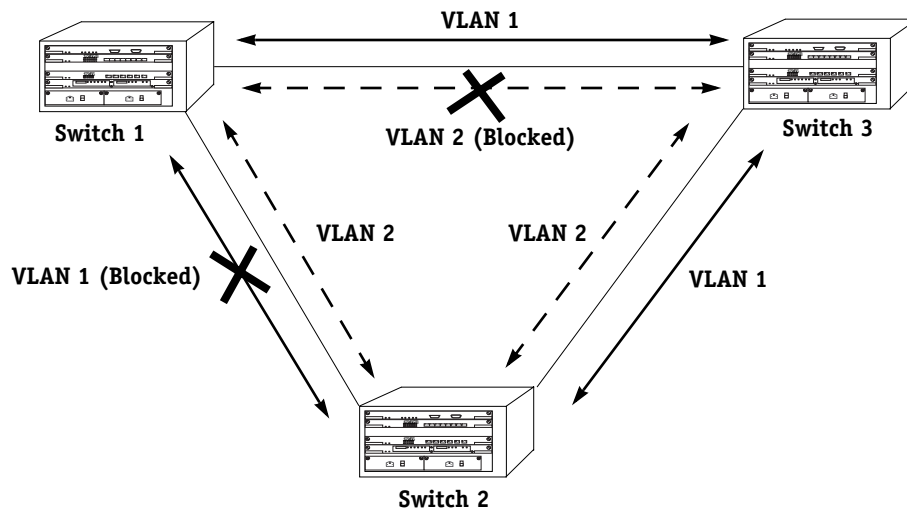
In previous releases of the Omni Switch/Router software (4.0 and earlier), spanning tree support was done on a per port basis. In other words, a physical port could only participate in one instance of a spanning tree on the network. If a network is passing both untagged and IEEE tagged frames, single spanning tree support could lead to packets being lost. Lost packets could occur if a port specifically assigned to handle one type of traffic (e.g., IEEE 802.1Q) is blocked by spanning tree, forcing traffic for that port to move to a port not assigned to handle IEEE 802.1Q traffic.



Port Based Spanning Tree

In the above diagram, the physical connection between Switch 1 and Switch 2 is blocked by spanning tree. No traffic can pass over the connected ports.

Release 4.1 (and later) of the Omni Switch/Router allows for multiple spanning tree instances on a single port. Put another way, a port can be part of separate spanning trees, with no impact on packet delivery. This is done by basing spanning tree configuration on groups rather than physical ports.




Group Based Spanning Tree

The above diagram shows how traffic on VLAN 1 is blocked between Switch 1 and Switch 2,

while VLAN 2 traffic is allowed to pass. The reverse is true for Switch 1 and Switch 3 (i.e., VLAN 2 traffic is blocked, while VLAN 1 traffic is allowed to pass).

Service commands are used in Ethernet modules to assign groups to 10/100 and Gigabit ports. The **cas**, **das**, **mas**, and **vas** commands create, delete, modify, and view trunk services created to handle 802.1Q traffic over an Ethernet backbone. This trunk service, coupled with the default bridging service, allows you to pass both tagged and untagged frames over the same port.

The following diagram shows the logical structure of the trunked 802.1Q groups:



Logical Configuration of Multiple Groups on a Single Port

In the above diagram, Groups 2 and 3 have been trunked to the physical port with an 802.1Q trunking service.

Since spanning tree is group based, the physical port in the above diagram participates in three spanning tree instances: one for untagged traffic and two for 802.1Q tagged traffic. Both types of frames can now pass through the same port.

◆ **Important Notes** ◆

Since a trunk is a service, and Alcatel switches have a 16 (10/100) or 15 (Gigabit) services per port limit, only 15 or 14 802.1Q groups can be added to the same port. In both cases, a default bridge service occupies one of the service slots.

For Kodiak ASIC-based Fast Ethernet and Gigabit Ethernet modules, up to 64 groups are supported using multiple spanning tree on an 802.1Q link. To support 64 groups, the following lines should be added into the `mpx.cmd` file :

```
MaxEthQGroups=64  
MaxGigaQGroups=64
```

See Chapter 7, “Managing Files,” for more information on editing text files.

Assigning an 802.1Q Group to a Port

Previous versions of the Omni Switch/Router (version 4.0 and earlier) only allowed for single spanning tree configured 802.1Q groups using the **addqgp**, **viqgp**, and **delqgp** menu commands. These commands were invalidated in the 4.1 release and replaced by the **cas**, **mas**, **vas**, and **das** service commands.

The procedure for assigning an 802.1Q group to a port is slightly different, depending on whether the port is a 10/100 or Gigabit Ethernet module port. (For additional information on Gigabit and Kodiak-based Ethernet modules, see Chapter 15, “Managing Ethernet Modules.”) Up to 64 groups can be supported using multiple spanning tree on an 802.1Q link for Kodiak ASIC-based Fast Ethernet and Gigabit Ethernet modules.

◆ Important Notes ◆

For Release 4.4 and later, Kodiak ASIC-based 10/100 Ethernet modules support 802.1Q traffic over OmniChannel in multiple spanning tree mode. However, you must first create an OmniChannel before creating 802.1Q groups. See Chapter 15, “Managing Ethernet Modules” for information about OmniChannel.

For information about the differences between single and multiple spanning tree, see *Single vs. Multiple Spanning Tree* on page 16-4.

In most of the procedures described in this section, the screens displayed vary, depending on what type of board and ASIC you are using. By viewing the front panel of your module, it should be easy to determine which procedure applies to you.

Ethernet modules are designated by ESX-K. Gigabit modules are designated by either GSX-K. Modules with a **K** on the front panel are Kodiak ASIC-based modules. For example, a module with designation **GSX-K** is a Gigabit module using a Kodiak ASIC.

For information on assigning an 802.1Q group to a 10/100 port, see *Configuring 802.1Q on 10/100 Ethernet Ports* on page 16-8. For information on assigning an 802.1Q group to a Gigabit port, see *Configuring 802.1Q on Gigabit Ethernet Ports* on page 16-11.

◆ Note ◆

802.1Q Omni Switch/Router tagging does not work with OmniCore 5200 tagging unless the OmniCore software is version 3.0.19 or later.

Configuring 802.1Q on 10/100 Ethernet Ports

Use the **cas** command to assign 802.1Q groups to 10/100 ports. To use this command, follow the steps below.

1. Enter **cas** at the system prompt, as shown:

```
cas <slot>/<port>
```

where **<slot>** is the slot of the module, and **<port>** is the port number that is to be added to the group. For example, to add port 3 on slot 5, you would enter:

```
cas 5/3
```

2. If you have a legacy 10/100 board, the following screen displays:

```
Slot 3 Port 5 Ethernet 802.1Q Service
1) Description          :
2) Group ID            :
3) Tag                 :
4) Priority            :
5) Mode
   Multiple Spanning Tree (3)
   Single Spanning Tree (4) :
```

If you have a Kodiak 10/100 board, the following screen displays:

```
Slot 3 Port 5 Ethernet 802.1Q Service
1) Description          :
2) Group ID            :
3) Tag                 :
5) Mode
   Multiple Spanning Tree (3)
   Single Spanning Tree (4) :
```

You can modify the parameters by entering the line number, an equal sign, and the value for the parameter. For example, to change the **Group ID** to **5**, you would enter **2** (the line number for **Group ID**), an equal sign (=), and a **5** (the group number), as shown:

```
2=5
```

3. Remember to save your changes by entering **save** at the system prompt when you have finished with the configuration.

◆ Important Notes ◆

Because 802.1Q support over OmniChannel is supported only in **Multiple Spanning Tree** mode on Kodiak 10/100 Ethernet boards, the **Mode** screen option is not configurable for this feature.

For 802.1Q support over OmniChannel, you must first create an OmniChannel before creating 802.1Q groups. See Chapter 15, “Managing Ethernet Modules” for information about OmniChannel.

The following sections describe the parameters shown in the screen on the preceding page.

Description

A textual description (up to thirty characters) for the service created when adding the port to a group.

Group ID

The number of the group to which the port is to be added.

Tag

A simple identifier that is added to 802.1Q packets for identification. This value can be any number between 1 and 4094.

Priority/Priority Remap Values

If the module uses a Kodiak ASIC, this field is labeled either **Priority** or **Priority Remap Values**. In single spanning tree mode, it is **Priority**. In multiple spanning tree mode, it is **Priority Remap Values**. See **Mode** below for more detailed information.

◆ Important Notes ◆

ESX-K and GSX-K Kodiak ASIC-based modules support 802.1p traffic prioritization. For chassis configurations that include only ESX-K, GSX-K and/or WSX series modules, 802.1p priority bits can be carried inbound on a tagged port (configured with multiple spanning tree 802.1Q) across the backplane. This priority information is used at the egress port to queue the packet, and is sent out in the packet whether the egress port is tagged or not.

The ESX-K and GSX-K modules can also remap incoming priority on an ingress port. If priority remapping has been configured, the new priority will be carried across the backplane. The priority information is used to queue the packet, and is sent out in the packet if the egress port is tagged.

Mode

This field allows you to choose either multiple or single spanning tree. This option only appears if the module uses 10/100 Ethernet ports. Once you select a type of spanning tree for a port, the port automatically retains the spanning tree selection for any other group it is added to.

Assigning an 802.1Q Group to a Port

For example, suppose that Port 3/1 is assigned to be in Group 2, and to use single spanning tree. If the port were to be assigned to another group, it would automatically set itself to use single spanning tree for that group as well.

When you set the **Mode** of the service, the **cas** screen changes to accommodate the selection and allows you to set the priority of the service. If you select single spanning tree, for example, the screen changes to the following display, as shown:

```
Slot 3 Port 5 Ethernet 802.1Q Service
1) Description           :
2) Group ID             :
3) Tag                  :
4) Priority              :
5) Mode                 : 4
```

If you select multiple spanning tree, the screen changes to the following display, as shown:

```
Slot 2 Port 1 Ethernet 802.1Q Service
1. Description (30 chars max) :
2. Group ID                   : 0
3. Tag                         : 0
4. Priority Remap Values      :
   40. 0 - 0
   41. 1 - 1
   42. 2 - 2
   43. 3 - 3
   44. 4 - 4
   45. 5 - 5
   46. 6 - 6
   47. 7 - 7
5. Mode                       : 3
```

The incoming priority level of the packet can be remapped to any value between **0** and **7**, with **7** being the highest priority. To set a value of **5** for an incoming priority value of **4**, for example, you would enter **44=5**.

For more information on single vs. multiple spanning tree, see *Single vs. Multiple Spanning Tree* on page 16-4.

Configuring 802.1Q on Gigabit Ethernet Ports

Use the **cas** command to assign 802.1Q groups to Gigabit ports. To use this command, follow the steps below.

1. Enter **cas** at the system prompt, as shown:

```
cas <slot>/<port>
```

where **<slot>** is the slot of the module, and **<port>** is the port number that is to be added to the group. For example, to add port 3 on slot 5, you would enter:

```
cas 5/3
```

2. If you have a Kodiak Gigabit module, the following prompt displays:

```
Slot 3 Port 5 Ethernet 802.1Q Service
1. Description (30 chars max)  :
2. Group ID                   : 0
3. Tag                         : 0
4. Priority Remap Values      :
   40. 0 - 0
   41. 1 - 1
   42. 2 - 2
   43. 3 - 3
   44. 4 - 4
   45. 5 - 5
   46. 6 - 6
   47. 7 - 7
```

You can modify the parameters by entering the line number, an equal sign, and the value for the parameter. For example, to change the **Group ID** to **5**, you would enter **2** (the line number for **Group ID**), an equal sign (=), and a **5** (the group number), as shown:

```
2=5
```

3. Remember to save your changes by typing **save** at the system prompt when you have finished with the configuration.

Most of the fields are the same as described in *Configuring 802.1Q on 10/100 Ethernet Ports* on page 16-8.

Modifying 802.1Q Groups

802.1Q groups for both 10/100 and Gigabit Ethernet ports can be modified using the **mas** command. The procedure is slightly different in each case. The screens for the **mas** command change, depending on whether you have a legacy Ethernet board or a Kodiak ASIC-based Ethernet board.

Modifying 802.1Q Groups for 10/100 Ports

To modify the configuration of an 802.1Q group for 10/100 ports, use the **mas** command as shown:

```
mas <slot>/<port> <instance>
```

where **<slot>** is the slot number of the module on the switch, **<port>** is the port number where the service was created, and **<instance>** is the identifier for the service on this port. For example, to modify 802.1Q service instance 1 on port 5 of slot 2, enter:

```
mas 2/5 1
```

If this is a legacy Ethernet module, the screen appears as shown:

Slot 2 Port 5 Ethernet 802.1Q Service

```
1) Tag           : 3
2) Priority       : 0
```

If this is a Kodiak ASIC-based module, the screen appears as shown:

Slot 2 Port 5 Ethernet 802.1Q Service

```
1. Description (30 chars max) :
2. Tag                         : 0
3. Priority Remap Values       :
   30. 0 - 0
   31. 1 - 1
   32. 2 - 2
   33. 3 - 3
   34. 4 - 4
   35. 5 - 5
   36. 6 - 6
   37. 7 - 7
```

To change a field setting, enter the line number, an equal sign, and the new value. For example, to change the **Priority** setting to **7**, you would enter a **3** (the line number for priority), an equal sign (=), and a **37**, as shown:

```
3=37
```

◆ Important Notes ◆

ESX-K and GSX-K Kodiak ASIC-based modules support 802.1p traffic prioritization. For chassis configurations that include only ESX-K, GSX-K and/or WSX series modules, 802.1p priority bits can be carried inbound on a tagged port (configured with multiple spanning tree 802.1Q) across the backplane. This priority information is used at the egress port to queue the packet, and is sent out in the packet whether the egress port is tagged or not.

The ESX-K and GSX-K modules can also remap incoming priority on an ingress port. If priority remapping has been configured, the new priority will be carried across the backplane. The priority information is used to queue the packet, and is sent out in the packet if the egress port is tagged.

Remember to save the changes to the service by entering **save** at the system prompt when finished.

To find the instance of a port service, use the **vas** command. See *Viewing 802.1Q Groups in a Port* on page 16-16 for more information.

Modifying 802.1Q Groups for Gigabit Ethernet Ports

To modify the configuration of an 802.1Q group for Gigabit ports, use the **mas** command as shown:

```
mas <slot>/<port> <instance>
```

where **<slot>** is the slot number of the module on the switch, **<port>** is the port number where the service was created, and **<instance>** is the identifier for the service on this port. For example, to modify 802.1Q service instance 1 on port 5 of slot 2, enter:

```
mas 2/5 1
```

If this is a legacy Ethernet module, the screen appears as shown:

Slot 2 Port 5 Ethernet 802.1Q Service

```
1) Tag           : 3
2) Priority      : 0
```

If this is a Kodiak ASIC-based module, the screen appears as shown:

Slot 2 Port 5 Ethernet 802.1Q Service

```
1. Description (30 chars max) :
2. Tag                       : 0
3. Priority Remap Values     :
   30. 0 - 0
   31. 1 - 1
   32. 2 - 2
   33. 3 - 3
   34. 4 - 4
   35. 5 - 5
   36. 6 - 6
   37. 7 - 7
```

To change a field setting, enter the line number, an equal sign, and the new value. For example, to change the **Priority** setting to **7**, you would enter a **3** (the line number for priority), an equal sign (=), and a **37**, as shown:

```
3=37
```

◆ Important Notes ◆

ESX-K and GSX-K Kodiak ASIC-based modules support 802.1p traffic prioritization. For chassis configurations that include only ESX-K, GSX-K and/or WSX series modules, 802.1p priority bits can be carried inbound on a tagged port (configured with multiple spanning tree 802.1Q) across the backplane. This priority information is used at the egress port to queue the packet, and is sent out in the packet whether the egress port is tagged or not.

The ESX-K and GSX-K modules can also remap incoming priority on an ingress port. If priority remapping has been configured, the new priority will be carried across the backplane. The priority information is used to queue the packet, and is sent out in the packet if the egress port is tagged.

Remember to save the changes to the service by entering **save** at the system prompt when finished.

To find the instance of a port service, use the **vas** command. See *Viewing 802.1Q Groups in a Port* on page 16-16 for more information.

◆ **Note** ◆

Tags (field number **1**) do not apply if proprietary tagging is used on this port.

Viewing 802.1Q Groups in a Port

To view which ports use which 802.1Q groups, enter the **vas** command at the system prompt, as shown:

```
vas <slot>/<port>
```

where **<slot>** is the slot number of the module on the switch and **<port>** is the port number where the service was created. For example, to view an 802.1Q service on port 5 of slot 2, enter:

```
vas 2/5
```

A screen similar to the following is displayed:

Slot/Port/Inst	Vport	Group	Tag	Priority or PriorityRemap	Tagging Mode	Description
2 5 1	33	2	2	4	Mult STree	

As a variation of this command, it is possible to enter **vas** without a slot or port number. This will display all services configured for the switch.

◆ Note ◆

The above screen is for Gigabit ports. The display is slightly different for 10/100 ports. See descriptions below for more details.

The following section describes the fields displayed using the **vas** command.

Slot. The slot number of the switch on which the service is located.

Port. The port number of the slot on which the service is located.

Instance. The service identifier for the 802.1Q service. This is assigned when the service is created.

Vport. The virtual port number that the service uses.

Group. The group identifier for the group attached to this service.

Tag. The tag information entered into tagged frames, as specified when creating the service.

Priority or PriorityRemap. The priority number assigned to packets from this service.

Tagging Mode. This field displays different information depending on whether the switch ports are 10/100 or Gigabit. If the ports are 10/100 or Kodiak-based Gigabit, this field shows either multiple or single spanning tree. For 802.1Q support over OmniChannel on Kodiak 10/100 Ethernet boards, this field will display as **Mult S Tree**.

Description. A textual description used to identify the service.

For more information on single vs. multiple spanning tree, see *Single vs. Multiple Spanning Tree* on page 16-4.

Viewing 802.1Q Statistics for 10/100 Ports

The **viqs** command provides a display of statistics for 802.1Q groups assigned to 10/100 ports. Enter the **viqs** command, as shown:

```
viqs <slot>/<port> <groupid>
```

where **<slot>** is the slot number of the module on the switch, **<port>** is the port number where the service was created, and **<groupid>** is the number of the group that the port belongs to. For example, to view an 802.1Q service for group 2 on port 5 of slot 2, enter:

```
viqs 2/5 2
```

A screen similar to the following displays:

Physical Port	Group Id (802.1Q)	Transmit Pkts	Received Pkts	Transmit Octets	Received Octets
2/5	2	29	0	41	0

Physical Port. The slot and port number for this port.

Group Id (802.1Q). The 802.1Q group to which this port was assigned.

Transmit/Received Pkts. The number of packets transmitted and received on this port.

Transmit/Received Octets. The number of bytes transmitted and received on this port.

Deleting 802.1Q Groups from a Port

802.1Q groups for both 10/100 and Gigabit Ethernet ports can be deleted using the **das** command. The procedure is slightly different in each case.

To delete an 802.1Q group from a 10/100 port using single spanning tree, use the **das** command, as shown:

```
das <slot>/<port> <instance> <groupid>
```

where **<slot>** is the slot number of the module on the switch, **<port>** is the port number where the service was created, **<instance>** is the identifier for the service on this port, and **<groupid>** is the number of the group that the port belongs to. For example, to delete an 802.1Q service for group 2, instance 1 on port 5 of slot 2, enter:

```
das 2/5 1 2
```

To delete 802.1Q groups from a Gigabit port or 10/100 ports using multiple spanning tree, enter the **das** command, as shown:

```
das <slot>/<port> <instance>
```

where **<slot>** is the slot number of the module on the switch, **<port>** is the port number where the service was created, and **<instance>** is the identifier for the service on this port. For example, to delete 802.1Q service instance 1 on port 5 of slot 2, enter:

```
das 2/5 1
```

In either case, a message will appear, confirming the delete operation:

```
802.1Q service deleted for Group ID 3 on 3/9 (slot/Port)
```

◆ Important Notes ◆

You must delete X802.1Q groups in the same order on both ends of the link. For example, if you delete groups 1, 2, 3, 4, and 5 on the local switch, you must delete the same five groups in the same order on the remote switch. *If groups are not deleted in this manner, X802.1Q packets will not be routed correctly.*


To delete 802.1Q support over OmniChannel, you must first delete the 802.1Q service before you delete the OmniChannel.

17 Configuring Bridging Parameters

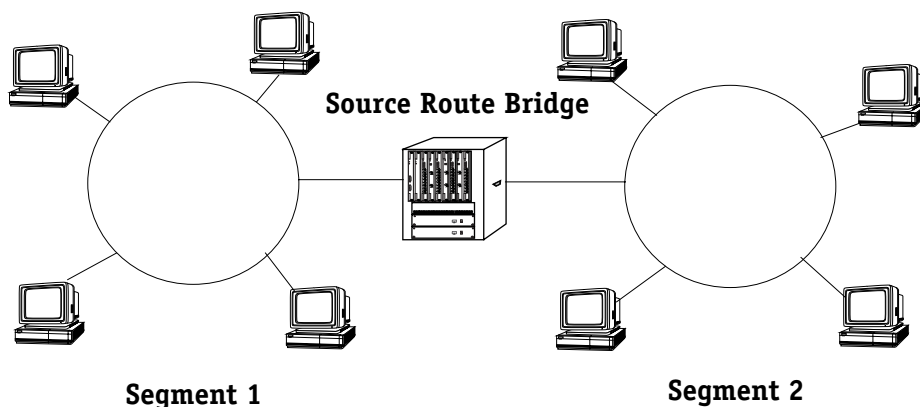
This chapter describes how to configure and maintain bridging parameters. Bridges are devices that interconnect LANs using one (or more) of the available standards such as transparent bridging, source route bridging, or source route to transparent bridging. Bridges primarily operate at Layer 2 of the OSI reference model, which controls data flow, transmission errors, physical addressing, and access to physical medium.

There are different types of bridging that are used to manage networks:


- **Transparent Bridging.** Used mainly in Ethernet environments, packets are usually forwarded without any changes being made to the packet. An ethernet environment is shown in the diagram below:



- **Source Route Bridging.** Used mainly in Token Ring environments, packets are transmitted along routes predetermined by explorer frames sent along multiple paths. Source Route Bridging modifies the routing information of the packet as it traverses the network. A token ring environment is shown in the diagram below:



- **Source Route to Transparent Bridging.** Used in mixed Ethernet and Token Ring environments, this protocol provides easy translation between transparent and source route bridging. A mixed ethernet and token ring environment is shown in the diagram below:



Spanning tree and fast spanning tree are also used to prevent physical loops in the network from creating excess traffic by blocking packet transmission on one or more ports.

This chapter describes the commands used for configuring various bridging commands for the above mentioned protocols, as well as diagnostic, spanning tree and fast spanning tree information.

◆ Important Notes ◆

In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. See Chapter 4, "The User Interface," for documentation on changing from CLI mode to UI mode.

Beginning with Release 4.4, FDDI is no longer supported. Beginning with Release 4.5, Token Ring and ATM are no longer supported.

Configuration Overview

When configuring bridging parameters, you will need to perform at least some of the following steps:

Step 1. Select a group

The bridging menu commands operate only on the currently selected group (or, for certain commands, VLAN). You can select a group with the **selgp** command. For information on using these commands, see *Selecting a Default Group* on page 17-7.

Step 2. Configure Bridging Parameters

There are several commands that allow you to configure and view basic bridging functions such as static MAC addresses, bridge forwarding tables, MAC information and statistics, and remote Trunking stations. Many of these commands are useful in diagnosing network problems, as they allow you to find specific MAC addresses and the port on which they were learned. For information on these commands, see *Bridging Commands* on page 17-8.

Step 3. Enable Spanning Tree (Optional)

Spanning tree is an algorithm that helps prevent broadcast storms by blocking ports in the network from transmitting data. If you plan to use spanning tree, you can use the spanning tree commands to configure and view IEEE and IBM Spanning Tree. For information on using spanning tree commands, see *Configuring Spanning Tree* on page 17-23.

Step 4. Enable Fast Spanning Tree (Optional)

Fast Spanning Tree is an algorithm that helps provide quick recovery from link, port and device failures on a network, by bringing blocked secondary links into forwarding mode as quickly as possible. You can the Fast Spanning Tree commands in the Bridge Management Menu to view and enable/disable Fast Spanning Tree parameters on a selected group or VLAN. For information on using Fast Spanning Tree commands, see *Configuring Fast Spanning Tree* on page 17-34.

Bridge Management Menu

To view the Bridge Management Menu, enter the **br** command at the system prompt. If you are in verbose mode, the following table appears outlining the commands available to you. If you are not in verbose mode, enter a **?** at the prompt to display the Bridge Management Menu.

Command	Bridge Management Menu
fls	Display Flood Limit of selected Group
flc	Configure Flood Limit on selected Group
sts	Display Spanning Tree parameters on selected Group
fstps	Display Fast Spanning Tree port parameters on selected VLAN
actfstps	Activate Fast Spanning Tree port parameters on selected VLAN
stc	Configure Spanning Tree parameters on selected Group
stps	Display Spanning Tree Port parameters on selected VLAN
stpc	Configure Spanning Tree Port parameters on selected VLAN
srs	Display Source Routing parameters on selected Group
src	Configure Source Routing parameters on selected Group
srsf	Enable or disable Source Routing SAP Filter Support
srtbcfg	View and configure Source Route to Transparent Bridging
srtbrif	View learned RIF from Source Route to Transparent Bridging Table
srtbclrrif	View and Clear learned RIF from Source Route to Transparent Bridging Table
fwf	Display Bridge Forward table on selected VLAN
fs	Display Bridge Static Address
fc	Configure Bridge Static Address
bps	Display Bridge Port Statistics on selected VLAN
macinfo	Locate learned Bridge MAC address in this chassis
macstat	Show statistics of Bridge MAC address
macclrstat	Clear statistics of Bridge MAC address
selgp	A Group can be selected for the bridge operations or to generate MIB reports
rts	Display remote Trunking Stations discovered
dbrmap	View the Domain Bridge Mapping table
+ / -	Select next / previous VLAN

Details on commands included in the Bridge Management Menu commands are given in the following sections:

Setting the Default Group. These commands allow you to choose which group you are modifying or viewing, and include the **selgp**, **+**, and **-** commands. For more information, see:

- *Selecting a Default Group* on page 17-7
- *Using the + or - to Change Groups* on page 17-7 for more information.

Bridging Commands. These commands allow you to view bridge forward tables, create and view static address tables, display bridge port statistics, view MAC address information, view remote trunking stations, and view the domain bridge mapping table. Commands in this section include **fw**, **fs**, **fc**, **bps**, **macinfo**, **macstat**, **macclrstat**, **rts**, and **dbrmap**. For more information, see:

- *Displaying Bridge Forwarding Table* on page 17-8
- *Configuring a Static Bridge Address* on page 17-10
- *Displaying Static Bridge Addresses* on page 17-13
- *Displaying Bridge Port Statistics* on page 17-14
- *Displaying Media Access Control (MAC) Information for a Specific MAC address* on page 17-16
- *Display Statistics of Bridge MAC Addresses* on page 17-17
- *Clear Statistics of Bridge MAC Addresses* on page 17-18
- *Display Remote Trunking Stations* on page 17-18
- *View the Domain Bridge Mapping Table* on page 17-19

Setting Flood Limits. These commands allow you to configure and view flood limits for a specific group using the **flc** and **fls** commands. For more information, see:

- *Setting Flood Limits for a Group* on page 17-21
- *Displaying Group Flood Limits* on page 17-22

Configuring Spanning Tree. These commands allow you to configure and view IEEE and IBM Spanning Tree for a specific group, and include the **stc**, **sts**, **stpc** and **stps** commands. (The **stc** and **sts** commands can also be used to configure and view Fast Spanning Tree for a selected VLAN.) For more information, see:

- *Configuring Spanning Tree Parameters* on page 17-25
- *Display Spanning Tree Bridge Parameters* on page 17-28
- *Configuring Spanning Tree Port Parameters* on page 17-30
- *Displaying Spanning Tree Port Parameters* on page 17-32

Configuring Fast Spanning Tree. These commands allow you to configure and view Fast Spanning Tree for a specific group or VLAN, and include the **actfstps** and **fstps** commands. Information is also included on configuring the Truncating Tree Timing and Speedy Tree Protocol features. For more information, see:

- *Configuring Truncating Tree Timing & Speedy Tree Protocol* on page 17-35
- *Displaying Fast Spanning Tree Port Parameters* on page 17-36
- *Enabling Fast Spanning Tree Port Parameters* on page 17-38
- *Disabling Fast Spanning Tree Port Parameters* on page 17-39

Selecting a Default Group

Most commands in the Bridge Management Menu allow you to specify a group when entering the command at the system prompt. If you do not specify a group when entering a command, the bridge operations are performed on the currently selected group.

◆ Note ◆

You can view the current groups in the switch by entering **gp** at any prompt.

To select a group, enter the **selgp** command as follows:

```
selgp <group number>
```

where **<group number>** is the number of the group you wish to modify or view. For example, to select Group 2 you would enter **selgp** and the number **2** as shown:

```
selgp 2
```

A message confirming the selection of the new group ID followed by the group description.

```
Group number: 2 is now selected (New GROUP (#1)).
```

Using the + or - to Change Groups

At any time from the system prompt, you can select a different group by typing a plus (+) to move up one group, or a minus (-) to move back one group. For example, if you are currently working on Group 4 and wish to change to Group 3, you would enter a - at the system prompt. The following message displays to confirm the change:

```
Currently GROUP 3 is selected (New GROUP (#3))
```

Bridging Commands

The Bridge Management menu provides several commands that are useful in pinpointing problems in the network. The commands allow you to lookup specific MAC addresses and where they were learned, create and view static bridge addresses, view information on remote trunking stations, view MAC address statistics for a group or a port, or look up information on domain mappings. Many times a network problem can be tracked down by viewing MAC address information, finding out where it came from, and where it forwards data.

The following sections detail the specific bridging commands that perform these functions.

Displaying Bridge Forwarding Table

You can display the MAC addresses and their forwarding and filtering information for a given group. The information in the table is used by the transparent bridging function in determining how to propagate a received frame.

To display the information for a group in the switch follow these steps:

1. Enter the **fw**t command at the system prompt as follows:

```
fw t <group number>
```

where **<group number>** is the number of the group for which you want to view MAC addresses. For example, to view MAC addresses for group 2, you would enter:

```
fw t 2
```

As a variation of this command, you can enter the **fw**t command without a group ID. This will display MAC addresses for the currently selected group in this switch. For information on selecting a group, see *Selecting a Default Group* on page 17-7.

2. Once you have entered the group number you will be prompted for a slot and port, as shown:

```
Enter Slot/Interface (return for all ports):
```

3. Enter the slot and interface (port) number and press **<return>**. For example, to view MAC addresses for port 2 on slot 3, enter 3/2 as shown:

```
Enter Slot/Interface (return for all ports): 3/2
```

The following screen appears listing the MAC addresses on this port:

Total number of MAC addresses learned for VLAN 2: 8										
Sl/If/Srv/In	MAC Address	Non-Canonical MAC Address	T	Group ID	CAM Indx	S	Last Seen	Exp Timer	ATM VCI	
3/1/ Brg/ 1	0020DA:A373B0	00045B:C5CE0D	E	2	305A	T	11	300		
3/1/ Brg/ 1	0020DA:8656F0	00045B:616A0F	E	2	3060	T	11	300		
3/1/ Brg/ 1	00045B:ED48C0	00045B:2251A1	E	2	3080	T	29	300		
3/1/ Brg/ 1	000077:8DDBB9	00045B:65EE22	E	2	3010	T	29	300		
3/1/ Brg/ 1	000039:F5520C	0009E4:3ED444	E	2	300E	T	35	300		
3/1/ Brg/ 1	009027:17F7EB	00045B:2D43EF	E	2	3018	T	59	300		
3/1/ Brg/ 1	0020DA:0C41E5	00045B:ED48C0	E	2	3078	T	26	300		
3/1/ Brg/ 1	0020DA:9645A1	0000EE:B1DB9B	E	2	304E	T	18	300		

Field Descriptions

The following section explains the fields displayed with the **fw**t command.

Sl/In/Srvc/In. The slot number (**Sl**), interface (port) number (**In**), type of service (**Srvc**), and service instance (**In**). For example, a bridge service on port 1 of slot 3 would be:

3/1/Brg/1

Services provide connection options for switches in a LAN, between LANs, or in a WAN. Other possible services include trunking, routing, and LANE. It is possible to have more than one instance of a service if there are more than one connections on a single port.

MAC Address. The learned MAC address for this port.

Non-Canonical MAC address. The non-canonical version of the learned MAC address. The non-canonical MAC address is different from a canonical MAC address in that the order in which the address information is sent is different. Ethernet uses canonical address, while other media (e.g., token ring, FDDI) use non-canonical.

T. The protocol type of this MAC address. There are two possibilities:

E	Ethernet
F	FDDI
T	Token Ring

Group ID. The associated group ID for this learned MAC address.

CAM Indx. The index number to the Content-Addressable Memory (CAM), where the MAC addresses are stored, in hexadecimal form.

S. The source of the MAC address (how it was learned). There are two possibilities:

T	Transparent Bridge
S	Source Route Frame.

Last Seen. The time in seconds since this MAC address was last seen on this port.

Exp. Timer. There are three possibilities for this column:

Value	The configured ageing timer, in seconds, for this MAC address is shown. Once this time period is exceeded, the MAC address is removed from the CAM.
STATIC	This MAC address was manually assigned to this group and will not age out.
OPSWT	This MAC address was learned on an optimized switch port and will not age out.

ATM VCI. The ATM Virtual Channel Identifier (VCI) for this MAC address entry. The VCI is shown for any media that uses Virtual Circuits (ATM, LANE).

Configuring a Static Bridge Address

You can configure static bridge address information by entering the **fc** command. A static bridge address is a fixed MAC address bridge that does not change or age out.

To configure a static MAC address follow these steps:

1. Enter the **fc** command as follows:

```
fc <groupNumber>
```

where **<groupNumber>** is the number of the group for which you want to create a static bridge MAC address. For example, to set up a static bridge address for Group 2, you would enter the following:

```
fc 2
```

As a variation of this command, you can enter the **fc** command at the system prompt with no group number. This will allow you to set up a static bridge address on the currently selected group. For information on selecting a group, see *Selecting a Default Group* on page 17-7.

The system displays the following:

```

      Bridge Static Address for Group 2 (New GROUP (#2))
-----
Index   MAC Address   Slot/Intf/Service/Inst   Static Status
              (A)              (B)
-----
   1    21A33E:00B001   3/ 1/  Brg/1           permanent
  
```

The entries can be modified by specifying the index and column.

For Static Status, use 2 to delete, 3 for Permanent,

4 for Delete on Reset, 5 for Delete on Timeout

To add an entry: Use command 'add MAC addr, receiving port, static status'.

Receiving port and Status must be provided.

Port could either be slot/intf or virtual port begin with v.

For non-canonical MAC format add 'nc' before MAC.

ie: add 123456:7890AB, 2/3, 3 or add nc001122:334455, v99, 3

NOTE: add command will be executed immediately.

save|cancel|next only applies to existing entry.

add|save|cancel|next :

2. To add an entry, use the format as described in the above screen:

```
add [MAC Addr], [Slot/Intf], [Static Status]
```

For example, to add a permanent non-canonical MAC address of 123456:123456 to port 2 of slot 3, you would enter the following:

```
add nc123456:123456, 3/2, 3
```

When you complete the operation by pressing **<return>**, an entry with MAC address 123456:123456, on slot 2, port 3, with a **Static Status** of **Permanent** is created.

3. Type **save** at the **fc** command prompt to save the entry. If you do not save the entry before exiting the **fc** command, the static bridge address is not created.

◆ Note ◆

The newly created static bridge address will not show up in the **fc** command table until you have exited the **fc** command by typing **cancel** at the command prompt.

Field Descriptions

The following section describes the fields in the **fc** command table.

Index. A number assigned to the row to identify a previously created static bridge address, when modifying the address.

MAC address. The canonical MAC address for this static bridge.

Slot/Intf/Service/Inst. The slot number, interface (port) number, type of service, and service instance. For example, a bridge service on port 1 of slot 3 would be:

3/1/Brg/1

Static Status. The status of the static MAC address as determined when created. The **Status** will be one of the following:

Invalid	This entry was deleted within the current session.
Permanent	This entry is in use and will remain so until it is deleted from the table. See <i>Deleting a Static Bridge Address</i> on page 17-12 for specific information.
deleteOnReset	This entry is in use and will remain so until the bridge is reset.
deleteOnTimeOut	This entry is currently in use and will remain so until it is aged out.

Modifying a Static Bridge Address

Once you have created a static bridge address, you can modify its interface assignment or its status. To modify a static bridge address:

1. Enter the **fc** command as documented above. The Bridge Static Address table will display as shown:

Bridge Static Address for Group 2 (Default GROUP (#2))

Index	MAC Address	Slot/Intf/Service/Inst (A)	Static Status (B)
1	21A33E:00B001	3/ 1/ Brg/1	permanent
2	001122:223344	3/ 2/ Brg/1	deleteOnReset

The entries can be modified by specifying the index and column.

For Static Status, use 2 to delete, 3 for Permanent,

4 for Delete on Reset, 5 for Delete on Timeout

To add an entry: Use command 'add MAC addr, receiving port, static status'.

Receiving port and Status must be provided.

Port could either be slot/intf or virtual port begin with v.

For non-canonical MAC format add 'nc' before MAC.

ie: add 123456:7890AB, 2/3, 3 or add nc001122:334455, v99, 3

NOTE: add command will be executed immediately.

save|cancel|next only applies to existing entry.

add|save|cancel|next :

- To modify an entry, use the index number for the specific static bridge address (listed in the leftmost column), the column letter for the column you want to change, an equal sign, and a new value. For example, to change the **Static Status** of the first address's in the table from **permanent** to **deleteOnReset**, you would enter a **1** (the static bridge address **Index** number), a **b** (the column letter for **Static Status**), an equal sign (=), and the number **4** (the value for **deleteOnReset**), as shown:

1b=4

- Press **<return>** to complete the operation.
- Type **save** at the **fc** command prompt to save the changes.

Deleting a Static Bridge Address

Deleting a previously created static bridge address is much the same process as modifying a Static Bridge Address. To delete a Static Bridge Address, follow these steps:

- Enter the **fc** command as documented above. The Bridge Static Address table will display as shown:

Bridge Static Address for Group 2 (Default GROUP (#2))

Index	MAC Address	Slot/Intf/Service/Inst (A)	Static Status (B)
1	21A33E:00B001	3/ 1/ Brg/1	permanent
2	001122:223344	3/ 2/ Brg/1	deleteOnReset

The entries can be modified by specifying the index and column.

For Static Status, use 2 to delete, 3 for Permanent,

4 for Delete on Reset, 5 for Delete on Timeout

To add an entry: Use command 'add MAC addr, receiving port, static status'.

Receiving port and Status must be provided.

Port could either be slot/intf or virtual port begin with v.

For non-canonical MAC format add 'nc' before MAC.

ie: add 123456:7890AB, 2/3, 3 or add nc001122:334455, v99, 3

NOTE: add command will be executed immediately.

save|cancel|next only applies to existing entry.

add|save|cancel|next :

- To delete an entry, use the index number for the specific static bridge address, the column letter **b** (the column letter for **Static Status**), an equal sign (=), and a **2** (the value for **Delete**).

For example, to delete the first address in the table, you would enter a **1** (the static bridge address **Index** number), a **b** (the column letter for **Static Status**), an equal sign (=), and the number **2** (the value for **Delete**), as shown:

1b=2

- Press **<return>** to complete the operation.
- Type **save** at the **fc** command prompt to save the changes. The **Static Status** will change to **Invalid**. Once you exit the **fc** command, the Static Bridge Address is removed from the table.

Displaying Static Bridge Addresses

You can view static bridge address information by entering the **fs** command. To display the information, enter the **fs** command as follows:

```
fs <group number>
```

where **<group number>** is the number of the group for which you want to view static bridge MAC addresses. For example, to view MAC addresses for Group 1, you would enter the following:

```
fs 1
```

This command will display a table similar to the following:

Bridge Static Address Summary for Group 1 (Default GROUP (#1))

MAC Address	Slot/Intf/Service/Inst	Static Status
002A3113:0012EA	3/ 1/ Brg/ 1	permanent

As a variation of this command, you can enter the **fs** command at the system prompt with no group number. This will allow you to view the static bridge addresses on the currently selected group. For information on selecting a group, see *Selecting a Default Group* on page 17-7.

The descriptions for the variables in the table displayed with the **fs** command are the same as those in the table displayed with the **fc** command. For details on these variables, see *Configuring a Static Bridge Address* on page 17-10.

Displaying Bridge Port Statistics

You can display statistics on bridge ports with the **bps** command. To view bridge port statistics enter the **bps** command as follows:

```
bps <group number>
```

where **<group number>** is the number of the group for which you want to view bridge port statistics. For example, to view statistics for Group 1, you would enter the following:

```
bps 1
```

This command will display a table similar to the following:

Frames discarded due to full Forwarding Database:0

Port Statistics for Group 1

Slot/Intf Service/Inst	Frames In	Frames Out	In Frames Discards	MTU Exceeded Discards	Delay Exceeded Discards	Flood Limit Discards
=====	=====	=====	=====	=====	=====	=====
2/ 1/ Brg/ 1	0	0	0	0	0	0
2/ 2/ Brg/ 1	0	0	0	0	0	0
3/ 1/ Brg/ 1	3354	85	0	0	0	0
3/ 2/ Brg/ 1	0	0	0	0	0	0
3/ 3/ Brg/ 1	0	0	0	0	0	0
3/ 4/ Brg/ 1	0	0	0	0	0	0
3/ 5/ Brg/ 1	0	0	0	0	0	0
3/ 6/ Brg/ 1	0	0	0	0	0	0
3/ 7/ Brg/ 1	0	0	0	0	0	0
3/ 8/ Brg/ 1	0	0	0	0	0	0
/VLAN/Bridge %						

As a variation on this command, you can enter **bps** at the prompt without a group number. This will display the port statistics for the currently selected group. For information on selecting a group, see *Selecting a Default Group* on page 17-7.

Field descriptions

The following section describes the fields displayed in the above table.

Frames discarded to full Forwarding Database. The number of frames that were not transmitted because the forwarding database is full. The forwarding database holds all known MAC address for this bridge and is used to learn the next hop MAC address for the packet(s) in question.

Slot/Intf/Service/Inst. The slot number (**SI**), interface (port) number (**Intf**), type of service (**Service**), and service instance (**Inst**). For example, a bridge service on port 1 of slot 3 would be:

3/1/Brg/1

Services provide connection options for switches in a LAN, between LANs, or in a WAN. Other possible services include trunking, routing, and LANE. It is possible to have more than one instance of a service if there are more than one connections on a single port.

Frames In. The number of frames received on the associated port.

Frames Out. The number of frames sent on the associated port.

In Frames Discards. The number of received frames discarded due to error.

MTU Exceeded Discards. The number of frames that were discarded because they exceeded the Maximum Transmission Unit (MTU) size. The MTU is set to the default of the media type (Ethernet, Token Ring, etc.) and is not configurable.

Delay Exceeded Discards. Frames that were delayed, usually due to collisions, but that were ultimately transmitted.

Flood Limit Discards. The number of frames that were discarded because they exceeded the flood limit set for the port or the group in which this port is a member. This flood limit is set with the **flc** command for groups or the **modvp** command for ports. For more information on setting flood limits, see *Setting Flood Limits* on page 17-21 for the **flc** command. For details on using the **modvp** command, see Chapter 19, “Managing Groups and Ports.”

Displaying Media Access Control (MAC) Information for a Specific MAC address

Media Access Control (MAC) information for the switch can be examined by using the **macinfo** command. You can view specific MAC address information, or choose a slot and view all MAC addresses associated with the selected slot.

To view MAC information for a specific address:

1. Enter **macinfo** at the system prompt and press **<return>**.
2. You will be prompted with the following message:

Enter MAC address ([XXYYZZ:AABBCC] or return for none):

Enter the MAC address you are interested in viewing, and press **<return>**.

3. You will be prompted with the following message:

Is this MAC in Canonical or Non-Canonical form (C or N) [C]:

Enter **c** for Canonical or **n** for Non-Canonical (the default is at the end of the prompt in brackets) and press **<return>**. A table similar to the following is shown:

Slot/Intf/Srvc/Inst	Group ID	CAM Index	Set by	MAC Type	Last Seen	Exp Timer	ATM VCI	Protocol
3/ 1/ Brg/ 1	1	1	0346	TB	ETH	11	15	

Field Descriptions

The following section explains the fields displayed using the **macinfo** command that are not previously explained in other sections.

Set by. This field lists what type of bridging was used to learn this MAC address. There are two possibilities:

- TB** This MAC address was learned using Transparent Bridging.
- SR** This MAC address was learned using Source Routing.

MAC Type. The media type of this MAC address. There are two possibilities:

- E** Ethernet
- F** FDDI
- T** Token Ring

Protocol. If Group Mobility is enabled, this field will list the type of packet encapsulation used when this MAC address was learned. For additional information on Group Mobility, see Chapter 19, “Managing Groups and Ports.”

Displaying Media Access Control (MAC) Information for all MAC addresses

Media Access Control (MAC) information for the switch can be examined by using the **macinfo** command. You can view all MAC addresses associated with the selected slot.

To view MAC information for all addresses:

1. Enter **macinfo** at the system prompt and press **<return>**. You will be prompted with the following message:

Enter MAC address ([XXYYZZ:AABBCC] or return for none):

2. Press **<return>**. You will be prompted with the following message:

Enter Slot Number (1-3):

Enter the slot number for the slot for which you are interested in viewing MAC addresses. The possible options are displayed on the right in parenthesis. A screen similar to the following is shown:

```
Total number of MAC addresses learned for VLAN 2: 8
```

Sl/If/Srvcln	MAC Address	Non-Canonical MAC Address	Group T	ID	CAM Indx	S	Last Seen	Exp Timer
3/1/ Brg/ 1	0020DA:A373B0	00045B:C5CE0D	E	2	305A	T	11	300
3/1/ Brg/ 1	0020DA:8656F0	00045B:616A0F	E	2	3060	T	11	300
3/1/ Brg/ 1	00045B:ED48C0	00045B:2251A1	E	2	3080	T	29	300
3/1/ Brg/ 1	000077:8DDBB9	00045B:65EE22	E	2	3010	T	29	300
3/1/ Brg/ 1	000039:F5520C	0009E4:3ED444	E	2	300E	T	35	300
3/1/ Brg/ 1	009027:17F7EB	00045B:2D43EF	E	2	3018	T	59	300
3/1/ Brg/ 1	0020DA:0C41E5	00045B:ED48C0	E	2	3078	T	26	300
3/1/ Brg/ 1	0020DA:9645A1	0000EE:B1DB9B	E	2	304E	T	18	300

Descriptions of the fields displayed with the **macinfo** command are identical to those displayed using the **fw** command. See *Displaying Bridge Forwarding Table* on page 17-8 for more information.

Display Statistics of Bridge MAC Addresses

The **macstat** command allows you to view a list of MAC address statistics for this switch on a slot-by-slot basis. To view MAC address statistics, enter the **macstat** command at the system prompt as shown:

macstat <slot>

where **<slot>** is the slot number on the switch for which you want to see statistics. For example, to view statistics for MAC addresses on slot 3, you would enter:

macstat 3

A table similar to the following is shown:

Slot	Discarded	Aged	Learned	in CAM
3	0	4	7	37

As a variation of this command, you can enter **macstat** at the prompt with no slot specified. This will display the statistics for all slots in the switch.

Field Descriptions

The following section describes the fields displayed using the **macstat** command.

Slot. The slot number of the switch to which the MAC address statistics apply.

Discarded. The number of MAC addresses that have been discarded on this slot due to the CAM being full.

Aged. The number of MAC addresses that have exceeded the age limit and been removed from the CAM by this slot.

Learned. The number of MAC address that have been learned on this slot.

in CAM. The total number of MAC addresses currently stored in the Content-Addressable Memory (CAM) of this module.

Clear Statistics of Bridge MAC Addresses

MAC address statistics for a slot can be cleared using the **macclrstat** command. To clear statistics, enter the **macclrstat** command at the system prompt as shown:

```
macclrstat <slot>
```

where **<slot>** is the slot number of the switch for which you want to clear MAC address statistics. For example, to clear statistics for slot 3, you would enter:

```
macclrstat 3
```

Once you have enter the command, a message appears to confirm the action.

As a variation of this command, you can enter **macclrstat** without specifying a slot. This will clear MAC statistics for all slots.

Display Remote Trunking Stations

The **rts** command displays a table of the remote trunking stations learned by this switch. A remote trunking station is a switch that has set up a trunking service to convey media through a network. Trunking services allow for media to be masked so that it appears to be a different type (for example, trunking ethernet over an ATM backbone). To display the remote trunking stations this switch has learned, follow these steps:

1. Enter the **rts** command as shown

```
rts <groupNumber>
```

where **<groupNumber>** is the number of the group on the local switch for which you want to view known trunking stations. For example, to view remote trunking stations for Group 1, you would enter the following:

```
rts 1
```

As a variation of this command, you can enter the **rts** command without a group number. This will show all the remote trunking stations for all groups in this switch.

- The following prompt is shown:

Enter service's Slot/Station (return for all services):

Enter the slot and station (port) number for the local switch for which you wish to view remote trunking services. For example, to list the trunking station at port 1 of slot 3, you would enter:

3/1

If you do not enter a specific slot and station, the system automatically sends information on all services for the remote trunking stations associated with this group.

- Once you have entered a slot and station, a table similar to the following is shown:

Remote Trunking Stations		
Slot/Station	Group ID	Remote MAC
=====	=====	=====
3/ 1	1	0020DA:022061
3/ 1	1	0020DA:05EAD1

Field Descriptions

The following sections describes the fields displayed by the **rts** command.

Slot/Station. The slot number and station (port) number associated with the remote trunking station.

Group ID. The group number of the switch that is associated with this remote trunking station.

Remote MAC. The Media Access Control address of the remote trunking service.

View the Domain Bridge Mapping Table

The **dbrmap** command allows you to display the mapping between a packet's destination MAC address and the remote Domain Bridge behind which it originated. To view this table:

- Enter the **dbrmap** command as shown:

dbrmap <groupNumber>

where **<groupNumber>** is the number of the group for which you want to see domain mappings of MAC addresses. For example, to view the mapping table for group 2, you would enter:

dbrmap 2

As a variation of this command, you can enter the **dbrmap** command without specifying a group. This will display mapping information for all groups on this switch.

- A prompt asking for a canonical MAC address is displayed, as shown:

Enter canonical MAC address ([XYZZ:AABBCC] or return to display everything):

Enter the MAC address you want to see the Domain Mapping for, or press **<return>** without entering a MAC address to see the mappings for all MAC addresses associated with this group.

3. A screen similar to the following is shown:

DOMAIN BRIDGE MAPPING				
Group 2				
Destination MAC	Group ID	Age	Slot / Intf	Domain MAC
00:20:da:7d:ef:44	2	14	8 / 1	00:20:da:6c:fb:85
00:20:da:7d:ef:45	2	120	8 / 1	00:20:da:6c:fb:85
00:20:da:7d:ef:46	2	220	8 / 1	00:20:da:6c:fb:86

Field Descriptions

The fields displayed by the **dbrmap** command are described below.

Destination MAC. The destination MAC address learned from a domain bridge port.

Group ID. The destination MAC's group number.

Age. The time, in seconds, since the destination MAC address was last seen.

Slot/Intf. The slot and interface number on this switch where the destination MAC address was learned.

Domain MAC. The remote domain MAC address behind which this destination MAC address was learned.

Setting Flood Limits

The flood limit is the number of bytes per second of flooded data that may be transmitted on a port on a group. This limit is a mechanism for controlling broadcast storms on the network.

The default flood limit for a port, regardless of the media type, is 192,000 bytes per second. You can change this default by configuring the flood limit on a per port or a per Group basis.

The **modvp** command (described in Chapter 19, “Managing Groups and Ports”) allows you to set the flood limit on a per port basis. The **flc** command (described in the following section) allows you to set the flood limit on a per Group basis. Configuring the flood limit for a Group is particularly useful when you need to disable flood limits for all ports in a single Group.

Setting Flood Limits for a Group

The **flc** command allows you to set flood limits for a Group. To set the flood limit for a Group

1. Enter the following at the system prompt follow these steps:

```
flc <groupNumber>
```

where **<groupNumber>** is the number of the group for which you are setting the flood limit. For example, to set the flood limit on Group 2 you would specify:

```
flc 2
```

As a variation of this command, you can enter the **dbmap** command without specifying a group. This will display mapping information for all groups on this switch.

The following prompt displays:

```
Enter flood limit override value (bytes/second) for Group 2 (192000):
```

2. Enter the flood limit for this Group and press **<Return>**.

◆ **Note** ◆

A value of negative one (-1) disables flood limits for the Group.

When new ports are added to a group, they will use the flood limit specified through **flc**. If a value has not been specified through **flc** for this Group, then the default port value (192000) is used.

◆ **Note** ◆

Flood limits set through **modvp** (set on a per-port basis) override the flood limit set through **flc**.

Displaying Group Flood Limits

The **fls** command allows you to view the current flood limits set for groups. The limits are set using the **flc** command. To display flood limits for all Groups, enter

```
fls <groupNumber>
```

where **<groupNumber>** is the number of the group for which you are viewing the flood limit. For example, to set the flood limit on Group 2 you would specify:

```
flc 2
```

A message similar to following is shown:

```
Flood Limit Override for Group 2(Group Name 1) is 190000 bytes per second.
```

A value will only be displayed for a Group on which **flc** has been used to set a flood limit.

As a variation of this command, you can enter **fls** at the system prompt without specifying a group number. This will return flood limit information for each group configured for this switch.

Configuring Spanning Tree

Spanning Tree is an algorithm developed to help prevent the occurrence of broadcast storms in a network. A packet can be broadcast multiple times in a network if the network is physically configured with loops.

If packets are broadcast to all ports (or flooded) in an attempt to deliver the data, networks with physical loops will rebroadcast packets repeatedly and cause a network to become severely congested. This congestion will adversely affect network performance.

Spanning Tree prevents broadcast storms by establishing a loop-free topology throughout the network. This is done by blocking ports in the physical topology that could result in flooded traffic being looped.

Both the IEEE and IBM versions of spanning tree are supported in the OmniSwitch/Router. The IBM Spanning Tree protocol is only supported by IBM Token Ring environments that make use of functional addresses for the transmission of Bridge Protocol Data Units (BPDUs). The following are the primary differences between the IEEE 802.1d and IBM Spanning Tree algorithms:

- The Hello BPDUs in IBM Spanning Tree are sent to the bridge functional address, X'C00000000100'. In the IEEE 802.1d Spanning Tree, they are sent to the Group address X'800143000000'.
- The Port ID in IBM Spanning Tree consists of a ring identifier and a bridge number. In 802.1d, it consists of a port priority and port number.
- IBM Spanning Tree has no learning process. Therefore, a port can be in one of three states—blocking, listening, or forwarding.
- IBM Spanning Tree does not support the Topology Change Notification (TCN) protocol.
- When you enable IBM Spanning Tree, the switch automatically sets defaults for the maximum age, forward delay, and hello time. In the interests of screen consistency, it is possible to change these defaults with the UI. In IBM Spanning Tree specification, these values are fixed, and should remain at the set defaults.
- When you enable IBM Spanning Tree, some additional defaults are set:
 - All virtual ports attached to the group with a physical port speed of 4 or 16 Mb are set to use Functional Addresses rather than Group Addresses.
 - All virtual ports attached to the group with a physical port speed that is not 4 or 16 Mb are set to manual forwarding.
 - As other virtual ports are attached to the group, the above two rules are applied.

Virtual ports in a manual forwarding state do not participate in either the IEEE or IBM versions of spanning tree. Any IEEE Spanning Tree frame received on a port in a manual forwarding state is forwarded to all other virtual ports in the same group also in a manual forwarding state. This is done to prevent loops from occurring in the network topology that could arise from applying the second default condition automatically.

- IBM SRT bridges send an IEEE-style STE RIF over Token Ring networks. The Omni Switch/Router does not support this frame, and any frame of this type received by the switch is discarded.
- The OmniSwitch/Router does not support using the same Functional Address (FA) for both data and spanning tree frames. The FA for IBM Spanning Tree is programmed into the MPX CAM, and all data frames with this FA are claimed by the MPX. Therefore, any data with the same FA as the IBM Spanning Tree FA will not be able to pass through the switch. There are two workarounds for this situation:
 - If you are *not* using IBM Spanning Tree and you want to prevent the specific FA from being programmed into the MPX CAM, then enter the command *faBpGrpDisable* into the mpx.cmd file, before the *cmInIt* command, with a value of 1.
 - If you are using IBM Spanning Tree and need the FA (0300 0000 0800), and you are using all Alcatel equipment (or other third party switch that allows you to change the IBM Spanning Tree FA), you can enter the command *faBpGrpOverride* into the mpx.cmd file with a new value for the lower 32-bit part of the address (0000 0800).

◆ **Note** ◆

If you change a group to IBM Spanning Tree, all non-Token Ring ports are put into manual forwarding state. Messages are displayed indicating these port state changes; in addition, SNMP traps are sent to indicate these changes. (Manual forwarding state is where the port is put into forwarding state and the Spanning Tree algorithm is disabled.) Token Ring ports will be set to use functional addresses.

The following sections provide specific information on using the spanning tree commands.

Configuring Spanning Tree Parameters

The **stc** command allows you to configure parameters for the spanning tree, and enable or disable the Fast Spanning Tree feature for a VLAN. To configure these parameters:

1. Enter the **stc** command as follows:

```
stc <groupNumber>
```

where **<groupNumber>** is the number of the group in the switch for which you are configuring spanning tree. For example, to configure spanning tree for Group 2, you would enter:

```
stc 2
```

2. The system shows you the current values and allows you to change them through a series of prompts, the first of which is shown below:

```
Spanning Tree Parameters for Group 2 (New GROUP (#2))
```

```
Spanning Tree is OFF for this Group, set to ON ? (y/n) :
```

Enter **y** to enable spanning tree or **n** to leave it disabled and press **<return>**. This field allows you to toggle spanning tree On or OFF by typing the appropriate response. Answering Yes (**y**) selects the option opposite the currently selected option.

◆ Important Note ◆

Remember to read the prompt carefully before responding. If spanning tree has already been activated for this group, this prompt will ask you if you would like to turn it *off*.

3. The following prompt is displayed asking whether you would like to use IEEE or IBM Spanning Tree:

```
IEEE spanning tree for this Group, set to IBM ? (y/n) :
```

Enter **n** to use IEEE Spanning Tree, or **y** to use IBM Spanning Tree, and press **<return>**. Select either the IEEE 802.1d Spanning Tree or IBM Spanning Tree. Answering Yes (**y**) changes the spanning tree type to the type not currently in use for this Group. The system automatically sets defaults for later **stc** prompts, such as **Bridge Hello Time** and **Bridge Max Age**, based on the spanning tree type you select here.

◆ Important Note ◆

Remember to read the prompt carefully before responding. If IEEE Spanning Tree is what you would like to use, the correct response to this prompt is *no*. A yes response changes it to IBM Spanning Tree.

- The following prompt is displayed asking whether you would like to use the Fast Spanning Tree feature:

Fast Spanning Tree is OFF for this Group, set to ON? (y/n) :

Enter **n** to leave Fast Spanning Tree disabled, or **y** to enable Fast Spanning Tree, and press **<return>**. Answering Yes (**y**) changes the setting of Fast Spanning Tree to the status not currently in use for this Group.

◆ **Important Note** ◆

Read the prompt carefully before responding. If Fast Spanning Tree is what you would like to use, the correct response to this prompt is *yes*. A *no* response leaves the Fast Spanning Tree feature disabled.

- The following prompt is shown allowing you to set the priority:

New Priority (0..65535) (current value is 32768[0x8000]) :

Enter the **Priority** value as a number between 0 and 65535, or press **<return>** to accept the default listed in parenthesis. A value of 0 is the highest priority. Bridge priority is utilized by the spanning tree algorithm to decide which bridge will be the root bridge. You can set the bridge priority by entering a decimal number from 0 to 65,535. 0 is the highest priority.

◆ **Note** ◆

To make sure that the proper negotiation occurs for the switch to become the Spanning Tree root bridge, always set the priority for the switch accordingly. Do not rely on MAC addresses to determine which switch becomes the root bridge.

- The following prompt is displayed allowing you to set the Bridge Hello Time:

New Bridge Hello Time (1..10 secs) (current value is 2) :

Enter the **Bridge Hello Time** as a number between 1 and 10, or press **<return>** to accept the default listed in parenthesis. The amount of time between the transmission of Configuration Bridge Protocol Data Units (BPDUs) on any designated port. Enter a value between 1 and 10 seconds. Shortening the time will make the protocol more robust, while lengthening the time lowers the overhead of the algorithm as the interval between transmission of configuration messages is larger.

- The following prompt is displayed allowing you to set the Bridge Maximum Age:

New Bridge Max Age (6..40 secs) (current value is 6) :

Enter the **Bridge Max Age Time** as a number between 6 and 40, or press **<return>** to accept the default listed in parenthesis. The maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded, in seconds. Enter a value between 6 and 40 seconds. A smaller value causes Spanning Tree to reconfigure more often.

8. The following prompt is displayed allowing you to set the Bridge Forward Delay:

New Bridge Forward Delay (4..30 secs) (current value is 4) :

Enter the **Forward Delay Time** as a number between 4 and 30, or press **<return>** to accept the default listed in parenthesis. This time value controls how fast a port changes its spanning state when moving toward the Forwarding state. The value determines how long the port stays in each of the Listening and Learning states, which precede the Forwarding state. This value is also used when a topology change has been detected and is underway to age out all dynamic entries in the Forwarding Database. Enter a value between 4 and 30 seconds. A value that is too small can cause temporary loops in the network due to data being forwarded before the reconfiguration message has reached all nodes on the network.

9. The following prompt is displayed allowing you to set the Ageing Time:

Ageing Time (10..1000000 sec) (current value is 300) :

Enter the **Ageing Time** as a number between 10 and 1000000, or press **<return>** to accept the default listed in parenthesis. The timeout period in seconds for aging out dynamically learned forwarding information. Enter a new Ageing Time between 10 and 1000000 seconds.

10. The following prompt is displayed allowing you to set the Auto-Tracker VLAN Ageing Time:

Auto-Tracker VLAN Ageing Time (10..1000000 sec) (current value is 1200) :

Enter the **Auto-Tracker VLAN Ageing Time** as a number between 10 and 1000000, or press **<return>** to accept the default listed in parenthesis. The length of time in seconds to remember which VLAN a port belonged to even after the port has been aged out of the Bridge Filtering Database. The MAC and port information are preserved for the set length of time. In the case of IPX it should be set to greater than the server Keep Alive Timer in order to prevent the server from losing communication with the station. The default is 1200 seconds.

11. The final prompt is displayed asking you if you would like to save the new parameters:

Save the new Spanning Tree Bridge parameters ? y/n :

Enter **y** to save the parameters, or **n** to discard them. If you chose to save the parameters, a confirmation message similar to the following is shown:

**Port 5/1 set to Forwarding!
Port 5/2 set to Forwarding!
Port 5/3 set to Forwarding!**

As a variation of this command you can enter the **stc** command without specifying a group. This will allow you to set up spanning tree for the previously selected group. For information on selecting a group see *Selecting a Default Group* on page 17-7.

Display Spanning Tree Bridge Parameters

The **sts** command allows you to display spanning tree bridge parameters. To display spanning tree parameters, enter the **sts** command as shown:

```
sts <groupNumber>
```

where **<groupNumber>** is the number of the group in the switch for which you want to view spanning tree bridge parameters. For example, to view parameters for Group 2, you would enter:

```
sts 2
```

A screen similar to the following is displayed:

```

Spanning Tree Parameters for Group 2 (New GROUP (#2))
Spanning Tree Status      :          ON
Fast Spanning Tree Status:          OFF
Bridge Protocol Use       :          IEE E 802.1D
Priority                   :          32768 (0x8000)
Bridge ID                  : 8000-0020DA:022860
Designated Root           : 8000-0020DA:022860
Cost to Root Bridge       :          0
Root Port                  :          None
Next Best Root Cost       :          0
Next Best Root Port       :          None
Hold Time                  :          1
Topology Changes          :          1
Last Topology Change      : 1 hours, 25 minutes, 54 seconds ago
Bridge Aging Timer        :          300

```

Current Parameters		Parameters system uses when attempt to become root	
Max Age	20 secs	System Max Age	20 secs
Forward Delay	15 sec	System Forward Delay	15 secs
Hello Time	2 secs	System Hello Time	2 secs

As a variation of this command, you can enter **sts** at the system prompt without specifying a group. This will display bridge parameters for the currently selected group. For information on selecting a group, see *Selecting a Default Group* on page 17-7.

Field Descriptions

The following sections describe the fields displayed using the **sts** command.

Spanning Tree Status. Spanning tree is either **ON** or **OFF**.

Fast Spanning Tree Status. Fast spanning tree is either **ON** or **OFF**.

Bridge Protocol Used. The bridge spanning tree protocol is set up through the **stc** command. This protocol can be IEEE 802.1D or IBM Spanning Tree. The type of spanning tree protocol used will affect other bridge parameters, such as **Maximum Age**, **Forwarding Delay**, and **Hello Time**. See *Configuring Spanning Tree Parameters* on page 17-25 for more information on the differences between IEEE and IBM Spanning Tree.

Priority. Bridge priority is utilized by the spanning tree algorithm to decide which bridge will be the root bridge. You can set the bridge priority by entering a decimal number from 0 to 65,535. Zero is the highest priority.

Bridge ID. The bridge identification number is a number created by concatenating the bridge **Priority** with its six-byte MAC address.

Designated Root. The bridge identifier of the root of the spanning tree as determined by the spanning tree protocol. It is created by concatenating the root bridge **Priority** with its six-byte MAC address.

Cost to Root Bridge. The cost of the path to the root bridge as seen from this bridge. Cost represents the distance of the group from the root bridge, in number of hops. If this is the root bridge, this number is 0.

Root Port. The slot number, port number, and service type of the root port. The root port is the bridge's preferred path to the root bridge.

Next Best Root Cost. The next-best available cost of the path to the root bridge as seen from this bridge. Cost represents the distance of the group from the root bridge, in number of hops. If this is the root bridge, this number is 0.

Next Best Root Port. The next-best available root port (slot number, port number, and service type). The root port is the bridge's preferred path to the root bridge.

Hold Time. This time value determines the interval length during which no more than two Configuration Bridge BPDUs shall be transmitted, in seconds.

Topology Changes. The total number of topology changes detected by this bridge since the management entity was last reset or initialized. Topology changes happen when spanning tree reconfigures to prevent logical loops from occurring.

Last Topology Change. The time since the last time a topology change was detected by the bridge entity.

Bridge Aging Timer. The timeout period in seconds for aging out dynamically learned forwarding information.

Max Age. The maximum age (in seconds) of spanning tree protocol information learned from the network on any port before it is discarded.

Forward Delay. This time value (in seconds) controls how fast a port changes its spanning tree state when moving toward the Forwarding state. The value determines how long the port stays in each of the Listening and Learning states, which precede the Forwarding state. This value is also used when a topology change has been detected and is underway to age out all dynamic entries in the Forwarding Database.

Hello Time. The amount of time (in seconds) between the transmission of Configuration Bridge Protocol Data Units (BPDUs) on any port when it is the root of the spanning tree, or trying to become so.

Configuring Spanning Tree Port Parameters

The **stpc** commands allows you to configure port parameters (as opposed to bridge parameters) for spanning tree. To configure port parameters

1. Enter the **stpc** command as shown:

```
stpc <groupNumber>
```

where **<groupNumber>** is the number of the group in the switch for which you want to configure spanning tree port parameters. For example, to configure parameters for Group 1, you would enter:

```
stpc 1
```

As a variation of this command, you can enter the **stpc** command without specifying a group. This will allow you to configure the port parameters on the currently selected group. For information on how to select a group, see *Selecting a Default Group* on page 17-7.

A screen similar to the following is displayed:

Spanning Tree Port Configuration for Group 1 (Default GROUP (#1))

Index	Slot/Intf/Service/Inst	Port Priority (a)	Path Cost (b)	Enable Spanning Tree (c)	tx FA (d)	Manual Mode (e)
1	2/ 1/ Brg/ 1	128	10	y	NA	n
2	2/ 2/ Brg/ 1	128	10	y	NA	n
3	3/ 1/ Brg/ 1	128	10	y	NA	n
4	3/ 2/ Brg/ 1	128	10	y	NA	n
5	3/ 3/ Brg/ 1	128	10	y	NA	n
6	3/ 4/ Brg/ 1	128	10	y	NA	n
7	3/ 5/ Brg/ 1	128	10	y	NA	n
8	3/ 6/ Brg/ 1	128	10	y	NA	n
9	3/ 7/ Brg/ 1	128	10	y	NA	n
10	3/ 8/ Brg/ 1	128	10	y	NA	n
11	3/ 9/ Brg/ 1	128	10	y	NA	n
12	3/ 10/ Brg/ 1	128	10	y	NA	n
13	3/ 11/ Brg/ 1	128	10	y	NA	n
14	3/ 12/ Brg/ 1	128	10	y	NA	n
15	3/ 13/ Brg/ 1	128	10	y	NA	n
16	3/ 14/ Brg/ 1	128	10	y	NA	n

save|cancel|next|prev :

2. To modify a parameter, enter the index (row) number, column letter (a, b, c, d, or e), an equal sign (=), and then the new parameter, as follows.

```
<index><column>=<new parameter>
```

For example, if you wanted to enable transmit Functional Address (**tx FA** in column **d**) for the slot identified by **index 10**, then you would enter:

```
10d=y
```

Field Descriptions

The following section explains the fields displayed by the **stpc** command.

Index

A number assigned as an identifier for the port.

Slot/Intf/Service/Inst

The slot number (**Slot**), interface (port) number (**Intf**), type of service (**Service**), and service instance (**Inst**). For example, a bridge service on port 1 of slot 3 would be:

3/1/Brg/1

Services provide connection options for switches in a LAN, between LANs, or in a WAN. Other possible services include trunking, routing, and LANE. It is possible to have more than one instance of a service if there are more than one connections on a single port.

Port Priority

The value of the priority field contained in the first (in network byte order) octet of the (2 octet long) Port ID. This value allows you to specify a particular port as more favorable if the bridge has more than one port connected in a loop.

Path Cost

The contribution of this port to the path cost towards the spanning tree root bridge that includes this port. 802.1D-1990 recommends that the default value of this parameter be in inverse proportion to the speed of the attached LAN. Path cost is a measure of the distance of the listed port from the root bridge, in number of hops.

Enable Spanning Tree

Whether or not spanning tree is enabled, either **y** or **n**.

tx FA

Transmit Functional Address. Values are:

- | | |
|-----------|---|
| NA | Function Addresses are not applicable because this port is not using spanning tree. |
| y | Transmit Functional Address instead of normal Spanning Tree Multicast Address. |
| n | Transmit normal Spanning Tree Multicast Address. This is the default setting. |

Manual Mode

Allows you to manually set the state for each port (forwarding or blocking) or defer the port's state configuration to the spanning tree protocol, which will either be IEEE 802.1d or IBM. This column is especially helpful if you are using the IBM Spanning Tree protocol with non-Token Ring (e.g., FDDI or Ethernet) ports that do not support this IBM Spanning Tree. In this situation you can manually set those ports to a forwarding (or blocking) state since the IBM Spanning Tree protocol will not be able to control these ports. The possible settings for this column are:

- f** The port is in forwarding state and remains so unless you change it.
- b** The port is in blocking state and remains so unless you change it.
- n** The state of the port is determined by the IEEE 802.1d Spanning Tree protocol. This option is not recommended because it means this Group will have a hybrid spanning tree algorithm that mixes the IEEE 802.1d and IBM Spanning Tree.

Displaying Spanning Tree Port Parameters

The **stps** command allows you to view the current spanning tree port parameters. To view the port parameters, enter the **stps** command as shown:

```
stps <groupNumber>
```

where **<groupNumber>** is the number of the group in the switch for which you want to view spanning tree port parameters. For example, to view parameters for Group 1, you would enter:

```
stps 1
```

A screen similar to the following is shown:

Spanning Tree Port Summary for Group 1 (Default GROUP (#1))

Slot Intf	Service Inst	Pri	State	MAC	Path Cost	Desig Cost	Des Pt	Rt Pt	SwT Pt	Fw Tx	Root Bridge ID Desig BridgeID
3/1	Brg/1	128	FORWD	C473C4	10	10	No	Yes	No	0	0010-0020DA:81D5B0 8000-0020DA:0C41E1

As a variation to this command, you can enter **stps** at the system prompt without specifying a group number. This will allow you to view the port parameters on the currently selected group. For information on how to select a group, see *Selecting a Default Group* on page 17-7.

Field Descriptions

The following section explains the fields displayed by the **stps** command.

Slot/Intf. The slot and interface (port) number of the port.

Service/Inst. The service type and instance of the service connected to the port.

Pri. The value (from 0 to 256) of the priority of the port, 0 being the highest priority.

State. The port's current state as defined by application of the spanning tree protocol. This state controls what action a port takes on reception of a frame. The **State** values are:

Disabled	This port has been disabled.
Blocking	This port is not participating in transmitting data to prevent loops.
Listening	This port is preparing to transmit data, but is temporarily disabled to prevent loops.
Learning	This port is preparing to transmit data, but is temporarily disabled to prevent loops. This is different from Listening in that the port is acquiring data to facilitate data transmission.
Forwarding	This port is transmitting data.

Some of these values are not available if you are using IBM Spanning Tree. For information on the differences between IEEE and IBM Spanning Tree, see *Configuring Spanning Tree Parameters* on page 17-25.

Path Cost. The contribution of this port to the path cost towards the spanning tree root. The spanning tree root will include this port.

Desig Cost. The path cost to the designated port of the segment connected to this port. If this is the root bridge this value is 0.

Des Port. The unique port identifier of the bridge port believed to be the designated port for the LAN associated with the port.

Rt Pt. This field indicates if this port is the root port. The root port is the port that offers the lowest cost path to the root bridge.

SwT Pt. This field indicates if this port is in Optimized Switch Mode. Optimized Switch Mode is appropriate for dedicated connections to a single workstation or server. For more information, see Chapter 19, "Managing Groups and Ports."

FWD Transition. The number of times this port has changed from the Learning state to the Forwarding state.

Root Bridge ID. The bridge identification number of the root bridge.

Desig Bridge ID. The unique bridge identifier of the designated bridge for this port (LAN).

Configuring Fast Spanning Tree

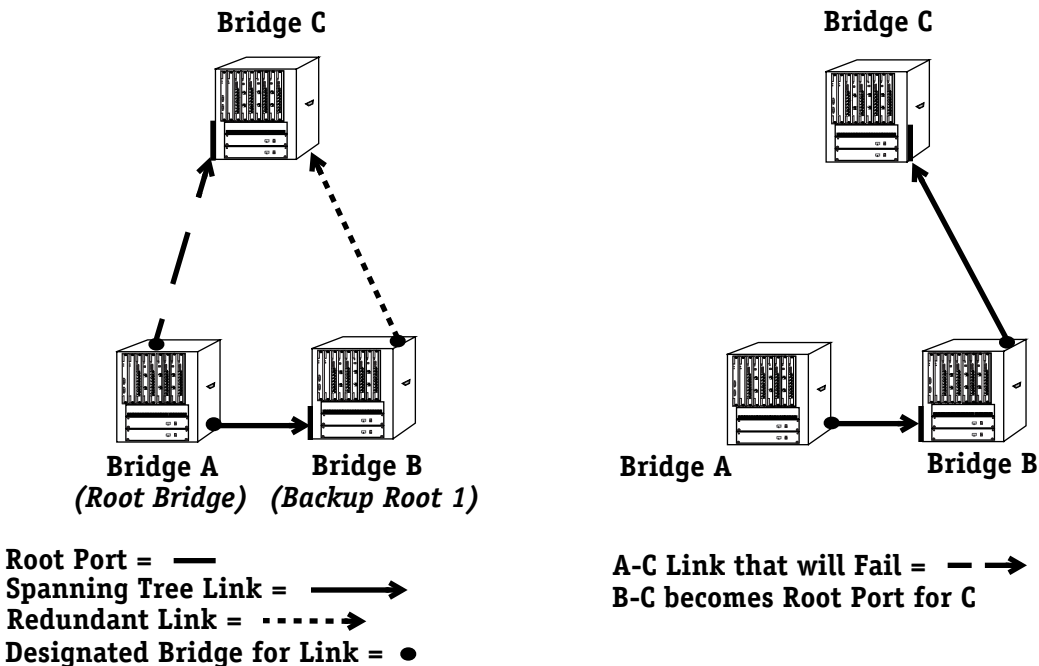
The Fast Spanning Tree (Rapid Reconfiguration) feature is designed to help provide an 802.1D standards-based method of quick recovery in the event of link, port and device failures in an Ethernet local area network. By automatically identifying and utilizing alternative secondary links, Fast Spanning Tree can rapidly converge backup connections between network devices within as little as 1 second. In addition, new Spanning Tree information can be processed faster.

If packets are broadcast to all ports (or flooded) in an attempt to deliver the data, networks with physical loops will rebroadcast packets repeatedly and cause a network to become severely congested. This congestion will adversely affect network performance.

While Spanning Tree prevents broadcast storms by blocking ports in the physical topology that could result in flooded traffic being looped, Fast Spanning Tree minimizes downtime by bringing these blocked secondary links into Forwarding mode as quickly as possible. If the Root Port is lost, an Alternate Port on the Bridge can be made the new Root Port, and placed into a Forwarding state immediately. The prior Root Port switches to a Listening state if it becomes a Designated Port; otherwise, it enters a Blocking state.

Similarly, any Designated Port on the Bridge can be made the new Root Port, and placed into a Forwarding state immediately. In this event, the existing (prior) Root Port changes to a Designated Port role, without a corresponding gain or loss of connectivity. A Backup Port can also be made the new Root Port and placed into Forwarding mode, resulting in the Designated Port assuming a Listening state.

The following diagram illustrates how a typical network connection can fail, such as the A-C Link shown below. Rapid Reconfiguration brings a blocked link - such as the B-C Link - into Forwarding state, helping achieve quick recovery from failure of networked devices.



Recovering from Linked Device Failure with Fast Spanning Tree

Truncating Tree Timing & Speedy Tree Protocol

Two additional enhancements are also included with the Fast Spanning Tree feature for improved performance: Truncating Tree Timing and Speedy Tree Protocol.

Truncating Tree Timing

Truncating Tree Timing allows Designated Ports attached to Point-to-Point links to change to Forwarding mode faster, by utilizing two extra bits in the Configuration BPDU for communication between neighboring bridges. This enhancement promotes quicker restoration of service between communicating stations and reduced flooding of traffic during relearning of station location information.

Speedy Tree Protocol

Speedy Tree Protocol significantly improves reconfiguration performance by allowing inferior information sent by the designated bridge for each LAN to be accepted, rather than timed out. Additionally, information previously received expires immediately on link failure. In both cases, spanning tree recomputation occurs, which can cause changes in both root and designated ports.

Configuring Truncating Tree Timing & Speedy Tree Protocol

Both Truncating Tree Timing and Speedy Tree Protocol are enabled by default. These features are configured by editing the following lines in the command file (`mpx.cmd`):

```
truncatingSt=1  
speedySt=1
```

To disable the Truncating Tree Timing feature, change the numeric entry for `truncatingSt` from **1** to **0**. (To re-enable the feature, change the numeric entry back to **1**.)

To disable the Speedy Tree Protocol feature, change the numeric entry for `speedySt` from **1** to **0**. (To re-enable the feature, change the numeric entry back to **1**.)

◆ Important Note ◆

Do not attempt to edit the command file (`mpx.cmd`) unless you have had significant experience working with files of this type. For additional information, see *Editing Text Files* in Chapter 7, “Managing Files.”

Displaying Fast Spanning Tree Port Parameters

The **fstps** command allows you to view the current Fast Spanning Tree port parameters on a selected group or VLAN. To view the port parameters, enter the **fstps** command as shown:

```
fstps <groupNumber>
```

where **<groupNumber>** is the number of the group in the switch for which you want to view Fast Spanning Tree port parameters. For example, to view parameters for Group 1, enter:

```
fstps 1
```

If Fast Spanning Tree is not enabled (default), a screen similar to the following will appear:

Fast Spanning Tree not enabled for Group 1 (Default GROUP (#1))

Slot Intf	Service		State	Role	Fwrds	Frwdr	FrgetRPs	PPs	Link Ups	Primary Port		
	Inst	Inst								Slot	Service Inst	
8/3	Brg/1		FORWD	ROOT	0	0	0	0	0	2		

As a variation on this command, you can enter **fstps** at the system prompt without specifying a group number. This will allow you to view the port parameters on the currently selected group. For information on how to select a group, see *Selecting a Default Group* on page 17-7.

The fields displayed by the **fstps** command include.

Slot/Intf. The slot and interface (port) number of the port.

Service/Inst. The service type and instance of the service connected to the port.

State. The port's current state as defined by application of the fast spanning tree protocol. This state controls what action a port takes on reception of a frame. The **State** values include:

- DSABL** Disabled - The port has been disabled.
- BLOCK** Blocking - The port is not participating in transmitting data in order to prevent loops.
- LISTN** Listening - The port is preparing to transmit data, but is temporarily disabled in order to prevent loops. BPDU processing does occur, but no user data is being passed.
- LEARN** Learning - The port is preparing to transmit data, adding source MAC addresses to the bridging table, but incoming data frames are dropped.
- FORWD** Forwarding - The port is transmitting data. This state applies to Root Ports and Designated Ports.
- FRWDS** Forwards - The port is transmitting data. This state applies to Designated Ports, and monitors old root ports for a period equivalent to two times the Forward Delay Timer default time period (default = 15 seconds).
- FRWDR** Forwarder - The port is transmitting data. This state applies to Designated Ports, and monitors old root ports for a period equivalent to the Forward Delay Timer default time period (default = 15 seconds).
- FRGET** Forgetting - The port is discarding frames, and is not learning source addresses. This state applies to prior Designated Ports that are placed into an Alternate Role. Forgetting State minimizes potential denial of service due to information races during extensive reconfigurations.

Role. The port's current role as defined by application of the fast spanning tree protocol. The **Role** values include:

- DISABLED** The port has been disabled.
- ROOT** The Root Port on a Bridge has the best path to the Root Bridge, and connects the Bridge to the Root Bridge.
- DESIGNATED** The Designated Port on a Bridge provides an attached LAN the best path to the Root Bridge, and connects the LAN through the Bridge to the Root Bridge, forwarding frames between them. (A Designated Port can be in a Listening, Learning, Forwards, Forwarder, or Forwarding state.)
- ALTERNATE** The Alternate Port is connected to a LAN with another bridge functioning as the Designated Bridge. (An Alternate Port may be in either a Forgetting state or a Blocking state.)
- BACKUP** The Backup Port is connected to a LAN with another port on the same Bridge functioning as the Designated Port. (Backup Ports are always in a Blocking state.)

Frwds. This counter records each instance when the port is in the Forwards state.

Frwdr. This counter records each instance when the port is in the Forwarder state.

Frget. This counter records each instance when the port is in the Forgetting state.

RPs. This counter records each instance when the Root Port is retired.

PPs. This counter records each instance when the Primary Port is retired.

Link Ups. This counter records each instance when the port is linked up.

Primary Port Slot Intf. The slot and interface (port) number of the Primary Port.

Primary Port Service Inst. The service type and instance of the service connected to the Primary Port.

Enabling Fast Spanning Tree Port Parameters

The **actfstps** command allows you to activate Fast Spanning Tree port parameters on a selected group or VLAN. To enable Fast Spanning Tree, enter the **actfstps** command as shown:

```
actfstps <groupNumber>
```

where **<groupNumber>** is the number of the group in the switch for which you want to view Fast Spanning Tree port parameters. For example, to view parameters for Group 1, enter:

```
actfstps 1
```

If Fast Spanning Tree is not enabled (default), a screen similar to the following will appear:

```
Fast Spanning Tree disabled for Group 1 (Default GROUP (#1))
```

```
Enable 1/ Disable 2 Fast Spanning Tree/ Return nothing?
```

To enable the Fast Spanning Tree feature, enter **1** at the prompt. (If you press the Enter key without typing anything, the setting will not be changed.)

No confirmation message will appear. To view the Fast Spanning Tree Port Summary, enter **fstps** at the prompt. For details about the Fast Spanning Tree Port Summary, see *Displaying Fast Spanning Tree Port Parameters* on page 17-36.

◆ Important Notes ◆

To determine whether Fast Spanning Tree is enabled on a VLAN, enter **sts** at the prompt.

To enable Fast Spanning Tree on a VLAN, enter **stc** at the prompt, then follow the onscreen instructions to enable it. For more details, see *Configuring Spanning Tree Parameters* on page 17-25.

Disabling Fast Spanning Tree Port Parameters

The **actfstps** command allows you to disable Fast Spanning Tree port parameters on a selected group or VLAN. To disable Fast Spanning Tree, enter the **actfstps** command as shown:

```
actfstps <groupNumber>
```

where **<groupNumber>** is the number of the group in the switch for which you want to view Fast Spanning Tree port parameters. For example, to view parameters for Group 1, enter:

```
actfstps 1
```

If Fast Spanning Tree is enabled, a screen similar to the following will appear:

```
Fast Spanning Tree Port Summary for Group 1 (Default GROUP (#1))
```

```
Enable 1/ Disable 2 Fast Spanning Tree/ Return nothing?
```

To disable the Fast Spanning Tree feature, enter **2** at the prompt. (If you press the Enter key without typing anything, the setting will not be changed.)

No confirmation message will appear. To view the Fast Spanning Tree Port Summary, enter **fstps** at the prompt. For details about the Fast Spanning Tree Port Summary, see *Displaying Fast Spanning Tree Port Parameters* on page 17-36.

◆ Important Notes ◆

To determine whether Fast Spanning Tree is enabled on a VLAN, enter **sts** at the prompt.

To disable Fast Spanning Tree on a VLAN, enter **stc** at the prompt, then follow the onscreen instructions to disable it. For more details, see *Configuring Spanning Tree Parameters* on page 17-25.

Configuring Source Routing

The **srs** and **src** commands allow you to display and configure the source routing parameters for the selected group.

SAP Filtering

The Service Advertising Protocol (SAP) filter is a method for allowing the user to decide what type of source routed packets are allowed to be transmitted out of the switch. When the filters are configured, they examine the DSAP (destination) and SSAP (source) fields in an outgoing packet, compare them to the filter values to see if they match, and then either allows or blocks packet transmission.

There are two types of filters that can be configured: a “permit” filter and a “deny” filter. If a packet matches the value in a deny filter, and the value is not 0, then the packet is discarded. If a permit filter is configured, and a packet does *not* match the filter value, then the packet is discarded. Only two of each type of filter can be configured.

To use this feature, it must first be enabled, then configured. Once a filter is enabled and configured, it can be viewed as part of the source routing statistics. These procedures are covered in the following sections:

- For information on enabling the SAP filter see *Enabling SAP Filtering* on page 17-40.
- For information on configuring SAP filters, see *Configuring SAP Filtering* on page 17-41.
- For information on viewing SAP filters, see *Viewing SAP Filtering* on page 17-42.

Enabling SAP Filtering

To use the **srsf** command to enable SAP filtering, follow the steps below:

1. Enter the **srsf** command at the system prompt.
2. The following message is displayed:

SAP Filter support is OFF, set it to ON? (n) :

Enter **y** and press **<return>**.

3. Another message is displayed confirming the activation of the SAP filtering feature:

SAP Filter Support is now “ON”

Disabling SAP filtering

To disable the SAP feature, use the **srsf** command as shown:

1. Enter the **srsf** at the system prompt.
2. The following message is displayed:

SAP Filter support is ON, set it to OFF? (n) :

Enter **y** and press **<return>**.

3. The following message is displayed:

Remove all SAP Filter values? (n) :

Enter a **y** to remove the configured filters, or an **n** to keep configured filters, and press **<return>**. See *Configuring SAP Filtering* on page 17-41 for information on how to set up a SAP filter.

4. Another message is displayed confirming the deactivation of the SAP filtering feature:

SAP Filter Support is now "OFF"

Configuring SAP Filtering

Once SAP filtering is activated, it is necessary to configure the filter value. This value is compared to the value of the packets DSAP and SSAP fields. Filters consist of 4 alphanumeric bits, 2 for the DSAP and 2 for SSAP. After enabling SAP filtering, another column is added to the **src** command, and four prompts are added to the ring configuration options.

To configure the filter value:

1. Enter the **src** command at the system prompt. The following screen is displayed:

Source Routing Parameters for Group 1 (Default GROUP (#1))

	Slot	Type/ Inst/Srvc	Ring Number	Bridge Number	Largest frame	HopCnt In Out	Port Type	Block ARE	SAP Filter
1.	2/1	Brg/ 1/ na	1 (0x001)	10 (0xA)	590	6 6	SRT	n	
2.	3/1	Brg/ 1/ na (V)	2 (0x002)	10 (0xA)	4472	7 7	SRT	n	
3.	3/2	Brg/ 1/ na	4 (0x004)	10 (0xA)	4472	7 7	SRT	n	
4.	3/3	Brg/ 1/ na	5 (0x005)	10 (0xA)	4472	6 6	SRT	n	
5.	3/4	Brg/ 1/ na	3 (0x003)	10 (0xA)	4472	7 7	SRT	n	
6.	3/5	Brg/ 1/ na (V)	2 (0x002)	10 (0xA)	4472	7 7	SRT	n	
7.	3/6	Brg/ 1/ na (V)	3 (0x003)	10 (0xA)	4472	7 7	SRT	n	

Enter index of the entry to configure (e.g. 1) **<RETURN>** to exit :

2. Enter the index number (on the far left) for the ring you want to filter.
3. Several prompts for configuring the ring are displayed. Follow the prompts and enter the values required, or accept the current values if the ring is already configured. The following prompt is shown:

Output SAP Deny Filter 1 (0000):

Enter the SAP value that the first deny filter should screen. Any packet matching this filter will be rejected. Excepting the default of **0000** is the same as not having a filter.

4. Press **<return>**. The second deny filter prompt is displayed:

Output SAP Deny Filter 2 (0000):

Enter the SAP value that the first deny filter should screen. Any packet matching this filter will be rejected. Excepting the default of **0000** is the same as not having a filter.

5. Press **<return>**. The first permit filter prompt is displayed:

Output SAP Permit Filter 1 (0000):

Enter the SAP value that the first permit filter should screen. Any packet *not* matching this filter will be rejected. Excepting the default of **0000** is the same as not having a filter.

6. Press **<return>**. The second permit filter prompt is displayed:

Output SAP Permit Filter 2 (0000):

Enter the SAP value that the first permit filter should screen. Any packet *not* matching this filter will be rejected. Excepting the default of **0000** is the same as not having a filter.

7. Press **<return>**. A final message asking to save the new configuration is displayed:

Save the new configuration? (y/n) :

Enter a **y** to save the configuration, or an **n** to cancel the operation.

Viewing SAP Filtering

To see how many SAP filters are configured for a specific ring, enter the **srs** command at the system prompt. A screen similar to the following appears:

Source Routing Parameters for Group 1 (Default GROUP (#1))

	Slot Intf	Type/ Inst/Srvc	Ring Number	Bridge Number	Largest frame	HopCnt In Out	Port Type	Block ARE	SAP Filter
1.	2/1	Brg/ 1/ na	1 (0x001)	10 (0xA)	590	6 6	SRT	n	1
2.	3/1	Brg/ 1/ na (V)	2 (0x002)	10 (0xA)	4472	7 7	SRT	n	2
3.	3/2	Brg/ 1/ na	4 (0x004)	10 (0xA)	4472	7 7	SRT	n	
4.	3/3	Brg/ 1/ na	5 (0x005)	10 (0xA)	4472	6 6	SRT	n	
5.	3/4	Brg/ 1/ na	3 (0x003)	10 (0xA)	4472	7 7	SRT	n	
6.	3/5	Brg/ 1/ na (V)	2 (0x002)	10 (0xA)	4472	7 7	SRT	n	
7.	3/6	Brg/ 1/ na (V)	3 (0x003)	10 (0xA)	4472	7 7	SRT	n	

Enter index of the entry to configure (e.g. 1) **<RETURN>** to exit :

The last column (**SAP Filter**) lists how many SAP filters are in place for the ring. See *Configuring SAP Filtering* on page 17-41 for information on configuring the SAP filter.

Configuring Source Route to Transparent Bridging

In order to provide switching between source-routed token ring networks supporting the IBM Spanning Tree, and transparently bridged networks (primarily Ethernet supporting 802.1d Spanning Tree), commands have been provided in the bridging menu to enable Source Route to Transparent Bridging (SRTB) on a configured group basis.

It is important not to confuse SRTB with source-route transparent (SRT) bridging. SRT bridging is the defined method for bridging on source-routed networks. In SRT bridging, all bridges run the 802.1d Spanning Tree. SRT bridges have the ability to forward a frame based on source-routing information if a Routing Information Field (RIF) is present. Frames without a RIF are bridged transparently. SRT does not provide the ability to switch between a pure source-routed network and a transparent network.

SRTB allows source-routed token ring networks and transparently bridged networks to exist in the same group, and supports connectivity between end systems on the token ring network and the end systems on the transparently bridged network.

The SRTB functions in the following network environments:

- Between token ring and Ethernet networks.
- Between token ring networks and Ethernet LAN emulation (LANE).
- Between token ring LAN emulation and Ethernet networks.

◆ **Note** ◆

Ethernet networks include 10Mbit, 10/100 MB, and Gigabit networks.

Enabling SRTB for a Group

The **srtbcfg** command allows you to display configured groups and the status of SRTB (either **on** or **off**), and to enable or disable SRTB for a specific group. To display groups and the status of SRTB:

1. Enter the **srtbcfg** command at the system prompt, as shown

```
srtbcfg
```

A screen similar to the following is displayed:

```
Group 1: SRTB is OFF
Group 2: SRTB is ON
      Default Explorer: STE Ethernet Ring ID: 291(x123)
Group 3: SRTB is ON
      Default Explorer: ARE Ethernet Ring ID: 561(x231)
```

```
/VLAN SRTB>
```

2. To enable SRTB for a group, enter the **srtbcfg** command at the system prompt, as shown:

```
srtbcfg <groupNumber>
```

where **<groupNumber>** is the number of the group for which SRTB is to be enabled. For example, to enable SRTB for Group 1, you would enter the following:

```
srtbcfg 1
```

3. Once you have entered the command, a screen similar to the following is displayed:

```
Group 1: SRTB is OFF
      Would you like to turn on SRTB ? (n) :
```

Enter **y** to enable SRTB for this group.

4. Once you have enabled SRTB, the following prompt appears:

```
Enter Ring ID for Ethernet segment(s) (0 - 0x0)? :
```

Create a ring ID for the Ethernet segment assigned to this group. This number can be in decimal or hexadecimal form, but it must be unique. For example, if you have a token ring segment with a ring ID of 2, then you could not assign the number 2 to an Ethernet ring ID.

5. Once you have assigned an Ethernet token ID, the following prompt appears:

```
Send Multicast/unknown frames as STE or ARE ? (STE) :
```

Choose to employ Spanning Tree Explorer (STE) frames or All Route Explorer (ARE) frames by entering **ste** or **are**. Explorer frames are sent to learn MAC addresses when there is no record in the RIF table. ARE frames ignore port blocks set up by spanning tree to avoid loops, while STE frames adhere to the spanning tree configuration. The default is **STE**.

6. Once you have selected the frame type, you are returned to the menu prompt. By reentering the **srtbcfg** command as you did in step 1, you can now see that SRTB has been activated for group 1, as shown:

```
Group 1: SRTB is ON
Default Explorer: STE Ethernet Ring ID: 871(x321)
Group 2: SRTB is ON
Default Explorer: STE Ethernet Ring ID: 291(x123)
Group 3: SRTB is ON
Default Explorer: ARE Ethernet Ring ID: 561(x231)
```

The ring ID and default explorer frame are shown as well.

Disabling SRTB for a Group

To turn SRTB off for a group, enter the **srtbcfg** command as shown

```
srtbcfg <groupNumber>
```

where **<groupNumber>** is the number of the group for which you want to disable SRTB. For example, to disable SRTB on Group 3, you would enter:

```
srtbcfg 3
```

The following prompt appears:

```
Group 3: SRTB is ON
Default Explorer: ARE Ethernet Ring ID: 561(x231)
Would you like to turn off SRTB ? (n) :
```

Enter **y** to disable SRTB. Once you have done this you are returned to the system prompt. To view the changes to the group, enter the **srtbcfg** command to display a screen similar to the following:

```
Group 1: SRTB is ON
Default Explorer: STE Ethernet Ring ID: 871(x321)
Group 2: SRTB is ON
Default Explorer: STE Ethernet Ring ID: 291(x123)
Group 3: SRTB is OFF
```

Viewing the RIF Table

A Routing Information Field (RIF) is stored for each MAC address learned on a token ring port. One RIF is stored for each MAC address. The maximum size of each RIF is 32 bytes (long enough to traverse 15 bridge hops)

Once a RIF is learned for a MAC address, it is maintained until the MAC address is aged out of the CAM. You can view a list of RIFs using the **srtbrif** command. To view the RIF table follow these steps:

1. Enter the **srtbrif** command at the menu prompt. The following prompt is displayed:

Enter MAC address ([XXYYZZ:AABBCC] or return for none) :

Enter the MAC address for which you want to see the RIF and press **<return>**, or enter a **<return>** without a MAC address to list all RIFs.

2. Once you enter a MAC address (or **<return>**), the following prompt appears:

Enter Group ID (return for all Group) :

Enter a group ID and press **<return>**, or enter a **<return>** without a group ID to list the RIFs for all groups.

3. Once you enter the group ID (or **<return>**), a screen similar to the following appears:

Port	Group ID	Non-Canonical MAC Address	CAM Indx	Len	RIF
4/ 1/Brg/ 1	2	10009E:4B7DE1	010E	6	0610:1231:0010:

Field Descriptions

The following section describes the fields shown using the **srtbrif** command.

Port. This field lists the slot, port number, service type, and instance number for where the RIF was learned for this MAC address.

Group ID. The group number with which this RIF is associated.

Non-Canonical MAC Address. The MAC address for this RIF. It is shown in non-canonical form.

CAM Indx. The index number in the Content-Addressable Memory (CAM), where the MAC addresses are stored, in hexadecimal form.

Len. The length of the RIF packet, in bytes.

RIF. The RIF address for this MAC address.

Clearing the RIF Table

If there is a topology change in your network, you most likely will need to clear one or more RIFs from the table so that SRTB can relearn them. You can clear specific entries for MAC addresses in the RIF table, or flush the entire table with the **srtbcrrif** command. To clear an entry in the RIF table:

1. Enter the **srtbcrrif** command at the system prompt. The following prompt appears:

Enter MAC address ([000000:000036] or return for none) :

Enter the MAC address for the RIF entry you wish to clear in canonical or non-canonical form, and press **<return>**. If you enter **<return>** without a MAC address, you will flush the entire table of RIF entries.

2. Once you have entered the MAC address, the following prompt appears:

Is this MAC in Canonical or Non-Canonical (C or N) [N] :

If you entered the MAC address in canonical form, enter a **c**. If you entered the MAC address in non-canonical form, enter an **n**. If you respond incorrectly, the RIF entry will not be deleted.

3. Once you entered the distinction of canonical or non-canonical, the following prompt appears to verify the deletion on the RIF entry:

RIF clear successfully!

18 Configuring Frame Translations

Any-to-Any Switching

Because the Omni Switch/Router is a LAN switch that carries frames from multiple media types on its backplane fabric, it offers the facility to switch frames from any media to any other media. For example, an Ethernet frame onto a Token Ring. This feature is referred to as Any to Any Switching.

Normally, the only way for data to get from one media type to another is via routing. Routing removes the media specific headers of a received frame and prepends the new media specific aspects of the destination port before the frame is retransmitted on the new media. In this process the frame itself is not transmitted from one media to another, only the information within it. This process involves heavy computation, requiring table lookups to guide the header deletion/creation and additional router-to-router protocols to set up and maintain these tables.

Routing is not restricted, nor even primarily intended, for moving data between unlike media but instead seeks to break networks down into a number of smaller networks, each of which is a broadcast domain. Historically, networks based on different technologies and media naturally form distinct broadcast domains.

The advent of LAN switching has rewritten these rules. Today, the formation of broadcast domains and the allocation of devices to them is driven by logical requirements such as Virtual LANs and LAN switches. They seek to break free of topology and network constraints imposed by mere media differences.

Within this new paradigm there is still a place for routing. The installed base of clients and servers must communicate by established routing protocols but the broadcast domains handled by a router need not now consist of a single media.

To support this paradigm a LAN switch must “transform” a frame on one media into a frame on the other media in such a way that the frame is still acceptable to the routing protocols. Unfortunately, the requirements for this “transformation” algorithm are specific to the various protocols that currently exist. There is no single, simple algorithm that will allow the frame to be switched between media transparently to the higher level protocols and frame formats. This leads to a fairly complex set of configuration options and limitations on the applicability of the any to any switching features.

In order to understand why these options and limitations arise and to better understand the configuration options available, it is advisable to understand as background the theory of operation of any to any switching. This material is also required if you are trying to determine the applicability of any to any switching to a protocol not described in the reference material.


◆ Important Notes ◆

In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. See Chapter 4, "The User Interface," for documentation on changing from CLI mode to UI mode.

Beginning with Release 4.4, FDDI is no longer supported. Beginning with Release 4.5, Token Ring and ATM are no longer supported.

Translating the Frame

In order to discuss these issues independent of particular media and protocols, consider that every frame, of any protocol, on any media, consists of the following parts.



The Essential Parts of Frame

MAC Header

Consists of a source and destination address specifying the transmitting station in the broadcast domain and the intended recipient(s), as well as other media specific fields. For example, AC and FC fields in Token Ring, FC in FDDI, etc.

RIF (Router Information Field)

If present, it is defined by the source routing standard and is only found on Token Ring and FDDI media.

Encapsulation

Defined by the various standards for the media, many of which reference common standards. For example, on Ethernet media, as defined by Ethernet II, this is a 16 bit type field. On Ethernet media, as defined by the IEEE 802.3 committee, this is a length field together with any encapsulation defined by the IEEE 802.2 Logical Link Control (LLC) committee. On Token Ring and FDDI, it is any encapsulation defined by the IEEE 802.2 LLC committee.

Network Header


Defined by the organization responsible for the particular routing protocol whose data is being carried within the frame. The values of fields defined in the Encapsulation area allow the recipient to identify which protocol standard to use to decode the Network Header part of the frame.

Data

The payload being carried between the end-stations.

In a routing implementation the first three fields (i.e., MAC header, RIF, and Encapsulation) are the ones stripped and rebuilt when the frame is forwarded. These are the three areas that have to be manipulated. The next sections examine each of these frame packet areas further to see the media and protocol dependencies. We can also examine their interactions.

The MAC Header



The format and values defined for the MAC header are covered in the media standards but even here a variety of choices which are dictated by the upper layer protocol can be found.

Canonical versus Non-Canonical

The first requirement of the switch transformation is the bit ordering of the address fields. For Token Ring and FDDI, this is the so called *non-canonical* ordering or most significant bit first. For Ethernet, this is *canonical* or least significant bit first. Thus, when a frame is moved between these media, the addresses must be bit-swapped.

Abbreviated Addresses


The FDDI and 802.5 Token Ring media allow for the use of small 16 bit addresses or full 48 bit addresses. The Omni Switch/Router *only* supports 48 bit MAC address LANs thus abbreviated address based protocols cannot be supported.

Functional Addresses and Multicasts

The 802.5 media also have different rules for the formation of multicast addresses or group addresses. In Ethernet a single bit defines the address as a multicast. In 802.5 a single bit also indicates a multicast but the remaining bits are structured into so called Functional Address groups with pre-assigned meanings and functions.

The Omni Switch/Router *does not* map MCASTs and Functional Addresses; thus protocols dependent on these features may not be switchable any to any.

The RIF Field



The same source routing standard is supported by FDDI and Token Ring so the RIF fields can be switched without problems between these media.

Ethernet does not support source routing thus frames with RIF fields cannot be switched onto these media. However, if you enable "RIF Stripping" you can switch source route frames with RIFs less than 2 bytes long.

The alternative of stripping fields, remembering them and reinserting them on replies, i.e. to terminate a source routed connection and act as a proxy to a transparent device is not well standardized and is difficult to execute and manage.

Source Route Termination by Proxy Not Supported

The Omni Switch/Router will not therefore allow RIF based frames onto Ethernet media unless RIF Stripping is enabled.

Ethernet frames are allowed onto rings if they support transparent bridging, i.e. the port is configured as either Transparent or Source Route/Transparent. Otherwise all communication between SR configured ring ports and transparent Ethernet ports is barred.

Encapsulation



Encapsulation is the biggest problem for implementing a transformation algorithm in support of any to any switching. All of the media provide a choice of more than one encapsulation and not all encapsulations are available on all media. Additionally, the methodology of these encapsulations vary from protocol to protocol.

An ideal protocol would dictate a single encapsulation which would be the same on all media.

Most protocols make use of more than one encapsulation. For example, IP uses Ethertype most of the time on Ethernet and SNAP (an instance of an 802.2 LLC) on FDDI and Token Ring. In this case, there may be clearly established rules for transforming from one encapsulation to another as media are traversed.

Some protocols may allow more than one encapsulation even on a single media type. Some might use the encapsulation to separate functional parts of the protocols, for example, routing table updating protocols from user data forwarding protocols. Others, like IPX may simply allow the user to arbitrarily choose them.

Some, most notably IPX, may entangle the notion of encapsulation with the notion of the network level broadcast domain to create multiple logical networks over a single physical broadcast domain.

Clearly, then there is no single algorithmic rule by which the any to any transformation function can switch arbitrary protocols. There are two choices available to address this situation.

1. The switch must be configurable, per device, per protocol, per media to select the transformation of encapsulations.
2. The switch performs a single transformation and the user must configure all end-stations and routers to use this single choice made by the switch.

The Omni Switch/Router uses the first approach for IP and IPX as the dominant protocols in the market. It uses the second approach for all other protocols.

Protocols other than IP and IPX

For protocols other than IP or IPX three encapsulations are possible on Ethernet media:

- Ethertype
- IEEE 802.2 LLC
- IEEE 802.2 SNAP (This is an instance of an LLC encapsulation defined by the 802.2 committee to support the transformation of Ethertype Ethernet frames to media which don't support that encapsulation.)

On Token Ring and FDDI, two encapsulations are permitted by the standards:

- IEEE 802.2 LLC
- IEEE 802.2 SNAP.

The SNAP Conversion

The intent of the 802.2 committee is that Ethertype frames are transformed to SNAP on crossing from Ethernet media to 802 media and restored to Ethertype in the reverse direction.

The Omni Switch/Router could follow this rule for all protocols including IP; however, this would prevent AppleTalk interworking between Ethernet and FDDI. The Omni Switch/Router explicitly checks for the AppleTalk protocol. If found, the rule is not applied. In addition, the Omni Switch/Router checks for the Banyan Vines protocol and translates according to the media type (see *Banyan Vines* on page 18-13).

As there may be other protocols with this problem, the SNAP-to-Ethertype transformation is configurable for all protocols other than AppleTalk.

Other Conversions

There are no equivalent algorithmic approaches which the transformation function can adopt for dealing with protocols which require Ethertype on Ethernet and some form of LLC encapsulation on FDDI and/or Token Ring. The mapping between Ethertype values and LLC values is arbitrary requiring tables indexed by protocol.

The approach followed in the Omni Switch/Router is therefore to simply pass LLC encodings between Ethernet, FDDI and Token Ring with no changes other than to insert/strip the length field required by IEEE 802.3 on Ethernet.

This leaves protocols which require transformations between Ethertype and LLC encapsulations as unswitchable unless the clients and servers can be configured to use SNAP.

Summary of Non-IPX Encapsulation Transformation Rules

To summarize:

- Ethertype/SNAP transformations are configurable for all protocols except AppleTalk and Banyan Vines. Ethertype frames going to FDDI or Token Ring are translated to SNAP unconditionally. SNAP frames going to Ethernet are translated to Ethertype or left as SNAP as per configuration, unless the protocol is AppleTalk in which case they are left as SNAP.
- LLC frames are passed unchanged in value but with the length field required on Ethernet media stripped/inserted.

IPX Encapsulation Transformation Rules

For IPX the encapsulation problems described above are compounded by the introduction of a fourth encapsulation on Ethernet media. Novell introduced a frame format when the IEEE 802.3 standards committee produced its version of Ethernet which was incompatible with Ethernet.

Novell places its network header and data within a raw IEEE 802.3 Ethernet frame with no intervening IEEE 802.2 LLC header. This is in direct contravention of the standards but has become a de facto standard encapsulation.

Novell refers to this encapsulation types as ETHERNET_802.3. It is also widely known as Novell Proprietary, Novell Raw, Raw 802.3, etc. Such frames are identifiable only by the fact that the Novell Network header starts with a two byte field called the *checksum*, which is never used and assumes the value 0xFFFF.

Routers, bridges and switches therefore check for the checksum after an 802.3 length field. In effect, Novell has usurped the value 0xFF for the Destination and Source SAP addresses (DSAP/SSAP) of an LLC header.

Thus on Ethernet media there are four encapsulations for IPX

- Ethertype - value 0x8137
- Novell Proprietary
- LLC - SAP value 0xE0
- SNAP - Protocol Identifier 0x0000008137

On Token Ring and FDDI, the same LLC and SNAP encapsulations are found as on Ethernet (without the length field.)

This leaves an aggregate of four encapsulations across all media with only two being universal (LLC and SNAP).

Unfortunately, the SNAP conversion rule isn't applicable and there is no algorithmic determination for the use of particular encapsulations on any media - it's purely the choice of the network administrator. Worse, multiple encapsulations can be found on a single media to create multiple logical networks over a single physical broadcast domain.

The Omni Switch/Router therefore allows configuration of the encapsulation transformations of IPX frames. Before transmission of a frame occurs the switch determines first the current encapsulation of the frame. Then, it consults configuration information to determine which of the permitted encapsulations for the media the frame is to be transmitted on is required. Thus, the administrator can choose not only a single output option but an option per possible received encapsulation.

For example, over FDDI media, LLC and SNAP are permissible so the administrator might configure one of the following:

- LLC and SNAP encapsulations received from other FDDI, Token Ring or Ethernet media are translated to SNAP.
- Ethertype and Proprietary encapsulations from Ethernet are translated to LLC.

Essentially, for each encapsulation, transformation to each of the other three encapsulations is available, but may simply be left as is. This choice may be further constrained by the output media type, for example, Ethertype is not a valid option on FDDI or Token Ring.

The Network Header



There are essentially two requirements for the any to any switching transformation function to address the network header fields:

- Network Address to MAC Address Mapping
In every protocol there is a mechanism for mapping global network wide addresses to the MAC addresses required in the local broadcast domain.
- Frame Size Requirements of the Media
Different media have different minimum and maximum frame sizes leading to the issues of padding insertion/stripping and fragmentation/reassembly or maximum frame size negotiation protocols at the network level.

Address Mapping

There are almost as many ways to map a global network level address to a local subnetwork MAC address as there are routing protocols. These may or may not be affected by any to any switching.

Some may construct MAC addresses algorithmically, for example, DECNET model. Some may involve table lookups with an additional protocol to build and maintain these tables, for example, the IP/ARP model. Others may involve some form of building the network address around the MAC address as in the IPX model.

In all cases these mechanisms are susceptible, without good design and forethought, to the problem of canonical versus non-canonical representation of addresses in the network header area.

Address Mapping in IP: ARP

To map a 32-bit IP network address into the MAC address of a locally connected station a router uses the Address Resolution Protocol (ARP) to build an ARP Table. The router broadcasts a request containing the IP address in the body of the frame. The station with that IP address responds with its MAC address *in the body* of an ARP reply frame. The router inserts these two addresses in its ARP table and can then use the MAC address received to transmit any frames addressed to that IP address.

Since a router can have interfaces to Ethernet ports (canonical MAC addresses) and FDDI and Token Ring (non-canonical MAC addresses), it is crucial that the router keeps track of what media type it receives on each port.

If IP ARP were defined such that all MAC addresses, *when conveyed in the body of an ARP*, were in canonical format, switching would be easy. A router, when taking an address from the ARP table and using it as the destination MAC address on an Ethernet port would use the address as is. If sending to FDDI or Token Ring it would bit swap the address to non-canonical format as required by the media.

Given this model of implementation a station responding with an ARP on Ethernet which was switched to FDDI would result in the same representation of the MAC address in the ARP table of the router. The router would then use the bit swapped form in the MAC address of subsequent frames to the FDDI ring and the switch would bit swap these MAC header address as it transformed the frame onto Ethernet, resulting in the correct representation to be received by the original station.

Unfortunately, this model has only been defined in IP for Ethernet and FDDI. Token Ring stations place MAC addresses into the body of ARP frames in their native, non-canonical format and routers use addresses from the ARP table as is when sending to Token Ring ports.

To achieve any to any switching with IP it is therefore necessary for the Omni Switch/Router to be sensitive to ARP frames and to bit swap the MAC addresses *in the body of the ARP* when switching a frame between Token Ring and FDDI or Ethernet.

◆ Important Note ◆

Beginning with Release 4.4, FDDI is no longer supported.

Because IP is well designed, the issue of address mapping being confined to the ARP protocol, this is sufficient to isolate the problem allowing all subsequent IP frames to be switched any to any.

Address Mapping in IPX

A network address in IPX consists of three parts:

1. Network Number -- a globally unique identifier of a particular broadcast domain.
Strictly, because of the formation of logical networks using encapsulations, this is not equivalent to a physical broadcast domain but the distinction can be put aside for the purposes of this particular discussion.
2. Node Address -- the MAC address of a station on that domain.
3. Socket Number -- the task (process) within that station which should process the message.

Just as in IP, routers move a frame along hop by hop on the basis of the network number portion of the destination address. To do this, IP needs the MAC address of the next hop router. This address is obtained from the RIP table that is built up from the RIP updates sent out by all routers. When a router receives a RIP update frame it uses the source node address in the frame as the MAC address for the next hop router.

Although there is not an explicit ARP like protocol for mapping addresses in IPX, this same function is achieved by the use of source node addresses in RIP frames.

In IPX, as in IP, the canonical versus non-canonical representation of addresses in ARPs still applies. In switching, this needs to be considered for the source node address in IPX frames.

In IPX Ethernet and FDDI observe a convention of using MAC addresses in the IPX header in canonical format. For Token Ring these addresses are non-canonical.

Proprietary Token Ring IPX switching

The Omni Switch/Router offers the facility to modify IPX frames switching between Token Ring and FDDI or Ethernet. ARP bit swapping for IP is a de facto standard widely implemented in the industry. This is not the case with IPX. The switch must be able to co-exist with bridges that do not support any to any switching or applications where this feature is not required. Therefore this feature can be configured on or off.

Frame Size Requirements

The frame size requirement for the different media cause two problem areas which have to be addressed by the any to any switching transformation function.

- Ethernet has a minimum frame size requirement. This requires that padding is inserted on frames switched to it which are below the minimum size and stripped from frames switched from it.
- All media have different maximum frame size requirements. This gives rise to the problems of fragmenting large frames and/or negotiating maximum frame sizes.

Insertion of Frame Padding

Ethernet has a minimum frame size of 64 bytes. For frames smaller than 64 bytes it is a simple task for the Omni Switch/Router to perform padding. Stripping such padding from Ethernet frames when switching to FDDI or Token Ring is not so easy.

In most implementations of IP that we have tested the presence of padding on FDDI or Token Ring frames appears not to cause any problems. However, IPX implementations are adversely affected by its presence. Therefore the Omni Switch/Router takes a conservative approach *for all frames*, regardless of protocol type, and strips padding *where it can be detected*.

Stripping of Padding for all IEEE 802.3 Frames.

Ethernet frames in IEEE 802.3 format can be stripped of padding because of the presence of the length field. This includes all LLC and hence SNAP encapsulated protocols as well as Novell Proprietary format.

No stripping of non-IPX Ethertype Frames

Padding can only be detected for Ethertype encapsulated frames if the protocol is known and the protocol has some length information which can allow the valid data size to be inferred. This is protocol specific and is currently only performed for IPX frames. Thus, the Omni Switch/Router *does not* strip padding from non-IPX Ethertype encapsulated frames *including IP*.

IPX Specific Stripping

For IPX the Omni Switch/Router performs pad stripping for all frame types including Ethertype. This is possible because all IPX frames have a common header that includes the data length, allowing the frame size to be inferred.

In fact, for IPX, the length in the IPX header is used to strip padding in all frame encapsulations including the 802.3 based formats. This is because many IPX Ethernet implementations also pad frames to an even byte length. This single byte pad when performed on 802.3 based frames is included in the 802.3 length field. Thus the generic 802.3 based stripping technique is not sufficient to strip this odd-byte padding. When performing any to any switching FDDI implementations of IPX were found to be tolerant of this extra byte whereas Token Ring implementations would not work with it present. By adopting the single IPX stripping strategy of using the IPX header length these problems are avoided thus the Omni Switch/Router unconditionally strips all padding from IPX frames.

Also, it *does not* support odd-byte pad insertion when switching to Ethernet. This was a feature added to overcome limitations of some NIC cards which is now of only historical importance and in fact, Netware 4.1 servers provide this insertion as a port configuration option.

MTU Handling

Routers address the problem of maximum frame size limitations with the notion found in many protocols of a Maximum Transmission Unit (MTU) size. Protocols use this notion in two possible ways.

- PDU Fragmentation/Reassembly

The router is configured with the MTU of each port. If a frame that is too large is required to be sent on a port, the Protocol Data Unit (PDU) within the frame is fragmented into many smaller PDUs, each of which is re-encapsulated and sent as a frame that fits within the MTU.

- Connection-oriented end-to-end MTU negotiation

When an end-station enters into a protocol to communicate with another station the initial PDU exchanges are guaranteed to fit all possible MTUs. In the handshaking between end-stations to establish the connection a phase is entered where large frames are sent. If an intervening link has an MTU too small for these frames it will be dropped and the handshaking will time out. The end-stations send progressively smaller frames until the handshaking succeeds and hence establish the MTU to be used between the two stations for the remainder of their connection use.

IP supports the former mechanism and IPX the latter.

IP Fragmentation

The Omni Switch/Router Ethernet interfaces will use IP fragmentation if they are allowed to (i.e., if the *Don't Fragment* bit is not set.) Fragmentation by FDDI and Token Ring is not supported though technically the Token Ring could send frames larger than those supported by FDDI and LAN Emulation could generate frames larger than both.

ICMP Based MTU Discovery

IP uses the Don't Fragment bit to support an MTU discovery protocol that superficially resembles the negotiation of IPX. The difference is that when IP stations attempt to discover an MTU size for their use, which doesn't require fragmentation by intermediate routers, the protocol expects a protocol response *by the intermediate router*, this is an ICMP reporting that a frame was dropped because it couldn't be fragmented.

The Omni Switch/Router transformation function of any to any switching *does not* support this ICMP generation but just silently drops IP frames which can't be fragmented. The IP *router* in the Omni Switch/Router does honor this protocol and support ICMP. It is only the any to any switching which doesn't because it is not a router and may not even have an IP address with which to respond.

IPX Packet Size Negotiation

For IPX the requirement of intervening devices is simply to drop frames that are too large to be forwarded. This is what the Omni Switch/Router does.

Other Protocols

Dropping oversize frames is the approach for all protocols other than IP. If the protocol in question is modeled like IPX this will be the correct thing to do and will not cause problems. If the protocol is modeled like IP and expects fragmentation to occur or requires explicit response from the Omni Switch/Router then the protocol will not succeed in any to any switching.

Banyan Vines

Banyan Vines supports Ethernet, FDDI, and Token Ring networks. Each type of network generates a different frame format, so the Omni Switch/Router performs translations for frames moving from one network type to another. The Banyan Vines protocol only uses one frame format per network type—no user configuration of translations is necessary. This protocol uses Ethernet II frames on Ethernet, SNAP frames on FDDI, and IEEE 802.2 (LLC) frames on Token Ring. The Omni Switch/Router uses these frame formats when translating Banyan Vines frames.

Note

Checksums for Banyan Vines frames are automatically set to the null checksum, 0xFFFF, so that the checksum header does not require recalculation. Receiving stations will ignore this field and assume the sender is not using checksums.

Configuring Encapsulation Options

You will configure frame encapsulation based on the destination MAC address or the destination switch port. Whether a frame is encapsulated based on the destination MAC or the port depends whether the frame has a unicast, multicast, or broadcast destination.

Forwarding versus Flooding

Such frames will be handled in two ways:

- **Forwarded Frames.** If the frame has a unicast destination address which has been learned on a particular port, the encapsulation translation choices are driven by options associated with the destination MAC address.
- **Flooded/Multicast Frames.** If the frame has a unicast destination address which has not been learned on a particular port, or if the destination address is a multicast address, then the frame has to be transmitted on potentially many ports. In this case the encapsulation translation choices are driven by options associated with each destination port.

Port Based Translation Options

The translation options for ports allow configuration of IP and IPX protocols on a per encapsulation basis.

MAC Address Based Translation Options

The translation options for MACs arises from two possible sources.

- Inheritance from Port Options During Source Address Learning
- When a source MAC address is learned, the translation options of the port on which it is learned are copied into the MAC-based database.
- Automatic Determination by AutoTracker
- When a frame is processed by AutoTracker as part of determining the VLAN to be associated with the MAC the frames protocol type and encapsulation are also determined. This information is used to update/set the translation options in the MAC based database.

Which of these options is used is determined by setting the autoencaps option.

“Native” versus “Non-Native” on Ethernet

For the Ethernet one further distinction is made. If the frame received from the backplane is an Ethernet media type frame from another Ethernet switching module in the same chassis, then *no encapsulation translations are applied*. Such frames are referred to as Native frames.

If the frame is of an Ethernet media type but was put onto the backplane by some other type of switching module, for example, the frame came from a FDDI card via a trunk port, or from the MPX via routing, then *encapsulation translations are applied*. Such frames are referred to as Non-Native frames.

◆ Important Note ◆

The `.cmd` file contains a command called `hreXnative` that by default is set to `1`. If your switch uses multiple encapsulations (for example, VLAN 2:1 is 802.3 IPX and VLAN 3:1 is Ethernet II IPX) then the `hreXnative` command must be set to `0`. See Chapter 7, “Managing Files,” for more information on the `.cmd` file.

“Native” versus “Non-Native” on FDDI and Token Ring

For FDDI, Token Ring and LAN Emulation on ATM, a native/non-native distinction is not made. Instead, no encapsulation translations are applied by these switching modules to frames which are of their own media type.

No Translation on Trunk or PTOP ports

Switching modules which support encapsulation mechanisms, such as Trunking ports on FDDI and Token Ring, and Point to Point ports on ATM do not apply translation to frames destined to such ports.

All other aspects of the transformation process are driven by the media type of the frame, the media type of the port on which the frame is to be transmitted and the protocol type determined for the frame. Thus frame padding insertion/stripping, IP fragmentation, IP ARP bit swapping, etc., are all automatic.

The Proprietary Token Ring IPX Option

The one area which remains configurable is the bit swapping of source addresses for IPX in order to allow Token Ring to work with FDDI and Ethernet. This is the equivalent function to IP ARP bit swapping.

This option is configurable and by default is on.

The User Interface

This chapter documents User Interface (UI) commands to configure encapsulation options. For documentation on Command Line Interface (CLI) commands to configure encapsulation options, see the *Text-Based Configuration CLI Reference Guide*.

◆ Important Note ◆

In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. See Chapter 4, “The User Interface,” for documentation on changing from CLI mode to UI mode.

Simple encapsulation options can be configured through the **modvp**, **addvp**, **crgp** commands. More advanced encapsulation options can be found in the commands under the **Switch** menu.

Essentially, the forwarding code is now capable of applying the transformation function per protocol per encapsulation per port for flooded/mcast traffic and per protocol per encapsulation per destination MAC address for forwarded unicast traffic. The old interface provides a small subset of these possible port translation options.

The `advp`, `modvp` and `crgp` Commands

All of these commands include in their dialogue an Output Format question for ports and a subsidiary IEEE 802.2 Pass through option.

The options offered are:

- a default,
- Ethertype,
- SNAP and
- LLC.

Each of these represents a set of translation options for the IP and IPX protocols. The names chosen for these sets basically represent the translations for IPX with the translation for IP being implied.

For example, LLC represents a translation set where all IPX encapsulations are configured to translate to IEEE 802.2. This is not a valid encapsulation for IP which is therefore configured to a default appropriate to the media, Ethertype for Ethernet ports and SNAP for FDDI and Token Ring ports. The translation of all other protocol types and encapsulations is fixed by the Omni Switch/Router. Thus AppleTalk is never translated and Ethertype/SNAP based protocols follow the IP option.

For those options which imply a translation of IEEE 802.2 IPX frames to something else a subsidiary question is asked, "IEEE 802.2 IPX Pass Through(y/n):" An IEEE 802.2 pass through option is provided because 4.1 Novell servers use this encapsulation by default and it is becoming Novell's encapsulation of choice.

The Default Translation Option

The meaning of the default is determined separately for each media type and is fully configurable. The factory defaults are chosen so that the latest release is fully compliant with earlier ones. The default translation option is provided to allow a "single point of configuration of all ports" capability. When the default option for a media is changed all ports of that media type whose encapsulation is configured as default will inherit the new translation setting. All MAC address-based translation options which were inherited from those ports, as opposed to those set by AutoTracker, will also be updated. Ports which have an encapsulation setting other than default will be unaffected.

Ethernet Factory Default Translations

For Ethernet switching module ports the factory default is set to the following:

Ethernet Media - Default Mode
No translation is performed on outbound Ethernet frames where the inbound interface was Ethernet.
IP frames of any encapsulation are transmitted as Ethernet II frames.
IPX frames are transmitted as IEEE 802.3 Proprietary as the default setting. The only exception is when LLC passthrough mode is enabled, then the IEEE 802.2 (LLC) frames are forwarded as is.
No translation is performed on Appletalk frames, and we currently support only Appletalk Phase II (SNAP format).
Banyan Vines frames are transmitted as Ethernet II frames.
Other than IP and IPX, all other Ethernet II and SNAP encapsulated protocols are sent as Ethernet II frames.
All other IEEE 802.3 with LLC encapsulated protocols are not translated.

FDDI Factory Default Translations

For FDDI switching module ports the factory default is set to the following:

FDDI Media - Default Mode
IP of any encapsulation is encapsulated SNAP
IPX encapsulations are encapsulated SNAP except for IEEE 802.2 which is forwarded as is.
Banyan Vines of any type are transmitted as SNAP.
All other Ethertype and SNAP encapsulated protocols are sent as for IP.
All other LLC encapsulated protocols are forwarded as is.

Token Ring Factory Default Translations

For Token Ring switching module ports the factory default is set to the following:

Token Ring Media - Default Mode
IP of any encapsulation is encapsulated SNAP
IPX encapsulations are encapsulated SNAP except for IEEE 802.2 which is forwarded as is.
Banyan Vines of any type are transmitted as LLC.
All other Ethertype and SNAP encapsulated protocols are sent as for IP.
All other LLC encapsulated protocols are forwarded as is.

ATM LANE Factory Default Translations

For ATM LAN Emulation service ports the factory default is set to the following:

ATM LANE - Default Mode
No translations performed on Ethernet frames.
FDDI and Token Ring frames are translated to either SNAP or LLC and are transmitted as such on ATM LANE.
Banyan Vines Token Ring and FDDI frames are translated to Ethertype.

The Ethertype Option

This option can only be applied to Ethernet switching module ports. It is set to the following:

Ethernet Media - Ethernet II Mode
No translation is performed on outbound Ethernet frames where the inbound interface was Ethernet.
IP frames are transmitted as Ethernet II frames.
All IPX frames are transmitted as Ethernet II frames. The only exception is when LLC passthrough mode is enabled, then the IEEE 802.2 (LLC) frames are forwarded as is.
No translation is performed on Appletalk frames, and we currently support only Appletalk Phase II (SNAP format).
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as Ethernet II frames.
Other IEEE 802.3 with LLC are not translated.

ATM LANE - Ethernet II Mode
IPX frames from FDDI, Token Ring, and Ethernet SNAP frames are translated to Ethertype.
All other SNAP frames from FDDI, Token Ring, and Ethernet SNAP are translated to Ethertype. However, Appletalk ARP SNAP frames from Token Ring and FDDI are left as SNAP; Banyan Vines frames from FDDI are translated to Ethertype.
All other 802.2 frames from FDDI, Token Ring, and Ethernet are left as is. The exception are Banyan Vine frames from Token Ring, which are translated to Ethertype.
All Ethernet Ethertype frames are not translated.

The SNAP Option

This option can be applied to all media type ports and is set to the following:

Ethernet Media - SNAP Mode
No translation is performed on outbound Ethernet frames where the inbound interface was Ethernet.
IP frames are transmitted as SNAP frames.
All IPX frames are transmitted as SNAP frames.
No translation is performed on Appletalk frames, and we currently support only Appletalk Phase II (SNAP format).
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as SNAP frames.
Other IEEE 802.2 with LLC are not translated.

FDDI / Token Ring Media - SNAP Option
No translation is performed on outbound FDDI or Token Ring frames where the inbound interface was the same media type.
IP frames of any encapsulation type are transmitted as SNAP frames.
IPX frames received that do not have an IEEE 802.2 encapsulation type, are transmitted as SNAP.
IPX frames received that are of IEEE 802.2 encapsulation type are transmitted as SNAP if the LLC passthrough is disabled. If the LLC passthrough is enables, these frames will not be translated.
No translation is performed on Appletalk frames, and we currently support only Appletalk Phase II.
All other LLC encapsulated protocols are left as is.

In the **modvp** or **addvp** commands for FDDI and Token Ring the only choices other than default are SNAP or LLC and the default must be one of these. As the factory default is SNAP with IPX 802.2 Pass through and the SNAP does not imply pass through the additional question about pass through is not asked on FDDI and Token Ring ports as the preference can be expressed by choosing default or SNAP explicitly.

ATM LANE - SNAP Mode

All IPX frames are translated to SNAP unless they are already SNAP, in which case they are forwarded as is.

All Ethertype or SNAP frames from Ethernet and SNAP frames from Token Ring or FDDI are translated to SNAP or left as SNAP. The exception is Banyan Vines frames from FDDI, which are translated to Ethertype.

All other LLC frames are left as is. The exception is Banyan Vines from Token Ring, which is translated to Ethertype.

The LLC Option

This option can be applied to all media type ports and is set to the following:

Ethernet Media - LLC Mode
No translation is performed on outbound Ethernet frames where the inbound interface was Ethernet.
IP frames are transmitted as Ethernet II frames.
All IPX frames are transmitted as IEEE 802.2 (LLC) frames.
No translation is performed on Appletalk frames, and we currently support only Appletalk Phase II (SNAP format).
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as Ethernet II frames.
Other IEEE 802.2 with LLC are not translated.

FDDI / Token Ring Media - LLC Mode
No translation is performed on outbound FDDI or Token Ring frames where the inbound interface was the same media type.
IP frames are transmitted as SNAP frames.
All IPX frames are transmitted as IEEE 802.2 (LLC) frames.
No translation is performed on Appletalk frames, and we currently support only Appletalk Phase II (SNAP format).
Other than IP and IPX, all other Ethernet II or SNAP frames are transmitted as SNAP frames.
Other IEEE 802.2 with LLC are not translated.

In the **modvp** or **addvp** commands for FDDI and Token Ring the only choices other than default are SNAP or LLC and the default must be one of these. As the factory default is SNAP with IPX 802.2 Pass through and SNAP does not imply IPX 802.2 Pass through, the additional question about pass through is not asked on FDDI and Token Ring ports. By choosing SNAP, it is implied that there is no IPX 802.2 Pass through.

ATM LANE - LLC Mode
IPX frames are translated to 802.2 LLC.
All other SNAP frames from FDDI, Token Ring, and Ethernet SNAP are translated to Ethertype. However, Appletalk ARP SNAP frames from Token Ring and FDDI are left as SNAP; Banyan Vines frames from FDDI are translated to Ethertype.
All other LLC frames are not translated. The exception is Banyan Vines frames from Token Ring, which are translated to Ethertype

Interaction with the new interface

If the port to which these commands are being applied has been configured with the new interface commands its encapsulation will be displayed as **SWCH** in the **vi** command output. The user is alerted to this fact in these commands by the default response to the output format question in the **modvp** command being displayed as "*" instead of **d,e,s** or **l**. A simple return will leave the options unchanged in this case. If the port is currently one of **d,e,s** or **l** and the user types "*" in response the encapsulation is changed to **SWCH** and the options are set to a null translation set.

The "vi" Command

The encaps column displays the encapsulation subset options set for each port. If the port has been configured with the new interface this is indicated by displaying "SWCH." The "canned" subsets offered in this interface are displayed as follows:

- **DFLT.** This indicates that the port is using the default translation options applicable to the media type of this port. See above.
- **802.2.** This indicates that IPX frames of any encapsulation will be encapsulated with IEEE 802.2. Non-IPX frames other than AppleTalk will be transformed to Ethertype on Ethernet ports and SNAP on FDDI or Token Ring ports. AppleTalk frames are never transformed.
- **SNAP.** This indicates that Ethertype frames of all protocols and IPX proprietary frames will be translated to SNAP and all SNAP frames will be left as is.
IEEE 802.2 encapsulated IPX frames may be left as is if the IEEE 802.2 pass through option is in effect for this port. All other IEEE 802.2 encapsulated protocols are left as is.
- **ETH.** This indicates that SNAP frames of all protocols except AppleTalk will be translated to Ethertype.
SNAP and Proprietary IPX frames will be transformed to Ethertype.
IEEE 802.2 encapsulated IPX frames may be left as is if the IEEE 802.2 pass through option is in effect for this port.
All other IEEE 802.2 encapsulated protocols are left as is.

To discover whether IEEE 802.2 pass through is in effect on a port the user must either use the **swch** command from the switch menu or use **modvp** and observe the encapsulation offered and/or the default response for the pass through question.

The Switch Menu

The switch menu contains commands that allow you to set translation options discussed earlier in this chapter. It also contains commands to change the default values.

To view the switch menu, enter **switch** at the prompt. If you are in verbose mode, the following screen is displayed. Otherwise, type a **?** at the switch menu prompt to display the Switch Menu:

Command	Switch Menu
propipx	Configure Default Proprietary IPX Token Ring to any switching
facdef	Configure Defaults to Factory values
ethdef	Configure Default Ethernet Translation
fdiddef	Configure Default FDDI Translation
trdef	Configure Default TR Translation
swch	Configure Any To Any Switching Port Translations
swchmac	View per MAC Translation Options
autoencaps	Turn AutoTracker translations On or OFF

The commands above and their operations are described in the sections below.

Proprietary IPX Token Ring

The **propipx** command allows you to turn on or off the default proprietary IPX switch translation. (Refer to Appendix B, “Output Translation Options,” for information on the Proprietary IPX feature.)

To turn on the Proprietary IPX feature (the default), enter the following at the system prompt:

```
propipx on
```

A message is displayed to confirm the activation of the Proprietary IPX feature. Please note that the switch must be rebooted for the setting to take effect.

To turn off the Proprietary IPX feature type:

```
propipx off
```

Factory Defaults

You can reset all ports in the switch to their default factory settings. Any custom translations you configured through **modvp**, **ethdef**, **fdiddef**, **trdef**, or **swch** commands will be overridden by the default translation for the given media type (i.e., Ethernet, FDDI, etc.). Factory defaults for each media type are described earlier in this chapter.

To reset to factory defaults, enter the **facdef** command at the system prompt. The following screen displays:

```
This will reset the default translations for each media type to a factory default.  
It will then set all port translation options to inherit these defaults.  
It will then reset the forwarding table translation options for all addresses learnt on  
those ports to those port defaults.  
Do you want to do this? (no):
```

Enter a **Y** to reset all port settings.

Default Ethernet Translations

The **ethdef** allows you to set up default translations for all Ethernet ports. To do so:

1. Enter **ethdef** at the system prompt. The following screen displays:

```
This will reset the default translations for Ethernet media to a new value.
All Ethernet ports currently set to default will inherit these new translation options.
It will then reset the forwarding table translation options for all addresses learnt on
those ports to those port defaults.
Do you want to do this? (no):
```

2. Press **Y** at the **Do you wish to do this?** prompt to indicate that you want to change the defaults. The current settings for Ethernet ports are displayed, in a screen similar to the following:

```
Translation Options:
1      IP Ethertype          -> Ethertype
2      IP IEEE 802 SNAP     -> Ethertype

3      IPX ETHERNET_II      -> 802.3
4      IPX ETHERNET_802.3   -> 802.3
5      IPX ETHERNET_802.3/FDDI/TOKEN_RING -> 802.3
6      IPX ETHERNET_SNAP/FDDI_SNAP/TOKEN-RING_SNAP -> 802.3
```

There are six frame types for which you can set translation options. The frame type in the left column indicates the incoming frame, and the frame type in the right column (after the **->**) indicates the outgoing frame. You can configure the outgoing frame type for each incoming frame.

3. You change an outgoing frame type by entering its line number, an equal sign (=) and a frame type indicator (**e**, **s**, **2**, or **3**). The frame type indicators represent the following frames:

```
e      Ethernet II or Ethertype
s      SNAP
2      802.2 or LLC
3      Ethernet 802.3
```

For example, if you wanted to change incoming IPX Ethernet II frames to Ethernet 802.3 frames, then you would enter

```
3=3
```

Please note that the IP Translation Options accept only Ethertype (**e**) or SNAP (**s**).

4. When you are done changing translations, enter **save** to save all of your settings. If you enter **quit**, you will exit the **ethdef** command without saving your changes.

Default FDDI Translations

The **fdiddef** command allows you to set up default translations for all FDDI ports. To do this:

1. Enter the **fdiddef** command at the system prompt. The following screen displays:

```

This will reset the default translations for FDDI media to a new value.
All FDDI ports currently set to default will inherit these new translation options.
It will then reset the forwarding table translation options for all addresses learnt on
those ports to those port defaults.
Do you want to do this? (no):

```

2. Press **Y** at the **Do you wish to do this?** prompt to indicate that you want to change the defaults. The current settings for FDDI ports are displayed, in a screen similar to the following:

```

Translation Options:
1      IP Ethertype          -> Ethertype
2      IP IEEE 802 SNAP      -> Ethertype

3      IPX ETHERNET_II       -> 802.3
4      IPX ETHERNET_802.3    -> 802.3
5      IPX ETHERNET_802.3/FDDI/TOKEN_RING -> 802.3
6      IPX ETHERNET_SNAP/FDDI_SNAP/TOKEN-RING_SNAP -> 802.3

```

There are six frame types for which you can set translation options. The frame type in the left column indicates the incoming frame, and the frame type in the right column (after the **->**) indicates the outgoing frame. You can configure the outgoing frame type for each incoming frame.

3. You change an outgoing frame type by entering its line number, an equal sign (=) and a frame type indicator (**e**, **s**, **2**, or **3**). The frame type indicators represent the following frames:

```

e      Ethernet II or Ethertype
s      SNAP
2      802.2 or LLC
3      Ethernet 802.3

```

For example, if you wanted to translate incoming IPX Ethernet 802.3 frames to Ethernet 802.3 frames (FDDI raw), then you would enter

```
4=3
```

Please note that the IP Translation Options accept only Ethertype (**e**) or SNAP (**s**).

4. When you are done changing translations, enter **save** to save all of your settings. If you enter **quit**, you will exit the **ethdef** command without saving your changes.

◆ Important Note ◆

The IP Translation Options allow only SNAP (**s**). The IPX translations allow SNAP (**s**), and LLC (**2**) for all frame types. The Ethertype (**e**) translation is not allowed for FDDI. The Ethernet 802.3 translation (**3**) is allowed only on incoming Ethernet 802.3 frames, which referred to as “FDDI raw.”

The **fdidef** command will accept your input and will not return an error message if you try to change an IPX translation option to Ethertype or Ethernet 802.3. However, that does not mean that the IPX frames are being translated to Ethertype or 802.3. Regardless of what the **fdidef** screen displays, switch software does not translate FDDI frames to Ethertype for any frame or to 802.3 for any frame except incoming 802.3.

Default Token Ring Translations

The **trdef** command allows you to set up default translations for all Token Ring ports. To do so:

1. Enter the **trdef** command at the system prompt. The following screen displays:

```
This will reset the default translations for TR media to a new value.  
All TR ports currently set to default will inherit these new translation options.  
It will then reset the forwarding table translation options for all addresses learnt on  
those ports to those port defaults.  
Do you want to do this? (no):
```

2. Press **Y** at the **Do you wish to do this?** prompt to indicate that you want to change the defaults. The current settings for FDDI ports are displayed:

```
Translation Options:  
1      IP Ethertype           -> Ethertype  
2      IP IEEE 802 SNAP      -> Ethertype  
  
3      IPX ETHERNET_II       -> 802.3  
4      IPX ETHERNET_802.3    -> 802.3  
5      IPX ETHERNET_802.3/FDDI/TOKEN_RING -> 802.3  
6      IPX ETHERNET_SNAP/FDDI_SNAP/TOKEN-RING_SNAP -> 802.3
```

There are six frame types for which you can set translation options. The frame type in the left column indicates the incoming frame, and the frame type in the right column (after the **->**) indicates the outgoing frame. You can configure the outgoing frame type for each incoming frame.

3. You change an outgoing frame type by entering its line number, an equal sign (=) and a frame type indicator (**e**, **s**, **2**, or **3**). The frame type indicators represent the following frames:

e	Ethernet II or Ethertype
s	SNAP
2	802.2 or LLC
3	Ethernet 802.3

For example, if you wanted to translate incoming IPX SNAP frames to LLC frame, then you would enter

6=2

4. When you are done changing translations, enter **save** to save all of your settings. If you enter **quit**, you will exit the **trdef** command without saving your changes.

◆ **Important Note** ◆

The IP Translation Options allow only SNAP (**s**). The IPX translations allow only SNAP (**s**), and LLC (**2**) for all frame types. The Ethertype (**e**) and 802.3 translations are not allowed for Token Ring.

The **trdef** command will accept your input and will not return an error message if you try to change an IPX translation option to Ethertype or Ethernet 802.3. However, that does not mean that the IPX frames are being translated to Ethertype or 802.3. Regardless of what the **trdef** screen displays, switch software does not translate Token Ring frames to Ethertype or 802.3.

Port Translations

The **swch** command allows you configure translations on a port-by-port basis. Its translation options are similar to those for **ethdef**, **fddidef**, and **trdef**. However, instead of applying translations to all ports for a particular media type, **swch** applies translations only to the port you specify.

To specify translation for a single port:

1. Start the **swch** command by entering it at the prompt as shown:

```
swch <slot>/<port>
```

where **<slot>** is the board on which the port is located and **<port>** is the port number. For example, to set the translation for port 1 on slot 2, enter the following:

```
swch 2/1
```

2. Something like the following screen displays, showing the current translation settings for the port:

```

Port Translations for Ethernet port 2/1/brg/1

0      Framing Type: DFLT

Translation Options:
1      IP Ethertype           -> Ethertype
2      IP IEEE 802 SNAP       -> Ethertype

3      IPX ETHERNET_II        -> 802.3
4      IPX ETHERNET_802.3     -> 802.3
5      IPX ETHERNET_802.2/FDDI/TOKEN_RING -> 802.3
6      IPX ETHERNET_SNAP/FDDI_SNAP/TOKEN-RING_SNAP -> 802.3

```

The top line of the display indicates the media type of the port as well as the slot number, port number, service type, and service number. The next line, **Framing Type**, indicates the framing type applied to this port through the **modvp** command. If the framing type had been defined through the Switch menu, then this field would read **SWCH**.

3. The Translation Options section shows the six frame types for which you can set translation options. The frame type in the left column indicates the incoming frame, and the frame type in the right column (after the **->**) indicates the outgoing frame. You can configure the outgoing frame type for each incoming frame.

Note that the default option is a question mark (?). If you press **<Return>**, the help information will be redisplayed

4. You change an outgoing frame type by entering its line number, an equal sign (=) and a frame type indicator (**e**, **s**, **2**, or **3**). The frame type indicators represent the following frames:

e	Ethernet II or Ethertype
s	SNAP
2	802.2 or LLC
3	Ethernet 802.3

For example, if you wanted to translate incoming IPX SNAP frames to LLC frames, then you would enter

```
6=2
```

5. When are done changing translations, enter **save** to save all your settings. If you enter **quit**, you will exit the **swch** command without saving your changes.

Please note that valid translation options depend on the media type of the port. Ethernet ports allow all frame translation options, but FDDI and Token Ring ports have limitations. See *Default FDDI Translations* on page 18-27 and *Default Token Ring Translations* on page 18-28 for more information on media limitations.

Configuring Additional Ports

If you want to configure additional ports, you can use the **n** option of the **swch** command to configure the next port, or the **p** option of the **swch** command to configure the previous port. For example, if you want to configure translations on port 2 for the card in slot 4 after configuring Port 1 in Slot 4, enter

```
n
```

at the prompt. You are now ready to configure port 3 of slot 4.

If you want to configure translations on port 1 for the card in slot 5 after configuring Port 2 in Slot 5, enter

```
p
```

at the prompt. You are now ready to configure port 1 of slot 5.

When are done changing translations, enter **save** to save all your settings. If you enter **quit**, you will exit the **swch** command without saving your changes.

Displaying Ethernet Switch Statistics

The **swch** command can also be used to display basic statistics for Ethernet ports. These statistics are the lowest level, most primitive statistics maintained by an Ethernet board. The more familiar RMON and MIB II statistics are generated from these statistics. If you want to display the switch statistics for an Ethernet port, enter

```
swch <slot>/<port>
```

where **<slot>** is the slot number of the module, and **<port>** is the number of the port for which you want to view statistics. For example, to look at statistics for port 4 in slot 3, enter:

```
swch 3/4
```

A screen similar to the following is displayed:

```

Port Translations for Ethernet port 3/4/brg/1
0      Framing Type: DFLT
Translation Options:
1      IP Ethertype           -> Ethertype
2      IP IEEE 802 SNAP       -> Ethertype
3      IPX ETHERNET_II        -> 802.3
4      IPX ETHERNET_802.3     -> 802.3
5      IPX ETHERNET_802.2/FDDI/TOKEN_RING -> 802.3
6      IPX ETHERNET_SNAP/FDDI_SNAP/TOKEN-RING_SNAP -> 802.3
```

If this port is an Ethernet media port, enter **r** at the system prompt and then press **<Return>**. If you do this for a port other than an Ethernet port, this will be ignored.

If the port selected is an Ethernet based port, something like the following would be displayed:

```
Ethernet Statistics for Ethernet port 3/4/Brg/1
Received Good Octets      0      Transmitted Good Octets      0
Received Bad Octets      0
Total Octets              0
Received Unicasts        0      Transmitted Unicasts         0
Received Multicasts      0      Transmitted Multicasts       0
Received Broadcasts      0      Transmitted Broadcasts       0
Received Buffer Discards  0      Transmitted Buffer Discards    0
Received Collision Count  0      Transmitted Retry Count       0
Received Runt Count       0      Transmitted More Count        0
Received Error Discard   0      Transmitted Once Count        0
Drop Event Count         0      Transmitted Defer Count       0
Received Jabbers         0      Loss Carrier Count            0
Received Over Size       0      Transmitted Late Collisions   0
Received Late Collision  0      Transmit Underflow            0
Received 1024 +          0      Port Filtered                 0
Received 512 +           0      Vlan Filtered                 0
Received 256 +           0      Mtu Exceeded                  0
Received 128 +           0
Received 65 +            0
Received 64              0
vseTxDiscard             0
```

The fields displayed by the **r** option of the **swch** command are described below:

◆ **Note** ◆

The first group of statistics are the numbers of bytes transmitted and received. These are useful in working out bandwidth usage by the port. Bad octets are important to count in the total octets count as they consume bandwidth at the expense of useful traffic. To ignore them would lead to mysterious loss of bandwidth in any calculations performed.

Received Good Octets. The total number of bytes received in good frames.

Received Bad Octets. The total number of bytes received in bad frames.

Total Octets. The total number of octets transmitted or received in good or bad frames on this port.

Transmitted Good Octets. The total number of bytes successfully transmitted.

Received Unicasts. The number of frames received on this port whose destination address is a unicast format.

Transmitted Unicasts. The number of frames transmitted on this port whose destination address is a unicast format.

Received Multicasts. The number of frames received on this port whose destination address is a multicast format.

Transmitted Multicasts. The number of frames transmitted on this port whose destination address is a multicast format.

Received Broadcasts. The number of frames received on this port whose destination address is the broadcast address.

Transmitted Broadcasts. The number of frames transmitted on this port whose destination address is the broadcast address.

Note that these statistics merely indicate the format of the destination address of frames transmitted/received on this port, not that the addressed device and/or devices necessarily reside on that port. For example, unknown unicast addressed frames are flooded to many ports.

Received Buffer Discards. Due to congestion of traffic from multiple ports on the board, timely access to buffers was not available to receive a frame from the network port and the frame was discarded.

Transmitted Buffer Discards. Due to a shortage of buffers and/or congestion on the network port, frames received from the backplane destined to this port were dropped.

Transmit Underflow. Due to congestion of traffic from multiple ports on the board, timely access to the buffer containing the frame currently being transmitted by this port was not obtained and the frame had to be aborted and discarded.

vseTxDiscard. Due to congestion of traffic from multiple ports and boards in the system, traffic received from the network port could not be queued to the backplane due to buffer availability.

Received Collision Count, Received Runt Count. These counts may be considered normal on a shared segment (e.g., AUI and BNC connected Ethernet) where more than two stations exist. The first indicates that a frame which the port started to receive from a station was subjected to a collision from a third station. This is normal. Such collisions between third party stations may cause this port to see fragments of a frame which are discarded as runts. This too is normal on multiple station Ethernet segments. On point to point 10Base-T connections these events may be considered abnormal indicating a possible intermittent wiring problem (unless hubs which propagate fragments are in use.) These statistics do not indicate the loss of any frame but rather events associated with the attempts to finally successfully transfer the frame.

Transmitted Defer Count, Transmitted Once Count, Transmitted More Count, and Transmitted Retry Count. These statistics are all related to collisions and deferral where this port is actively trying to transmit a frame. The CSMA part of CSMA/CD, the protocol of Ethernet, requires that a station which wishes to transmit first listens to the media to see if a transmission is already in progress. If it is, then the station must defer transmission until the media is quiet. The Defer count is the number of times this happens and is normal. A high defer count, relative to total numbers of frames transmitted by the port, can be indicative of a busy segment. If a transmission is not in progress the station may begin to transmit. Due to propagation delays it is possible for a station to suffer a collision from another station trying to transmit, even though both listened for quiet media. When this occurs, both stations “back off” for a random time before attempting transmission again. In theory, subsequent collisions may occur on these retries. Once, More, and Retry indicate whether this is occurring. If a collision occurs but succeeds on the retry, the Once counter is incremented, i.e., we collided once. If more than one retry is required, the More count is incremented. If up to 16 retries are attempted and all collide, then the frame is dropped and the Retry count is incremented. Again, Once, More, and Retry are normal events on CSMA/CD media but high numbers, relative to total transmitted frames, are again indicative of a very busy segment whose throughput could be increased by further segmentation.

Received Error Discard. A frame was received with an FCS and/or alignment error. A high count here, relative to total received frames, is indicative of a noisy media subject to errors.

Loss Carrier Count. This is a count of transmitted frames which are lost due to a loss of carrier. This is indicative of poor quality/noisy wiring or adapter cards.

Received Late Collision, Transmitted Late Collisions. A late collision is a collision which occurs in a frame when more than 64 bytes have been received/transmitted. On a correctly configured network, which doesn't exceed physical limits of size, impedance, station spacing, etc., stations should always collide within 64 bytes due to propagation times. Late collisions indicate that the network is violating such restrictions or some stations are having a problem which prevents them correctly implementing the CSMA/CD protocol. For example, a station with a faulty receiver can not "hear" transmissions in progress and so may fail to defer its transmissions causing late collisions to be seen by other stations.

Received Jabbers, Received Over Size. The maximum frame size on Ethernet is 1518 bytes. Frames longer than this are illegal. When such a frame has a valid FCS it is counted as over-size. If it has an FCS error then it is counted as a Jabber. The former is indicative of a device with improper software, the latter of a device with some hardware fault on its transmitter. In both cases the faulty station causes other devices, such as this port, to see these errors.

Drop Event Count. When a frame is dropped, for example, frame reception is aborted because of lack of buffers, there may be only one or there may be many frames so affected. In either case there is a single occurrence of an "event" during which frames were lost. This is what this statistic counts. This statistic is used in RMON as follows. For example, at network start up there may be a huge amount of flooded traffic leading to much lost traffic. When a network administrator subsequently looks at the statistics they might see 2 million frames transmitted with 5000 frames lost. At that point they have no clue as to when and why those 5000 frames were lost. If drop event is 5000 it may indicate an intermittent problem where single frames are being lost. If drop event is 5 or 6 it might indicate a few events when large numbers of frames were lost such as in our example, the network restart.

Received 1024 +, Received 512 +, Received 256 +, Received 128 +, Received 65 +, and Received 64. These count the number of frames in the indicated frame sizes: **Received 64** counts 64 byte frames, **Received 65+** counts frames between 65 and 127 inclusive, **Received 128+** counts between 128 and 255, etc. These statistics are only applied to received frames.

◆ **Note** ◆

The **Received 1024 +, Received 512 +, Received 256 +, Received 128 +, Received 65 +, and Received 64** fields will always display zero for Gigabit ports.

Port Filtered. On shared media ports, Station A transmitting to Station B will be directly delivered. Therefore, the frame received by this port just needs to be dropped. This action is referred to as filtering and this counts the number of frames so filtered.

Vlan Filtered. The Omni Switch/Router restricts traffic above the normal Level 2 filtering by applying VLAN rules. Frames which are dropped because of VLAN rules are counted here.

Mtu Exceeded. This statistic is not currently supported and is always zero.

Displaying Token Ring Switch Statistics

In Release 3.4 and later, you can display statistics for the new generation of Token Ring modules known as “Bigfoot” (e.g., TSM-CD-16W, TSX-CD-16W, and TSX-C-32W). For example, if you want to display the switch statistics for a Token Ring port on Port 1 on Slot 4, enter:

```
swch 4/1
```

at the system prompt. Press **r** and then press **<Enter>** at the prompt. Something like the following displays:

```
n={e,s,2,3},quit,save,(?) : r
  Token Ring Statistics for 4/16 Mbit Token Ring port 4/1/Brg/1

Rx MAC Good Bytes           0      Rx LLC Good Bytes           0
Rx Total Mac Packets        0      Rx Total LLC Packets        0
Rx MAC Errored Bytes       0      Rx LLC Errored Bytes       0
Rx Unicast Packets         0      Tx Unicast Packets         0
Rx Multicast Packets       0      Tx Multicast Packets       0
Rx Broadcast Packets       0      Tx Broadcast Packets       0
Rx Buffer Discards         0      Tx Buffer Discards         0
Rx Error Discards          0      Tx Error Discards          0
Ring Purge Events          0      Ring Purge Packets         0
Beacon Events              0      Beacon Packets             0
Claim Token Events         0      Claim Token Packets        0
Internal Errors            0      Line Errors                 0
Burst Errors               0      AC Errors                   0
Abort Errors               0      LostFrame Errors           0
Congestion Errors          0      Frame Copied Errors        0
Frequency Errors           0      Token Errors                0
Soft Errors                 0      Ring Poll Events           0
Internal Errors            0      NAUN Changes               0
Received 18_63 byte Pkts   0      Received 64_127 byte Pkts  0
Received 128_255 byte Pkts 0      Received 256_511 byte Pkts 0
Received 512_1023 byte Pkts 0      Received 1024_2047 byte Pkts 0
Received 2048_4097 byte Pkts 0      Received 4096_8191 byte Pkts 0
Received 8K_18000 byte Pkts 0      Received 18000+ byte Pkts  0
n={e,s,2,3},quit,save,(?) : ?
```

Note that the default option is now **r**. If you press **<Enter>**, the switch statistics will be redisplayed.

The fields displayed by the **r** option of the **swch** command for Token Ring are described below.

The first group of statistics are the numbers of bytes transmitted and received. These are useful in working out bandwidth usage by the port. Bad octets are important to count in the total octets count as they consume bandwidth at the expense of useful traffic. To ignore them would lead to mysterious loss of bandwidth in any calculations performed.

Rx MAC Good Bytes. The total number of bytes received in good Media Access Control (MAC) packets. (MAC packets are used for management of the Token Ring network.)

Rx LLC Good Bytes. The total number of bytes received in good Logical Link Control (LLC) packets. (LLC packets are used to transfer data.)

Rx Total MAC Packets. The total number of bytes received in MAC packets.

Rx Total LLC Packets. The total number of bytes received in LLC packets.

Rx MAC Errored Bytes. The total number of bytes received in bad MAC packets.

Rx LLC Errored Octets. The total number of bytes received in bad LLC packets.

The next group of statistics are the types of packets being transmitted and received.

Rx Unicast Packets. The number of packets received on this port whose destination address is a unicast format.

Tx Unicast Packets. The number of packets transmitted on this port whose destination address is a unicast format.

Rx Multicast Packets. The number of packets received on this port whose destination address is a multicast format.

Tx Multicast Packets. The number of packets transmitted on this port whose destination address is a multicast format.

Rx Broadcast Packets. The number of packets received on this port whose destination address is the broadcast address.

Tx Broadcast Packets. The number of packets transmitted on this port whose destination address is the broadcast address.

Note that these statistics merely indicate the format of the destination address of packets transmitted/received on this port, not that the addressed device and/or devices necessarily reside on that port. For example, unknown unicast addressed packets are flooded to many ports.

The next group of statistics are the buffer resource related statistics. The NI board receives packets from the backplane to be transmitted to the network ports and receives packets from the network ports to be transmitted to the backplane. It requires buffers to store these packets in while being transferred across the board in this manner. Under heavy and congested traffic a shortage of buffers or lack of timely access to these buffers may occur. These statistics count these events which are more indicative of the amount of traffic on the board as opposed to this particular port.

Rx Buffer Discards. Due to congestion of traffic from multiple ports on the board, timely access to buffers was not available to receive a frame from the network port and the frame was discarded.

Tx Buffer Discards. Due to a shortage of buffers and/or congestion on the network port, packets received from the backplane destined to this port were dropped.

The next group are also indicative of network segment health but are indicative of ill health and indicate events where a frame is lost.

Rx Error Discards. The total number of errored packets (bad CRC, code violations, invalid frame length, etc.) received by this port that were discarded.

Tx Error Discards. The total number of errored packets exceeding the maximum frame length (MTU exceeded, FIFO underruns, etc.) by this port that were discarded.

The next group describe events that can occur when stations are inserted or removed from a ring.

Ring Purge Events. The total number of times this port enters the ring purge state from the normal ring state.

Ring Purge Packets. The total number of times that this port enters a beaconing state.

Beacon Events. The total number of beacon packets received and transmitted by this port.

Beacon Packets. The number of beacon MAC packets detected by this port.

Claim Token Events. The total number of times that this port enters the claim token state from

the normal ring state or ring purge state to elect a new active monitor.

Claim Token Packets. The total number of claim packets transmitted by this port.

The next group describe error statistics for token, MAC, and LLC packets.

Internal Errors. The total number of times this port detects a recoverable internal error.

Line Errors. The total number of errors caused by problems with the physical links (code violations, Frame Check Sequence (FCS) errors inside a frame).

Burst Errors. The total number errors when this port detects the absence of transmissions for five (5) half-bit timers (burst-five errors).

AC Errors. The total number of token packets with an invalid Access Control (AC) byte.

Abort Errors. The total number of times that this port detects an abort delimiter while transmitting a packet.

LostFrame Errors. The total number of packets that failed to reach their destination after the token ring rotation timer has expired.

Congestion Errors. The total number of packets lost due to the fact that no buffer was available at the destination station.

Frame Copied Errors. The total number of times that a frame has been incorrectly copied by another station on the ring or copied by a station with a duplicate address.

Frequency Errors. The total number of timing errors frames detected by this port that did not contain a proper ring-clock frequency.

Token Errors. The total number of times this port detects that a new token was generated by the Active Monitor on the ring due to a lost token.

Soft Errors. The total number of recoverable errors detected by this port.

The next group describe statistics for changes in ring topology.

Ring Poll Events. The total number of times that this port has learned its upstream neighbor's address and has broadcasted the inserting adapter's address to the port's downstream neighbor.

Internal Errors. The total number of insertion failures.

NAUN Changes. The number of times that the Nearest Active Upstream Neighbor (NAUN) for this port has changed.

The next set of statistics display information on network traffic. These statistics are only applied to received packets.

Received 18_63 byte Pkts. The total number of packets received on this port that were at least 18 bytes (octets) long and less than or equal to 63 bytes long.

Received 64_127 byte Pkts. The total number of packets received on this port that were at least 64 bytes (octets) long and less than or equal to 127 bytes long.

Received 128_255 byte Pkts. The total number of packets received on this port that were at least 128 bytes (octets) long and less than or equal to 255 bytes long.

Received 256_511 byte Pkts. The total number of packets received on this port that were at least 256 bytes (octets) long and less than or equal to 511 bytes long.

Received 512_1023 byte Pkts. The total number of packets received on this port that were at least 512 bytes (octets) long and less than or equal to 1023 bytes long.

Received 1024_2047 byte Pkts. The total number of packets received on this port that were at least 1024 bytes (octets) long and less than or equal to 2047 bytes long.

Received 2048_4097 byte Pkts. The total number of packets received on this port that were at least 2048 bytes (octets) long and less than or equal to 4095 bytes long. [check]

Received 4096_8191 byte Pkts. The total number of packets received on this port that were at least 4096 bytes (octets) long and less than or equal to 8191 bytes long.

Received 8k_18000 byte Pkts. The total number of packets received on this port that were at least 8192 bytes (octets) long and less than or equal to 18,000 bytes long.

Received 18000+ byte Pkts. The total number of packets received on this port that were more than 18,000 bytes long.

Any to Any MAC Translations

The **swchmac** command allows you to view the current frame translation settings for a given MAC address. Follow these steps:

1. Enter **swchmac** and the following prompt displays:

Enter MAC address ([XYZZ:AABBCC] or return for none :

2. Enter the MAC for which you want to view translations. The following prompt displays:

Is this MAC in Canonical or Non-Canonical (C or N) [C] :

3. Enter if the MAC address you entered is expressed in canonical (**C**) or non-canonical format. The default is canonical. A screen similar to the following displays:

Port Translations for Ethernet port 3/4/brg/1

Translation Options:

IP Ethertype	-> Ethertype
IP IEEE 802 SNAP	-> Ethertype
IPX ETHERNET_II	-> 802.3
IPX ETHERNET_802.3	-> 802.3
IPX ETHERNET_802.2/FDDI/TOKEN_RING	-> 802.3
IPX ETHERNET_SNAP/FDDI_SNAP/TOKEN-RING_SNAP	-> 802.3
Proprietary Token Ring IPX Switching	-> Off

The screen shows how each incoming frame type is translated. The frame type in the left column indicates the incoming frame type, and the frame type in the right column (after the ->) indicates the outgoing frame translation.

Default Autoencapsulation

Autoencapsulation is a technique employed by AutoTracker software to learn the protocol and encapsulation type used by a source MAC address and automatically translate frames bound to that MAC address to the appropriate encapsulation type.

Normally all devices attached to a switch port receive frames translated according to the translation options defined for that port. However, some devices attached to the same port may require different frame formats.

For example, one workstation may support IPX 802.3 frames and another may support IPX SNAP frames. The switch port may be configured to translate incoming IPX 802.3 frames to LLC frames, which would not satisfy either of the workstations. If autoencapsulation is on, then the switch would translate frames for the first workstation to IPX 802.3 and frames for the second workstation to IPX SNAP. The translation setting for the port is overridden for those ports that require a special translation.

Autoencapsulation operates only on learned unicast frames. It does not work for broadcast, multicast, or unlearned unicast frames. For this reason it is recommended only for ports attached to client devices. It is not recommended for ports attached to servers due to high volume of broadcast traffic on such a connection.

In addition, autoencapsulation is not supported for Banyan Vines frames. It operates only on IP and IPX frames.

To turn on autoencapsulation type the following at the prompt:

autoencaps on

To turn off autoencapsulation type the following at the prompt:

autoencaps off

Translational Bridging

Translational Bridging enables internetworking between FDDI, Ethernet, and Token Ring LANs. There is no standard which encompasses this. The Omni Switch/Router's features focus on bridging of frames between media and translating the MAC and LLC headers into the appropriate "native" frame formats. This provides media-independent internetworking.

Learning

For VLAN trunk frames, the switch will learn the source MAC address of the encapsulated frame and associate this with the source MAC address of the originating switch. When a frame arrives, the switch checks to see if the frame has been learned. If so, then the frame will be encapsulated and sent directly to the destination switch. If not, then the switch will learn the association of VLAN, trunk service, virtual port, source, and destination MACs. If the switch has no ports in the VLAN associated with the frame's destination, the frame is dropped.

Translations across Trunks

The Omni Switch/Router sends frames onto the trunk in the same format as the original LAN type. Any required translation is done at the destination switch.

Dissimilar LAN Switching Capabilities

Switching traffic between like media requires no changes to the frame, whereas switching traffic between unlike media requires some level of change to the frame. To fully explain the various changes possible we need to define the portion of the frame where changes could occur.

Media Specific fields and MAC address fields are different for Token Ring, FDDI, and Ethernet. For Token Ring and FDDI, the switch generates MAC addresses in non-canonical format, where Ethernet generates MAC addresses in canonical format. The Omni Switch/Router will perform media translations which means the media specific, source MAC and destination MAC will be changed for each frame which changes media.

The source routing field is optional, and use of this field is driven by endstations who wish to communicate using source routing. The Omni Switch/Router participates in source routing on FDDI and Token Ring interfaces when it is configured as a Source Route Bridge. The Omni Switch/Router will also forward source route frames transparently while performing standard switching of frames on Token Ring and FDDI interfaces as well as when using the virtual ring feature.

The encapsulation type field can be a number of different encapsulations, which really includes the Media Specific fields, source MAC address, and destination MAC address. The choices are Ethernet II, IEEE 802.2 (LLC), SNAP, and Novell 802.3 or FDDI proprietary formats. There are configuration options for Ethernet, FDDI, and Token Ring interfaces. The encapsulation type field may or may not be changed. This decision is made based on the incoming encapsulation type, the user configuration, and the topology that frame is traveling.


The data field is the remainder of the frame which is application dependent. This data field is not changed for switched traffic. Each frame is followed by a CRC.

Below are some examples when translation can occur.

Switching Between Similar LANs

Translations are not performed for switched traffic between similar LANs within one Omni Switch/Router. For example in the diagram below, if Station A on an Ethernet segment wants to talk to Station B on another Ethernet segment, the switched frames are not changed.

This is true for any two media where the originating media and the destination media are of the same type (i.e. Ethernet, FDDI, Token Ring).




Similar LANs

Switching Between Ethernet LANs Across a Trunked Backbone

Frames that are switched between like media across a Trunked backbone will only be translated at the egress port of the egress Omni Switch/Router. For example in the figure below, frames switched from Station A to Station B will be translated at point 4, where point 4 is the egress port of Switch 2. Frames switched from Station B to Station A will be translated only at point 1, where point 1 is the egress port of Switch 1.


This is true if the originating media and destination media are Ethernet. It is not true if the originating media and destination media are either Token Ring or FDDI.



Ethernet LANs Across a Trunked Backbone

Switching Between Similar LANs across a Native Backbone

Switched traffic between similar LANs across a non-trunked or native backbone will have translations performed at each egress point. In the figure below, for traffic originating from Station A destined to Station B, point 1 represents the ingress (input) port of Switch 1. Likewise, point 2 represents the egress (output) port of Switch 1, point 3 represents the ingress (input) port of Switch 2 and the point 4 represents the egress (output) port of Switch 2. Translations will occur at each egress port. For traffic from Station A to Station B, output translations will occur at points 2 and 4. For traffic from Station B to Station A, output translations will occur at points 3, and 1.



Similar LANs Across a Native Backbone

In the above example, the backbone could be of any media type other than Ethernet. If all three media types were Ethernet, then no translations would occur, because the traffic is being switched from like media to like media.

The following table shows interoperability between dissimilar LANs with two switches where the client and server are resident on like media types and the connection is switched over various LAN backbone types. This table is representative of the IP and IPX protocol only.

	Backbone			
	<i>Token Ring</i>	<i>FDDI</i>	<i>Ethernet</i>	<i>ATM</i>
<i>Token Ring to Token Ring</i>	No	Yes	Yes	No
<i>FDDI to FDDI</i>	Yes	No	Yes	No
<i>Ethernet to Ethernet</i>	Yes	Yes	No	No

Dissimilar LANs

19 Managing Groups and Ports

In a traditional hub-based network, a broadcast domain is confined to a single network interface, such as Ethernet, or even a specific physical location, such as a department or building floor. In a switch-based network, such as one comprised on Omni Switch/Routers, (OmniS/Rs) a broadcast domain—or *Group*— can span multiple physical switches and can include ports using multiple network interfaces. For example, a single OmniS/R Group could span three different switches located in different buildings and include Ethernet and WAN physical ports.

An unconfigured Omni Switch/Router contains one Group, or broadcast domain. It also contains one default Virtual Network, or VLAN, referred to as “default VLAN #1”. The default Group, Group #1, and its default VLAN contain all physical ports in the switch. When a switching module is added to the switch all of these additional physical ports are also assigned to Group #1, VLAN #1.

You can create Groups in addition to this default Group. When you add a new Group, you give it a name and number, optionally configure a virtual router port for its default VLAN, and then add switch ports to it. The switch ports you add to a new Group are moved from the default Group #1 to this new Group. (For more information on how ports are assigned to Groups, see *How Ports Are Assigned to Groups* on page 19-2.)

Up to 500 Groups can be configured on each OmniS/R. An entire OmniS/R network can contain up to 65,535 Groups. Each Group is treated as a separate entity.

There are three main types of Groups:

- 1. Mobile Groups.** These groups allow ports to be dynamically assigned to the Group based on AutoTracker polices. In contrast to non-mobile Groups, AutoTracker rules are assigned directly to a mobile Group. No AutoTracker VLANs are contained within a mobile Group. (However, mobile groups do contain a default VLAN 1 to which AutoTracker policies are assigned; policies assigned to this default VLAN apply to the entire mobile group.) Any AutoTracker policy may be used as criteria for membership in a mobile Group. Mobile groups are described in more detail in *Mobile Groups* on page 19-5.
- 2. Mobile Groups based on authentication.** Authenticated Groups are a special form of mobile Group. These Groups include devices that are dynamically assigned based on an authentication criteria. Typically the user will have to log in with a valid password before being included in an authenticated mobile Group. Group membership is based on users proving their identity rather than the physical location of user devices. Authenticated Groups are described in more detail in the *Switch Network Services User Manual*.
- 3. Non-mobile Groups.** These Groups are the original Group type used in previous releases. They contain statically assigned ports and may contain AutoTracker or Multicast VLANs. These VLANs within a non-mobile Group use AutoTracker policies to filter traffic. AutoTracker rules are not assigned to non-mobile Groups, they are assigned to the VLANs within the Group. Non-mobile groups are described in more detail in *Non-Mobile Groups and AutoTracker VLANs* on page 19-15.

All three types of Groups may co-exist on the same switch. However, a switch port cannot belong to a non-mobile group and a mobile group.

How Ports Are Assigned to Groups

There are two methods for assigning physical OmniS/R ports to a Group. One method is static and requires manual configuration by the network administrator; the other method is dynamic and requires only the configuration of AutoTracker rules for port assignment to occur. The two methods are described in this section.

Static Port Assignment

In the static method, the network administrator manually assigns a port to a Group through the **crgp** and **addvp** commands. The static method can be restrictive because it limits the mobility of users in a multi-Group network. Users can only move within their assigned Group. In addition, customized access for individual users is limited by this method. You can use the static method of port assignment with mobile and non-mobile groups. Static port assignment can be combined with dynamic port assignment for mobile groups, while static port assignment is the only method for assigning ports to non-mobile groups.

Dynamic Port Assignment (Group Mobility)

The dynamic method is available with the Group Mobility feature. Initially each port is part of the default Group #1 (only ports in the default Group and ports in mobile Groups are candidates for dynamic port assignment). Based on the nature of traffic and configured AutoTracker policies, ports are dynamically assigned to the appropriate Group.

For example, if a device attached to a port transmits traffic from the 140.0.0.0 subnet, AutoTracker will check to see if a policy exists for this IP address. If it does, then it will move the port from the default Group to the first Group using this policy. If this device detaches from the network the port will be re-assigned to a Group without intervention by the network administrator.

A port can belong to multiple mobile groups (up to 16) as long as devices attached to that port match policies of these mobile groups. However, an individual device, or MAC address, can only belong to one mobile group per protocol.

The dynamic method of port-to-Group assignment still requires the creation of Groups through the **crgp** command. The criteria for the dynamic assignment of ports to a Group are determined by AutoTracker policies that you can configure during the **crgp** procedure.

Only Ethernet ports can be dynamically assigned to Groups.

If more than one Group has the same type of rule, then ports matching that policy will be assigned to the first Group matching the policy. For example, if a device matched policies in both Groups 2 and 5, the port would be assigned to Group 2. To make the most out of Group Mobility it is best not to duplicate policies among Groups.

Configuring Dynamic Port Assignment

You can enable dynamic port assignment while creating a group through the **crgp** command. During the **crgp** procedure, you will be prompted

Enable Group Mobility on the Group ? [y/n] (n):

Answer **Yes** to this question to give this Group the capability of having ports and devices dynamically added to the Group. Port and devices will be dynamically assigned based on AutoTracker rules you define.

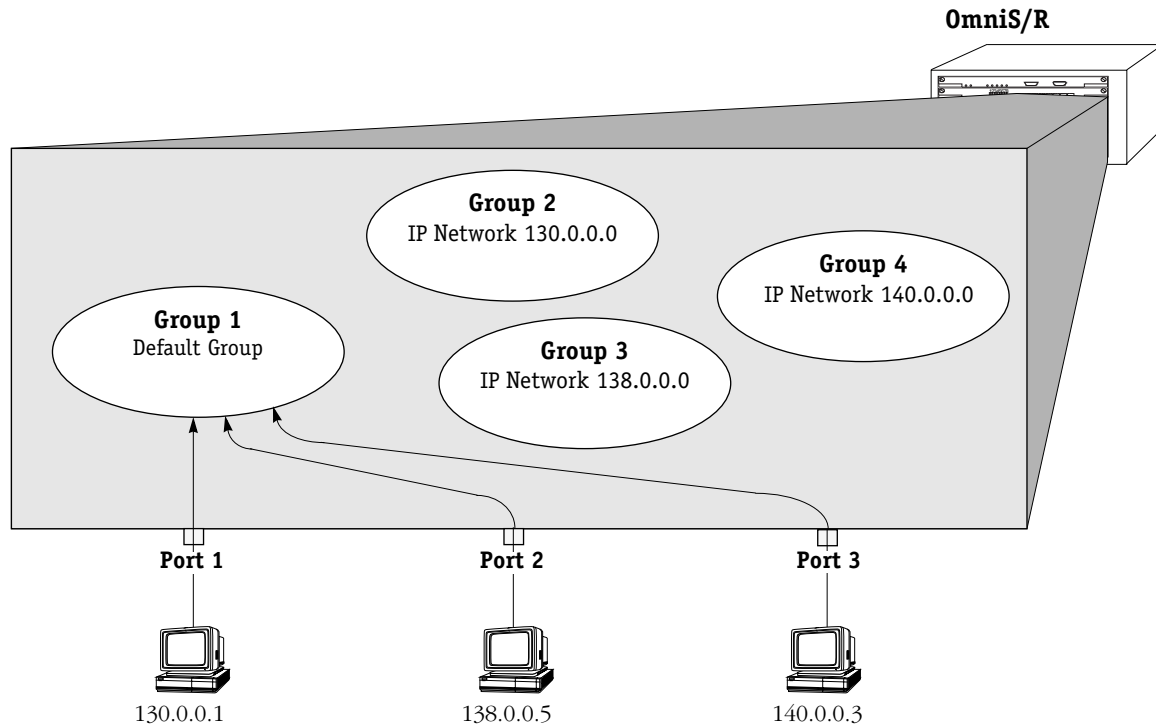
Service Ports and Group Mobility

These ports may be automatically added to the mobile group during the **crgp** procedure or through the **cats** command.

How Dynamic Port Assignment Works

Initially each port is assigned to the default Group. In this example, all three ports have workstations that belong to three different IP subnets (130.0.0.0, 138.0.0.0, and 140.0.0.0). All three ports start out in the default Group.

Group Mobility examines traffic coming from OmniS/R ports. Three mobile groups are defined on the switch and each uses a different IP policy. Traffic that matches IP policies for a Group will trigger the movement of the port to the matching Group.




Initial Configuration: All Ports in Default Group

As soon as the workstations start transmitting traffic, Group Mobility checks the source subnet of the frames and looks for a match with any configured IP policies. If a match is found—and in this example all three ports can be matched with a corresponding Group—the port is moved to the matching Group.

Devices matching a policy trigger the assignment of a port to a mobile group. Therefore, the device is moved to the mobile group at the same time as the port to which it is attached. If more than one device comes in on a port, then that port can belong to more than one mobile group. Similarly, if a device transmits more than one protocol—such as IP and IPX—then the port to which it is attached can belong to more than one mobile group.

How Ports Are Assigned to Groups

As the illustration below shows, the three ports are each moved from the default Group to a Group with a policy that matches the subnet address of the workstation attached to the port. AutoTracker IP address policies have been set up in Groups 2, 3, and 4. The ports are moved to the Group with policies matching the subnet of the workstation.



Ports Move to Groups With Matching Policies

Mobile Groups

Switch ports can be dynamically assigned to mobile groups through AutoTracker policies. Support for dynamic port assignment is one of the main differences between mobile groups and non-mobile groups. AutoTracker rules are assigned *directly* to a mobile group. In contrast, AutoTracker rules are assigned to the VLANs *within* a non-mobile group. No AutoTracker VLANs are contained within a mobile Group, and each mobile group constitutes a single spanning tree.

A switch port can belong to multiple mobile groups, whereas a switch port can belong to only one non-mobile group. However, a port can *not* belong to a mobile and a non-mobile group at the same time.

Ports can be assigned to mobile groups either statically or dynamically. A port is *statically* assigned to a mobile group when one of the following occurs:

- Port by default assigned to default group 1
- Port assigned to a group through **crgrp** or **addvp** commands

Although switch ports can belong to multiple mobile groups, it is not possible to assign a port to two different groups using the **addvp** command. However, a switch port could be assigned to one mobile group via the **addvp** command and then gain membership to another mobile group by matching the policy criteria for that group.

A switch port is *dynamically* assigned to a mobile group after one of its attached devices matches an AutoTracker policy for that mobile group. An overview of how ports and devices are dynamically assigned to mobile Groups can be found in *How Ports Are Assigned to Groups* on page 19-2.

Authenticated Groups

Mobile groups provide the added flexibility of user-authentication policies. Using Authentication Management Console (AMC) software, you can configure mobile groups to use log-in procedures as a means of assigning group membership. Mobile groups that use authentication are a special group type called an Authenticated Group. Authenticated Groups are described in more detail in the *Switch Network Services User Manual*.

Configuring Mobile Groups

You configure mobile Groups through the **crgrp** command. During the **crgrp** procedure you will receive a prompt asking if you want to create a mobile Group

Enable Group Mobility on this Group ? [y/n] (n):

You must answer **Yes** to this prompt to set up a mobile group. After this question, you will be asked to configure virtual ports and AutoTracker policies for the Group. Documentation for the full **crgrp** procedure can be found in *Creating a New Group* on page 19-18.

Turning Group Mobility On or Off

The **gmstat** command turns group mobility on or off for a Group that you specify. Essentially, you can change a non-mobile group into a mobile group and a mobile group back into a non-mobile group through **gmstat**. The group you specify must previously have been created through the **crgp** command.

Use the following syntax for the gmstat command:

```
gmstat <group number>
```

For example, if you wanted to change the group mobility status of group 2, you would enter:

```
gmstat 2
```

Mobile Group to Non-Mobile Group

If this group is already a mobile group, the following would display:

```
Group Mobility is ON for Group 2  
Change Group Mobility Status for Group 2 to OFF ? [y/n] (y):
```

If you wanted to change this mobile group back to a non-mobile group, you would press **<enter>** and the group would lose its mobile status. All AutoTracker policies you set up for the Group would no longer be valid.

If you decided not to turn off group mobility, enter **n** and the following prompt displays:

```
Group Mobility Status unchanged
```

Non-Mobile Group to Mobile Group

If this group is currently a non-mobile group, the following would display:

```
Group Mobility is OFF for Group 8  
Change Group Mobility Status for Group 8 to ON ? [y/n] (y):
```

If you wanted to turn on Group Mobility, you would press **<enter>** and would then be asked if you want to configure AutoTracker policies. If you answer yes, then the AutoTracker policies menu would display as follows:

```
Select rule type:  
1. Port Rule  
2. MAC Address Rule  
    21) MAC Address Range Rule  
3. Protocol Rule  
4. Network Address Rule  
5. User Defined Rule  
6. Binding Rule  
7. DHCP PORT Rule  
8. DHCP MAC Rule  
    81) DHCP MAC Range Rule
```

```
Enter rule type (1):
```

You define policies for a mobile Group. Non-mobile groups do not require policies. However, mobile Groups use policies to define membership. Instructions for specifying AutoTracker policies may be found in Chapter 22.

◆ Note ◆

As of the current release, the MAC Address Range Rule and DHCP MAC Range are not supported for AutoTracker VLANs

If you decided not to turn group mobility on, you would enter **n** at the group mobility prompt and the following message would display:

Group Mobility Status unchanged

Understanding Port Membership in Mobile Groups

Switch ports can belong to multiple mobile groups. A port becomes a member of a mobile group as long as one of its attached devices matches the policy criteria for that group. However, the movement of ports between groups and the status of port membership in groups can be affected by more than just whether or not devices match policy criteria.

Group mobility uses three variables that can affect a port's default group and whether or not a port ages out of a group. These variables are as follows: `def_group`, `move_from_def`, and `move_to_def`. The `def_group` and `move_to_def` variables can be configured through the **gmcfg** command, which is described on page 19-12. The `move_from_def` variable is enabled by default, but can be disabled by entering a statement in the **mpx.cmd** file. The effects of these three variables are described through diagrams on the following pages.

From the perspective of a device or switch port, there are three types of mobile group—default, primary, and secondary. Keep in mind that definitions of these three types are relative and can change for each port and device depending on the settings of the group mobility variables and traffic patterns of devices.

Default Group

The default group is the group a port or device is statically assigned to by “default.” Typically, a port's default group will be Group 1. A port can also be statically assigned to its default group through the **crgp** or **advp** commands. A port or device does not have to match a policy to gain membership into its default group.

The default group for a port or device is stored in memory; it can only be manually changed through the **advp** or **crgp** commands. Depending on the settings of other group mobility variables a device or port can age out of other mobile groups but still remain a member of its default group.

Primary Group

The primary group is the group upon which Spanning Tree operations converge. The primary group is similar to the default group. There are two main differences between a primary and a default group.

1. A primary group only contains devices that have matched one of its AutoTracker policies. In contrast, switch ports may end up in a default group without matching any policy.
2. It is possible for the primary group of a port or device to change through learning or aging. For example, if the `move_from_def` variable is enabled and a device matches the policies of a mobile group other than its default group, then this new mobile group becomes the primary group for the device and the port to which the device is attached (see diagram on page 19-10). In this case the default group and primary group will be different.


If the `move_from_def` is disabled, the port always remains in the default group (which can now also be the primary group).

In addition a port can age out of its primary group if the `move_to_def` variable is enabled (see diagram on page 19-11). A port cannot age out of its default group.

Secondary Group


Switch ports and devices may become members of multiple mobile groups. A switch port starts in its default group, which initially is also its primary group. The primary group may change if the `move_from_def` variable is enabled. Any subsequent mobile groups to which a port gains membership beyond the primary group are “secondary” mobile groups. A port can age out of these secondary groups if the `move_to_def` variable is enabled (see diagram on page 19-11).

How a Device Is Dropped from the Default Mobile Group (def_group)



Device sends traffic that is forwarded to the MPX for processing. If the traffic matches the policies of an existing mobile group, then it will become a member of that group. If the device does not match the policies of any mobile group, then the `def_group` variable determines whether that device becomes a member of the default group.

If `def_group` is enabled....




The device that does not match any policies becomes a member of the default group.

Why enable `def_group`?

- Ensure that all network devices will be a member of at least one mobile group.

If `def_group` is disabled....




All traffic from the device that does not match any policies is dropped. The device is not a member of any mobile group, including the default mobile group.


Why disable `move_from_def`?

- Reduces traffic to and from devices that do not satisfy any network policies.

How a Port's Primary Mobile Group Changes (move_from_def)



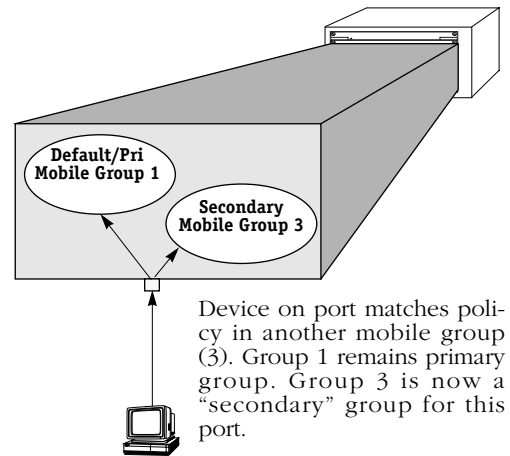
If move_from_def is enabled....



Helpful Hints:

- Reduces broadcasts to the default group.
- Best used when only one device is attached to each port.

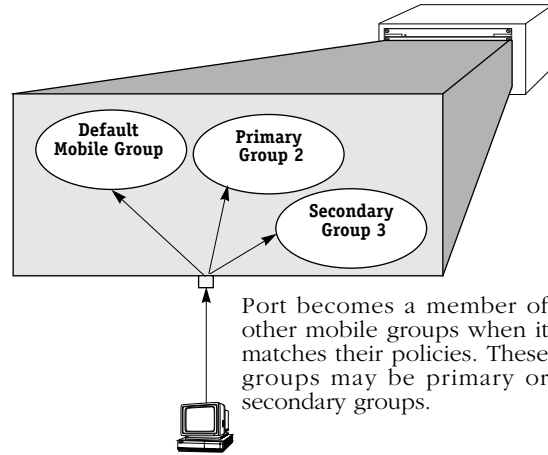

If move_from_def is disabled....




Why disable move_from_def?

- When multiple devices are attached to the switch port, the port must support multiple traffic in the default group as well as traffic in the secondary mobile groups.

How a Port Ages Out of a Mobile Group (move_to_def)




If move_to_def is enabled....



Why enable move_to_def?

- Security. Mobile groups only contain devices and ports that have recently matched policy criteria.

If move_to_def is disabled....



Why disable move_to_def?

- Switch ports retain group membership even when idle for some time. May be appropriate for silent devices, such as printers.

If the port is in “optimized mode,” then the MAC does not age out and the port would stay in the mobile group even if move_to_def is enabled.

Configuring Switch-Wide Group Mobility Variables

There are several switch-wide group mobility variables that you can configure through the **gmcfg** command. These variables control the status of group mobility on all groups in a switch as well as the use of the default group. These variables are illustrated through diagrams on pages 19-9 to 19-11.

Follow these steps to use the **gmcfg** command:

1. Enter **gmcfg**. You do not need to specify a group number as this command applies to all mobile groups in this switch.
2. The following prompt displays:

Group Mobility is Enabled. Disable Group Mobility ? [yes/no] (no) :

This prompt controls the status of group mobility in this switch. If you disable group mobility here then mobile groups will not be supported in this switch even if they are configured through the **crgp** command.

Default Group 1. When group mobility is enabled, default group 1 in the switch will be treated as a mobile group and you will not be able to create AutoTracker VLANs within this group. When group mobility is disabled, default Group 1 in the switch will be treated as a non-mobile group in which AutoTracker VLANs could be created.

The default is to turn Group Mobility off. If you want to enable group mobility, then you need to indicate that choice at this prompt. The prompt will always show the current status of Group Mobility and then ask if you want to change that status. If you want to change the current status, then enter a **y** at this prompt and press **<enter>**. To keep the current status, simply press **<enter>**.

3. The following prompt displays:

move_to_def is set to Disabled. Set to Enable ? [yes/no] (no) :

The **move_to_def** variable determines what happens to a port once the devices on that port age out of the filtering database. By default this variable is Disabled, which means that a port will remain a member of a mobile group as long as its attached device satisfied the criteria for membership in that mobile group at one point. If devices on a port stop transmitting, the port will still retain all its mobile group memberships.

If the **move_to_def** variable is Enabled, then a port will lose its membership in a mobile group if its devices age out of the filtering database for that mobile group (i.e., they stop transmitting traffic that satisfies the criteria for membership in the mobile group). Once a port loses membership in all criteria-based mobile groups, it will return to its default group. The effect of this variable is illustrated on page 19-11.

By default, the **move_to_def** variable is Disabled. If you want to enable it (ports lose mobile group membership when they age out), then you need to indicate that choice at this prompt. The prompt will always show the current status of **move_to_def** and then ask if you want to change that status. If you want to change the current status, then enter a **y** at this prompt and press **<enter>**. To keep the current status, simply press **<enter>**.

4. The following prompt displays:

def_group is set to Enable. Set to Disable ? [yes/no] (no) :

The **def_group** variable determines what happens to devices that do not match any mobile group policies. If **def_group** is Enabled (the default), then devices that do not match any mobile group policies will be part of the default group for that port. If the **def_group** variable is Disabled, then devices that do not match any mobile group policies will be dropped from their default group and will not be part of any mobile group.

By default the `def_group` variable is Enabled. If you want to disable it (devices that do not meet criteria for mobile group membership will not be part of any mobile group), then you need to indicate that choice at this prompt. The prompt will always show the current status of `def_group` and then ask if you want to change that status. If you want to change the current status, then enter a **y** at this prompt and press **<enter>**. To keep the current status, simply press **<enter>**.

The `move_from_def` Variable

The `move_from_def` variable controls whether or not a port's primary group can differ from the port's default mobile group. This variable is enabled by default, but can be changed to disabled in the `mpx.cmd` file.

The original default group for a port is group 1 or the group to which the port is assigned through the `crgp` or `addvp` commands. The primary group at this point is the same as the default group. However, if the `move_from_def` variable is enabled, the primary group can change as soon as a device on the port matches the policy criteria for another mobile group.

For example, Port 5 may start out in Group 1, its default group. The primary group in this case will also be Group 1. If the `move_from_def` variable is enabled and Port 5 matches AutoTracker policies for mobile group 3, then the new primary group for Port 5 will be Group 3. All further Spanning Tree operations for the port will converge on group 3 rather than group 1. The effects of the `move_from_def` variable are further illustrated through diagrams on page 19-10.

If you disable the `move_from_def` variable, then the primary group for a port will always match the default group regardless of the number of other mobile groups to which it gains membership. To disable the `move_from_def` variable, enter the following statement in the `mpx.cmd` file

```
move_from_def=0
```

For this new setting to take place you need to reboot the switch.

Viewing Ports in a Mobile Group

The **vpl** command lists all the Groups in the switch currently configured as mobile Groups and the ports currently assigned to those Groups. Since ports are assigned to mobile groups dynamically, this display is helpful to find out which ports the switch already sees in each group. Ports will only display in this screen for secondary groups (i.e., not default or primary groups). Enter **vpl** and a screen similar to the following displays:

```

=====
Group ID      Physical Port      Virtual Port
=====
Group ID: 2   4/2 4/3 4/4 4/5   12 13 14 15
Group ID: 3   3/1 5/2           8 20
Group ID: 6   NULL Port List
Group ID: 8   4/1 5/1           11 19

```

Group ID. The group number assigned to this mobile group during the **crgrp** procedure.

Physical Port. The physical switch ports that have been dynamically assigned to this group because they matched an AutoTracker policy. (Primary groups do not display in this screen. For a display of port-to-primary group mappings, use the **vi** command) If this column reads **NULL Port List**, then no physical ports have been assigned to the group yet.

Virtual Port. The virtual ports that are part of this mobile group. For Ethernet switch ports, there is a one-to-one relationship between physical and virtual ports.

Viewing a Port's Mobile Group Affiliations

The **vigl** command lists all the ports in the switch that have been assigned to mobile Groups. It is similar to the **vpl** command, but it lists ports first and then Groups. Since ports are assigned to mobile groups dynamically, this display is helpful to find out which ports the switch already sees in each group. Ports will only display in this screen for secondary groups (i.e., not default or primary groups). Enter **vigl** and a screen similar to the following displays:

```

=====
Virtual Port  Physical Port      Group ID
=====
12 13 14 15  4/2 4/3 4/4 4/5   Group ID: 2
8 20         3/1 5/2           Group ID: 3
NULL Port List
11 19        Physical Port      Group ID

```

Virtual Port. The virtual ports in this mobile group. For Ethernet switch ports, there is a one-to-one relationship between physical and virtual ports.

Physical Port. The physical switch ports that have been dynamically assigned to this secondary mobile group because they matched an AutoTracker policy. (Primary groups do not display in this screen. For a display of port-to-primary group mappings, use the **vi** command) If this column reads **NULL Port List**, then no physical ports have been assigned to the group yet.

Group ID. The group number assigned to this mobile group during the **crgrp** procedure.

Non-Mobile Groups and AutoTracker VLANs

Non-mobile Groups are comprised of *physical* entities—switch ports. Groups can span multiple switches, but they are still made up of physical ports that you can see and touch. But just as physically-based broadcast domains are limited, entirely port-based Groups can also be limiting. In a large, flat, switched network, broadcast traffic can overload the network. There needs to be a method for subdividing traffic even further. That's where virtual networks, or VLANs, come into play.

VLANs are created within a Group to subdivide network traffic based on specific criteria. The criteria you use to define a VLAN are called AutoTracker™ policies. AutoTracker policies can be defined by port, MAC address, protocol, network address, a user-defined policy, or a multi-cast policy. VLANs are described in more detail in Chapter 22, “Managing AutoTracker VLANs” and Chapter 23, “Multicast VLANs.”

Routing in a Non-Mobile Group

Communication within a Group containing only the default VLAN is switched; the ports are in the same broadcast domain and do not require routing to communicate. Communication between VLANs in the same Group or to VLANs in other Groups requires routing. That's why all VLANs—including the default VLAN within each Group—may contain their own virtual router port. A virtual router port for each VLAN can be configured to support IP and/or IPX routing. If you do not configure a virtual router port for a VLAN, the devices in that VLAN will not be able to communicate with devices in other VLANs unless there is an external router between the VLANs.

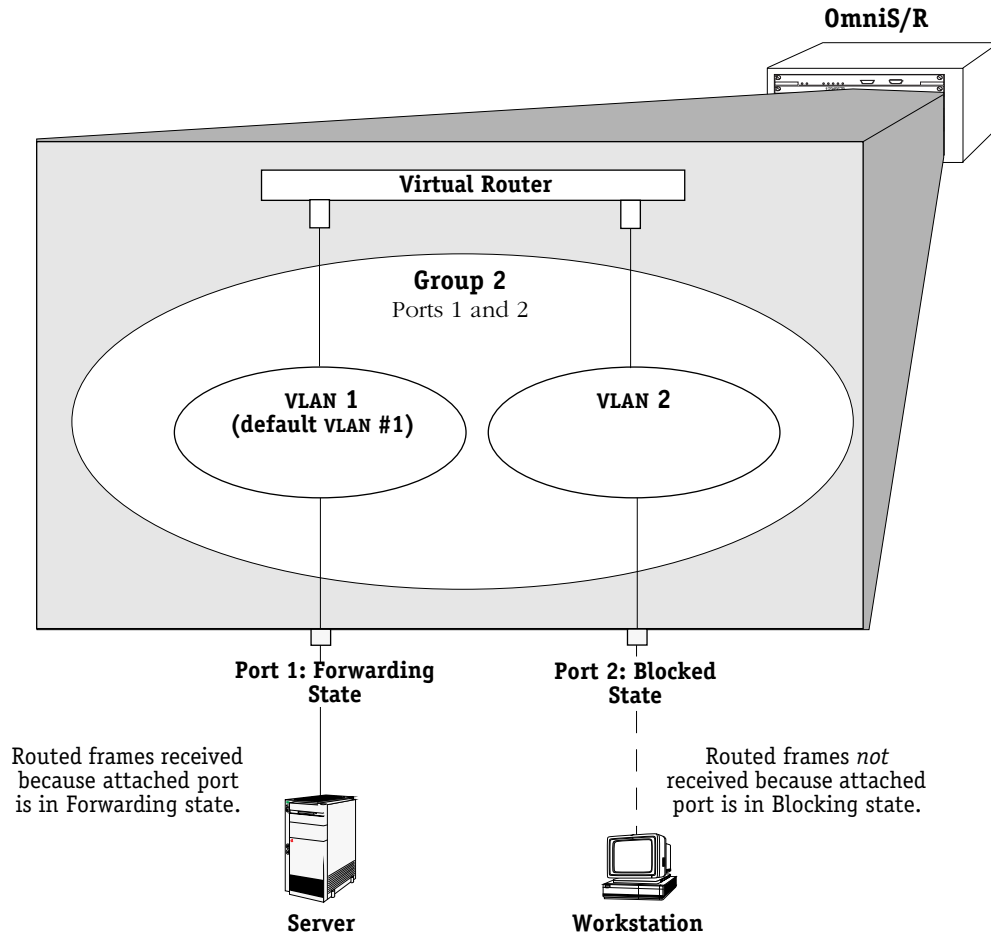
Each OmniS/R supports up to 32 virtual router ports. A single router port, using one MAC address, can support IP routing, IPX routing, or both types of routing. When you enable a router port for a default VLAN, you are actually creating a static route to that VLAN. Routing is covered in more detail in Chapters 25 and 27.

◆ Note ◆

For mobile, non-mobile groups and AutoTracker VLANs, the router port operational status is not active unless an active switch port is a member of the group or VLAN.

Spanning Tree and Non-Mobile Groups

Each Group uses one Spanning Tree for bridging. The OmniS/R supports both 802.1d and IBM Spanning Tree protocols. The Spanning Tree state for the port is Forwarding. Ports that are in Blocked state, or in another non-Forwarding state, will not receive frames from the router port. The figure below illustrates this concept.



Spanning Tree State and Routed Frames

Group and Port Software Commands

Group and Virtual Port commands are part of the VLAN menu within the User Interface. Entering **vlan** at any prompt displays the following menu:

<u>Command</u>	<u>VLAN Management Menu</u>
gp	View the list of Groups currently defined
crgp	Create a Group
modvl	Modify a VLANs configuration/availability
rmgp	Remove a Group
addqgp	Add 802.1q group/s to a port
delqgp	Delete 802.1q group/s from a port
viqgp	Display 802.1q groups on port/s
via	View ports assigned to the selected Group
vi	View info on a specific virtual port
vs	View statistics on a virtual port attachment
ve	View errors on a virtual port attachment
addvp	Add ports to a GROUP
modvp	Modify existing VPORT configuration information
rmvp	Remove ports from a Group
pmapcr	Create a Port Map
pmapdel	Delete a Port Map
pmapmod	Modify a Port Map
pmapv	View Port Mapping Configuration
br	Enter the Bridge Configuration/Parameter sub-menu
prty_mod	Modify the priority of a group
prty_disp	Display the priority of a group
at	Enter the AutoTracker sub-menu

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The VLAN menu commands are divided into four sets of commands. The first set, at the top of the menu beginning with **gp**, contains commands that create, modify, delete, and view Groups. The second set of commands, beginning with **addqgp** are obsolete and no longer control 802.1Q implementation. (See Chapter 16 for information on 802.1Q.) The third set, beginning with **addvp**, contains commands for adding, modifying, and deleting virtual ports. All of these commands are described in this chapter.

The final set of commands at the bottom of the menu, **br** and **at**, are actually entry points to the Bridging and AutoTracker submenus, respectively. Commands for the Bridge Management (**br**) sub-menu are documented in Chapter 17, “Configuring Bridging Parameters.” Commands for the AutoTracker (**at**) sub-menu are documented in this chapter and in Chapter 22, “Managing AutoTracker VLANs” and Chapter 23, “Multicast VLANs.” Some commands in the **at** sub-menu apply to mobile groups and authenticated groups; those commands are described in this chapter.

The **pmapcr**, **pmapdel**, **pmapmod**, and **pmapv** commands allow you to create port mapping configurations. The port mapping feature is documented in *Port Mapping* on page 19-66. The **prty_mod** and **prty_disp** commands allow you to modify and view the priority of a selected group. These commands are detailed in *Priority VLANs* on page 19-73.

Creating a New Group

There are several steps involved in creating a new Group. Note that some steps apply only to mobile groups. These steps are as follows:

1. Enter Basic Group Information, such as the Group number and type. This section starts on page 19-19.
2. Configure the Virtual Router Port (Optional). This section starts on page 19-21.
3. Enable/disable Group Mobility and User Authentication. This section starts on page 19-27.
4. Configure Virtual Ports. This section starts on page 19-28.
5. Configure AutoTracker policies (for mobile groups only). This section starts on page 19-34.

WAN Routing Groups follow a slightly different procedure for their creation. You will receive prompts during the procedure asking whether you want to create one of these special Groups.

Step 1. Entering Basic Group Information

- a. Type **crgp** at any prompt.
- b. The following prompt displays:

GROUP Number (5):

By default the Group number you entered or the next available Group number is displayed in parentheses. Enter the Group number or accept the number shown in parentheses. Each Group must have a unique number, which may range from 2 to 65,535. (Group 1 is the default switch Group. It does not need to be created and it cannot be deleted.) Press **<Enter>** after entering the Group number.

- c. The following prompt displays:

Description (no quotes) :

Enter a descriptive name for the new Group. Group names can consist of up to 30 alphanumeric characters. Press **<Enter>** after entering the Group name.

- d. The following prompt displays:

Enable WAN Routing? (n):

If you want to perform WAN Routing through this Group you must enter a **y** at this prompt. If you do not need to support WAN Routing, then answer **n** at this prompt and continue with Step e.

◆ **Note** ◆

You do not need to create a special WAN Routing Group to bridge or trunk traffic over a WAN connection. If you are just Bridging or Trunking on WAN, answer **n** to this prompt and continue with Step e.

A WAN Routing Group is different from other Groups; it must contain only WAN ports. In addition, the virtual router and virtual ports are configured differently. Please skip ahead to *Creating a WAN Routing Group* on page 19-35 to continue setting up this WAN Routing Group.

- e. The following prompt displays:

Enable ATM CIP? (n):

Answer **n** at this prompt and skip ahead to *Step 2. Configuring the Virtual Router Port (Optional)* on page 19-21.

◆ **Note** ◆

ATM is not supported in Release 4.5 and later.

- f. The following prompt displays:

Enable MPLS? (n):

Multi-Protocol Label Switching (MPLS) must be enabled if this group is going to be used for machines in the network that communicate via MPLS. Answer **n** at this prompt and skip ahead to *Step 2. Configuring the Virtual Router Port (Optional)* on page 19-21.

◆ **Note** ◆

MPLS is not supported in Release 4.5 and later.

Step 2. Configuring the Virtual Router Port (Optional)

You can now optionally configure the virtual router port that the default VLAN in this Group will use to communicate with other VLANs. When you define a virtual router, a virtual router port for the default VLAN in the Group is created. If you do not define a virtual router, no virtual router port is created and the default VLAN in the new Group will be “firewalled,” unable to communicate with other VLANs.

◆ Important Note ◆

Use caution when setting up routing on the default VLAN for a Group. In some configurations enabling routing on the default VLAN may not be necessary or desirable. You can always enable routing on other, non-default VLANs, within this Group. Refer to *AutoTracker Application Example 4* in Chapter 24 for more information.

You will have the choice of configuring IP, IPX, or both IP and IPX routing. Continue with the steps below:

- a. After answering **n** to the **Enable ATM CIP?** prompt, the following prompt displays:

Enable IP (y):

Press **<Enter>** if you want to enable IP Routing on this virtual router port. If you do not enable IP, then the default VLAN in this Group will not be able to route IP data. If you don't want to set up an IP router, enter **n**, press **<Enter>** and skip to Step j.

◆ Note ◆

You may enable routing of both IP and IPX traffic on this router port. If you set up dual-protocol routing, you must fill out information for both IP and IPX parameters.

- b. The following prompt displays:

IP Address:

Enter the IP address for this virtual router port in dotted decimal notation (e.g., 198.206.181.10). This IP address is assigned to the virtual router port of the default VLAN within this Group. After you enter the address, press **<Enter>**.

- c. The following prompt displays:

IP Subnet Mask (0xfffff00):

The default IP subnet mask (in parentheses) is automatically derived from the default VLAN IP address class. Press **<Enter>** to select the default subnet mask or enter a new subnet mask in dotted decimal notation or hexadecimal notation and press **<Enter>**.

- d. The following prompt displays:

IP Broadcast Address (198.200.10.255):

The default IP broadcast address (in parentheses) is automatically derived from the default VLAN IP address class. Press **<Enter>** to select the default address or enter a new address in dotted decimal notation and press **<Enter>**.

- e. The following prompt displays:

Description (30 chars max):

Enter a useful description for this virtual IP router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

- f. The following prompt displays:

Disable routing? (n) :

Indicate whether you want to disable routing in the group. You can enable routing later through the **modvl** command.

- g. The following prompt displays:

**IP RIP Mode {Deaf (d),
Silent (s),
Active (a),
Inactive (i)} (s):**

Define the RIP mode in which the virtual router port will operate. RIP (Router Information Protocol) is a network-layer protocol that enables the default VLAN in this Group to learn and advertise routes. The RIP mode can be set to one of the following:

Silent. The default setting shown in parentheses. RIP is active and receives routing information from other VLANs, but does not send out RIP updates. Other VLANs will not receive routing information concerning the default VLAN in this Group and will not include the VLAN in their routing tables. Simply press **<Enter>** to select Silent mode.

Deaf. RIP is active and sends routing information to other VLANs, but does not receive RIP updates from other VLANs. The default VLAN in this Group will not receive routing information from other VLANs and will not include other VLANs in its routing table. Enter **d** and press **<Enter>** to select Deaf mode.

Active. RIP is active and both sends and receives RIP updates. The default VLAN in this Group will receive routing information from other VLANs and will be included in the routing tables of other VLANs. Enter **a** and press **<Enter>** to select Active mode.

Inactive. RIP is inactive and neither sends nor receives RIP updates. The default VLAN in this Group will neither send nor receive routing information to/from other VLANs. Enter **i** and press **<Enter>** to select Inactive mode.

- h. If routing domains *are not* configured on the switch, go to the next step. If routing domains *are* configured on the switch, the following prompt displays:

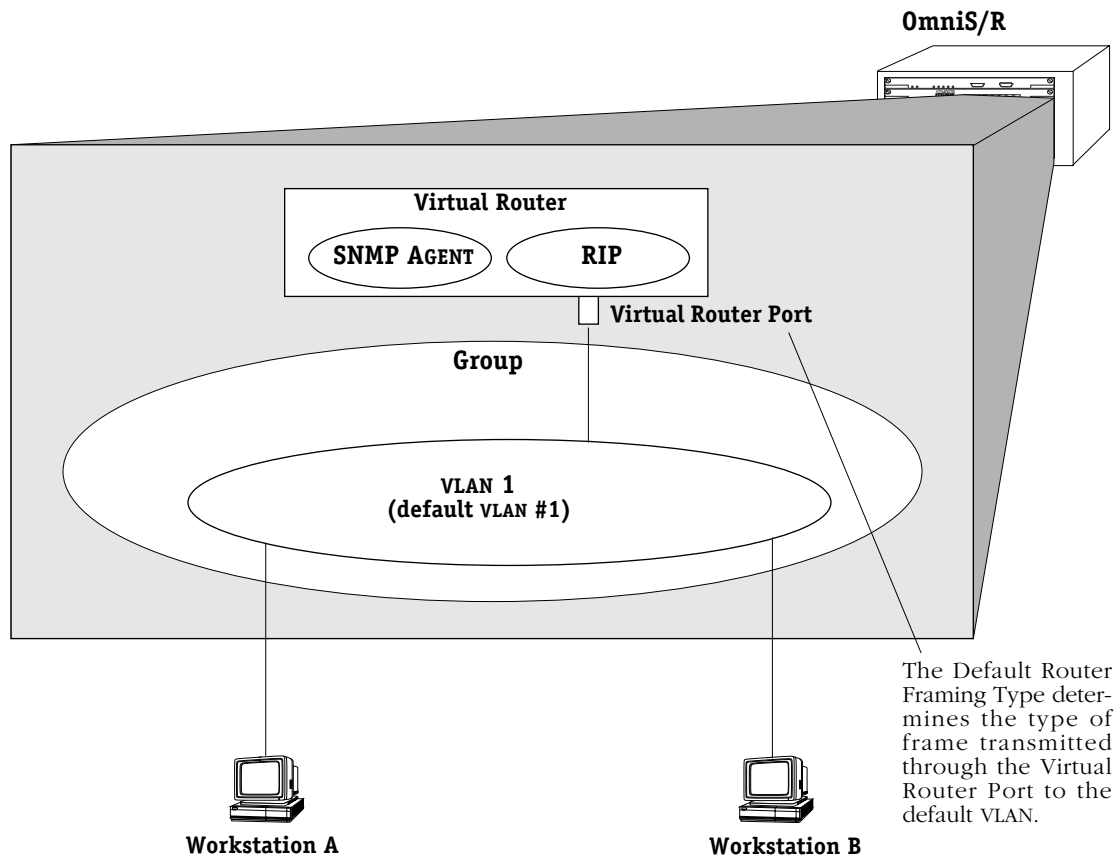
Apply to Routing Domain ID (none) :

Enter a routing domain in which this group should be included, or press **Enter**. A routing domain is a grouping of IP router interfaces that can forward packets only within the domain. Routing domains are part of Advanced Routing software and are not part of the base code. For more information about routing domains, see Chapter 14, "Routing Domains," in the *Advanced Routing User Manual*.

- i. After you enter the RIP mode, or after you enter a routing domain ID, the following prompt displays:

**Default framing type [Ethernet II(e),
fddi (f),
token ring (t),
Ethernet 802.3 SNAP (8),
source route token ring(s)} (e):**

Select the default framing type for the frames that will be generated by this router port and propagated over the default VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in the default VLAN. If the default VLAN contains devices using encapsulation types other than those defined here, the switching modules must translate those frames, which slows throughput. The figure on the next page illustrates the Default Framing Type and its relation to Virtual Router Port communications.



Default Framing Type and the Virtual Router Port

- j. You can now configure IPX routing on this port. The following message displays:

Enable IPX? (y) :

Press **<Enter>** if you want to enable IPX Routing on this virtual router port. If you do not enable IPX, then the default VLAN in this Group will not be able to route IPX data. You can set up a virtual router port to route both IP and IPX traffic.

If you don't want to set up an IPX router for the default VLAN in this Group, enter **n**, press **<Enter>**, and skip ahead to step **p** below. You can always set up IPX routing for other VLANs within this Group.

- k. After selecting to enable IPX, the following prompt displays:

IPX Network:

Enter the IPX network address. IPX addresses consist of eight hex digits and you can enter a minimum of one hex digit in this field. If you enter less than eight hex digits, the system prefixes your entry with zeros to create eight digits.

- l. The following prompt displays:

Description (30 chars max):

Enter a useful description for this virtual IPX router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

- m. The following prompt displays:

IPX Delay in ticks (0):

Enter the number of ticks you want for the IPX network. A tick is about 1/18th of a second. The default is 0.

- n. The following prompt displays:

**IPX RIP and SAP mode {RIP and SAP active (a)
RIP only active (r)
RIP and SAP inactive (i)} (a):**

Select how you want the IPX protocols, RIP (router information protocol) and SAP (service access protocol), to be configured for the default VLAN in this Group. RIP is a network-layer protocol that enables this VLAN to learn routes. SAP is also a network-layer protocol that allows network services, such as print and files services, to advertise themselves. The choices are:

RIP and SAP active. The default setting. The default VLAN to which this IPX router port is attached participates in both RIP and SAP updates. RIP and SAP updates are sent and received through this router port. Simply press **<Enter>** to select RIP and SAP active.

RIP only active. The default VLAN to which this IPX router port is attached participates in RIP updates only. RIP updates are sent and received through this router port. Enter an **r** and press **<Enter>** to select RIP only active.

RIP and SAP inactive. The IPX router port is active, but the default VLAN to which it is attached does not participate in either RIP nor SAP updates. Enter an **i** and press **<Enter>** to select RIP and SAP inactive.

- o. After selecting the RIP and SAP configuration, the following prompt displays the default router framing type options:

```

Default router framing type for : {
    Ethernet Media:
        Ethernet II (0),
        Ethernet 802.3 LLC (1),
        Ethernet 802.3 SNAP (2),
        Novell Ethernet 802.3 raw (3),

    FDDI Media:
        fddi SNAP (4),
        source route fddi SNAP (5),
        fddi LLC (6),
        source route fddi LLC (7),

    Token Ring Media:
        token ring SNAP (8),
        source route token ring SNAP (9),
        token ring LLC (a),
        source route token ring LLC (b) }      (0) :
    
```

Select the default framing type for the frames that will be generated by this router port and propagated over the default VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in the default VLAN. If the default VLAN contains devices using encapsulation types other than those defined here, the switching modules must translate those frames, which slows throughput. See the figure, *Default Framing Type and the Virtual Router Port* on page 19-23 for an illustration of the Default Framing Type and its relation to Virtual Router Port communications.

◆ Note ◆

The **.cmd** file contains a command called **hreXnative** that by default is set to 1. If physical ports in an end station are using a different encapsulation than the virtual router ports (for example, the **modvl** command shows router ports set to Ethernet II IPX, but the **swch** command shows that physical ports are using SNAP) then the **hreXnative** command *must* be set to 0. See Chapter 9, “Switch Wide Parameters,” for more information about the **.cmd** file.

- p. If you chose a Source Routing frame format in the last step (options 5, 7, 9, or b), an additional prompt displays:

```

Default source routing broadcast type : {
    ARE broadcasts(a), STE broadcasts(s)}      (a) :
    
```

Select how broadcasts will be handled for Source Routing. The choices are:

ARE broadcasts. All Routes Explorer, the default setting. Broadcasts are transmitted over every possible path on inter-connected source-routed rings. This setting maximizes the generality of the broadcast. Simply press **<Enter>** to select All Routes Explorer.

STE broadcasts. Spanning Tree Explorer. Broadcasts are transmitted only over Spanning Tree paths on inter-connected source-routed rings. This setting maximizes the efficiency of the broadcast. Enter an **s** and press **<Enter>** to select Spanning Tree Explorer.

- q. The following prompt displays:

Enter a priority level (0...7)(0):

Prioritizing VLANs allows you to set a value for traffic based on the destination VLAN of packets. Traffic with the higher priority destination will be delivered first. VLAN priority can be set from 0 to 7, with 7 being the level with the most priority.

Modifying and displaying a group's priority is described in *Priority VLANs* on page 19-73.

You have now completed the configuration of the virtual router port for this group. At this point, you will be asked whether you want to enable group mobility. The following prompt will display:

Enable Group Mobility on the Group ? [y/n] (n):

Mobile groups are discussed in detail in *Mobile Groups* on page 19-5. If you want to enable group mobility answer **Y** to this prompt, press **<enter>**, and go on to *Step 3. Set Up Group Mobility and User Authentication* on page 19-27.

If you do not want to configure group mobility answer **N** at the prompt, press **<enter>**, and go on to *Step 4. Configuring Virtual Ports* on page 19-28 for further instructions.

Step 3. Set Up Group Mobility and User Authentication

A mobile group offers more flexibility than a non-mobile group. With a mobile group, ports are assigned dynamically to the group based on AutoTracker policies that you configure. In a non-mobile group, ports are statically defined and AutoTracker policies are assigned to individual VLANs within the Group. In most cases, you will want to set up a mobile group. The following steps show you how.

- a. After configuring the virtual router port, you will receive the following prompt:

Enable Group Mobility on the Group ? [y/n] (n):

To create a mobile group, enter a **Y** as this prompt, press **<enter>**, and continue with step b. If you want to configure a non-mobile Group, enter **N**, press **<enter>**, and you will see the following prompt:

This Group will not participate in Group Mobility

If you are *not* creating a mobile group, go on to *Step 4. Configuring Virtual Ports* on page 19-28.

- b. The following prompt displays:

Enable User Authentication on the Group ? [y/n] (n):

An authenticated group is a special type of mobile group. It uses an authentication process as its criteria for group membership. Typically, users will be prompted for an id and password before gaining membership to an authenticated group. Authenticated groups require additional Windows NT server software. More detailed information on these groups can be found in the *Switch Network Services User Manual*. If you are not sure whether this is an authenticated group, simply press **<enter>** at this prompt.

- c. The following prompt displays:

Enable spanning tree for this group [y/n] (y):

Spanning Tree prevents broadcast storms by limiting logical loops in the network. For more information on Spanning Tree, see Chapter 17, titled “Configuring Bridging Parameters.” If you wish to enable Spanning Tree, enter **y** and press **<enter>**. Otherwise, enter **n**.

- d. The following prompt displays:

Do you wish to configure the interface group for this Virtual LAN at this time? (y)

You can assign physical ports to the new Group at this time. To begin assigning ports to the new Group, press **<Enter>** and go to Step 4.

To assign ports to the Group later, type **n** and **<Enter>**. The new Group is configured but does not yet contain any ports. You can use the **addvp** command later to assign ports to the Group (see *Adding Virtual Ports* on page 19-44). A message similar to the following displays confirming the creation of the new Group.

**GROUP 6 has been added to the system.
You may add interfaces to this group using the addvp command at a later date.
For now, the GROUP is inactive until you add interfaces.
Configure Auto-Activated LANE service ? [y/n] (y) :**

If you want to configure switch ports later (or simply rely on the dynamic port assignment capability’s of the mobile group) skip ahead to *Step 5. Configuring AutoTracker Policies (Mobile Groups Only)* on page 19-34.

Step 4. Configuring Virtual Ports

You can now enter configuration parameters for each switch port to be included in this Group. These configuration parameters include the bridging mode, output format type, and administrative state. In addition, if the port you are configuring is Ethernet (10/100 Mbps), you can also configure port mirroring.

Prompts for configuring virtual ports follow directly after Group Mobility prompts. You can choose to add ports now or add them later through the **addvp** command. Follow these steps:

- a. After you have stepped through the Routing and/or Group Mobility prompts, the following message displays:

Do you wish to configure the interface group for this Virtual LAN at this time? (y)

You can assign physical ports to the new Group at this time. To begin assigning ports to the new Group, press **<Enter>** and go to Step b.

To assign ports to the Group later, type **n** and **<Enter>**. The new Group is configured but does not yet contain any ports. You can use the **addvp** command later to assign ports to the Group (see *Adding Virtual Ports* on page 19-44). A message similar to the following displays confirming the creation of the new Group.

**GROUP 6 has been added to the system.
You may add interfaces to this group using the addvp command at a later date.
For now, the GROUP is inactive until you add interfaces.**

- b. After indicating that you want to set up ports, the following prompt displays:

Initial Vports (Slot/Phys Intf. Range) - For example, first I/O Module (slot 2), second interface would be 2/2. Specify a range of interfaces and/or a list as in: 2/1-3, 3/3, 3/5, 4/6-8

Enter the port or ports that you want to include in this new Group. The notation for adding a port to a group is

<slot number of module>/<port number on the module>

OmniS/R-3 are numbered from 1 to 3 top to bottom and OmniS/R-5 slots are numbered from 1 to 5 top to bottom. OmniS/R-9 slots are numbered 1-9, left to right. Port numbers are labelled on the front panel of switching modules.

You may enter multiple ports from multiple switching modules. For example, to add ports 1 through 3 on the module in slot 2, specify **2/1-3**. To additionally add the third and fifth port on the module in the third slot, specify **3/3, 3/5**. The complete slot port specification would be:

2/1-3, 3/3, 3/5

- c. If you enter a port that is already assigned to another Group, then you will be prompted on whether or not you want to change its assignment. A message similar to the following displays for each port that you enter:

**Initial Slot/Interface Assignments: 2/8
2/8 - This interface has already been assigned to GROUP 1 -
(Default GROUP #1).
Do you wish to remove it from that GROUP and assign it (with
new configuration values) to this GROUP (n)?**

Simply enter a **y** at each port prompt to change its Group assignment and begin setting port parameters. You could also enter a **c** at this prompt to accept all default port parameters and skip port configuration prompts. If you enter a **c**, *all* remaining ports are automatically added to the Group with default settings, and your work is complete.

- d. The virtual port configuration menu displays:

Modify Ether/8 Vport 2/8 Configuration

```

1) Vport                : 9
2) Description          :
3) Bridge Mode          : Auto-Switched
   31) Switch Timer     : 60
4) Flood Limit          : 192000
5) Output Format Type   : Default (IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status : Enabled, inactive
8) Mirrored Port Status : Disabled, available
9) MAC address          : 000000:000000

```

Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Redraw) :

Descriptions for each of the fields in this display follow. To change any default value, enter the line number for item, an equal sign (=), and then the value for the parameter. Enter **save** to save all configured settings and move onto the next step in the group creation process.

1) *Vport*

The virtual port number for this port. The next virtual port number available in the switch is shown by default in this field.

2) *Description*

Enter a useful description for this virtual port using alphanumeric characters. The description may be up to 30 characters long.

3) *Bridge Mode*

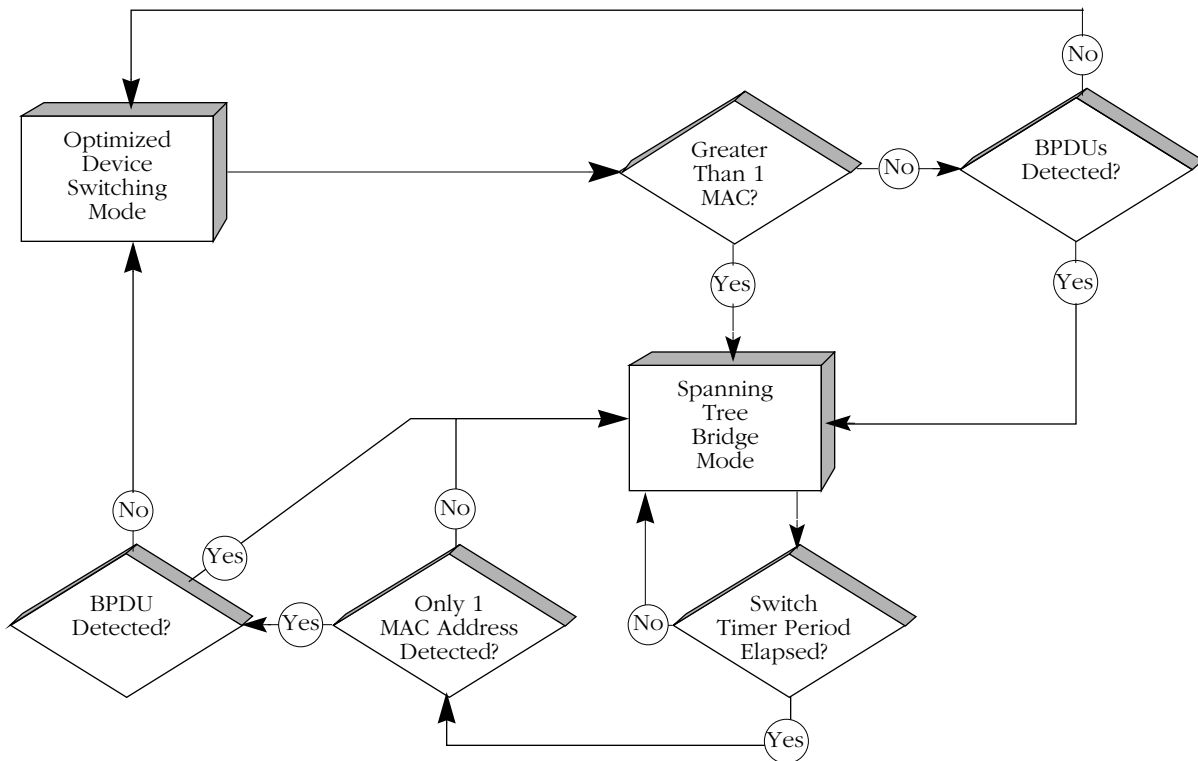
Select the bridge mode used by this port. The choices are:

Spanning Tree Bridge. The default setting for all non-Ethernet ports. This mode is appropriate for backbone and hub connections. The port acts as a standard 802.1d bridge port. It forwards BPDU frames out the port. When frames are received, Spanning Tree BPDUs are processed, and Spanning Tree dynamically controls the forwarding state. If flooding occurs, all frames destined for unknown MAC addresses, broadcast addresses, or multicast addresses will be sent to all ports in the same Group. Enter **3=b** and press **<Enter>** to select Spanning Tree Bridge mode.

Optimized Device Switching. This mode is appropriate for dedicated connections to a single workstation or server. Spanning Tree is turned off. No Spanning Tree BPDUs will be sent and the port will always be in the forwarding state. The port will stay in this mode even if a Spanning Tree BPDU is detected. In addition, all MACs learned will not be aged out (regardless of the Bridge Aging Timer setting) until the port is disconnected or configured to be administratively down. No flooding of packets with an unknown destination address is allowed after at least one MAC address has been learned. (An exception to this rule occurs on newer Mammoth-generation Ethernet modules, such as the ESM-100C-12, ESM-100F-8, and ESM-C-32. When these ports are in optimized mode, packets with unknown destination addresses will be flooded.) Packets with a broadcast or multicast destination will always be allowed. Enter **3=o** and press **<Enter>** to select Optimized Device Switching mode.

Auto-Switch. The default setting for all Ethernet ports. This mode is appropriate for dedicated connections requiring a switch-over to bridge mode when multiple devices are detected. A port in Auto-Switch mode will start in Optimized Device Switching mode (see description above). The port will remain in Optimized Device Switching mode until a Spanning Tree BPDU is detected or more than one MAC address transmits data. Once either of these conditions is met, the port will switch to Spanning Tree Bridge mode and Spanning Tree will start (if configured in the switch).

An Auto-Switch port will remain in Spanning Tree Bridge mode as long as there are BPDUs and multiple MACs. However, the port can revert back to Optimized Device Switching Mode if the time specified in the next field (**Switch Timer**) transpires without BPDUs and multiple MACs. Also, if the port is disconnected or configured to be administratively down, then an Auto-Switch port will revert back to Optimized Device Switching mode when it becomes operational again. Enter **3=a** and press **<Enter>** to select Auto-Switch mode.



How Auto-Switch Bridge Mode Works

31) Switch Timer

If you selected the Auto-Switch bridge mode, then you can configure this field. Enter the time-out period, in seconds, for an Auto-Switch port that has turned to Spanning Tree Bridge mode port to revert back to Optimized Switching mode. When in Auto-Switch mode, a port switches to Spanning Tree Bridge mode as soon as it detects a BPDU or more than one MAC address. The port will switch back to Optimized Switching mode after the time-out value you define here.

4) Flood Limit

The flood limit allows you to tune a virtual port to limit the flooding of broadcast, multi-cast, and unknown destination packets. This feature is useful for controlling broadcast storms on your network. While each network is different, in general the amount of flooded traffic represents a relatively small percentage of network traffic.

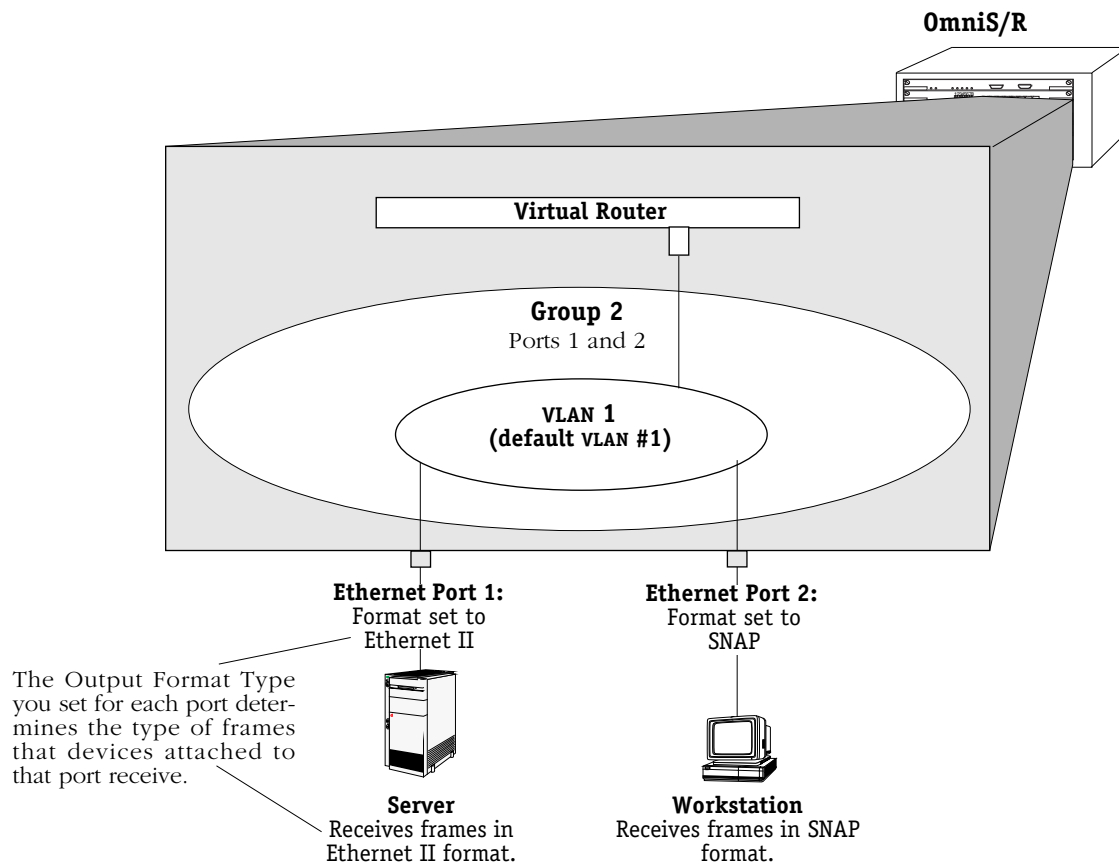
The flood limit is actually a “transmit credit” that is issued every five (5) seconds. When a packet is flooded on this port, the size of the packet, in bytes, is decremented from the current credit value. The credit value is the value you enter in this field multiplied by five. An additional credit, in the amount of the value you enter here multiplied by five, is allocated to each virtual port every five (5) seconds. If the credit value ever falls below zero, then all flooded packets are discarded until another credit is allocated. Flood limit checking is disabled if you enter a flood limit of zero (0). The flood limit default is 192,000 bytes per second, which equates to a transmit credit of 960,000 bytes every five seconds.

5) Output Format Type

The output format setting determines the kind of frame that will be sent out this physical port. If translation is necessary, then incoming frames will be translated to this format before being sent out this port. For example, on an Ethernet port incoming FDDI frames need to be translated to Ethernet. However, there are four types of Ethernet frames—Ethernet II, IPX 802.3, SNAP, and LLC. The format type you select here would determine the frame format to which non-Ethernet frames would be translated. The following figure illustrates how a port’s framing type affects communication with attached devices.

◆ Note ◆

This parameter differs from the router framing type selected during the configuration of the virtual router port. The router framing type is the encapsulation done on a router port, whereas this output format type applies only to translations on this virtual port.



Output Framing Type on Physical Ports

Note that for Ethernet, the default output format option is Ethernet II for IP frames and 802.3 for IPX frames.

You can customize your frame translation settings even further through the Switch menu. The Switch menu allows you to set translations at the frame format level (i.e., incoming SNAP frames could be translated one way, while incoming LLC frames could be translated another way) based on protocol type (IP or IPX). The Switch menu is explained in Chapter 18, "Configuring LAN Switch Translations."

6) Ethernet 802.2 Pass Through

For Ethernet ports only. If you answer **Yes** to this prompt, then frames received in the IEEE 802.2 format will not be translated according the Output Format Type chosen in line 5; they will be sent as is in their native IEEE 802.2 format. If you answer **No**, then 802.2 frames will be subject to the Output Format Type chosen in line 5.

7) Admin, Operational Status

Select whether to administratively enable or disable this port. When you enable the port, the port can transmit and receive data as long as a cable is connected and no physical or operational problems exist. When you disable a port, the port will not transmit or receive data even if a cable is connected and the physical connection is operational. If you disable the port at this point, you can enable it later through the **modvp** command (see *Modifying a Virtual Port* on page 19-45).

8) Mirrored Port Status

If the port you are configuring is Ethernet (10 or 10/100 Mbps), you can set up port mirroring. You can mirror traffic on this port to another like port. Port mirroring is a useful feature for monitoring traffic on particular ports. It is discussed in more detail later in this chapter in *Port Mirroring* on page 19-57.

If you want to mirror this port, enter a **8=e**, press **<Enter>** and you will be prompted for the slot and port number of the “mirroring” port (i.e., the port that can “see” all traffic for this port):

Mirroring vport slot/port ? () :

Enter the mirroring port’s slot and port number and press **<Enter>**.

If port mirroring is not supported on this port, then the following prompt will display:

mirroring not supported on this port type

9) MAC address

Enter the MAC address for this virtual port if it is known.

After the MAC address prompt, the switch confirms the addition of the port to the group with a message similar to the following:

Adding port 2/8 to Group 6. . .

Make configuration changes to the port until you are satisfied. If you have completed the final virtual port, then your work is complete. You can always alter Group parameters (including virtual router parameters for the default VLAN) later through the **modvl** command (see *Modifying a Group or VLAN* on page 19-40) and modify virtual port parameters through the **modvp** command (see *Modifying a Virtual Port* on page 19-45).

Step 5. Configuring AutoTracker Policies (Mobile Groups Only)

When you have completed configuring mobile group and auto-activated LANE services, you can begin configuring AutoTracker policies for this mobile group. Instructions for configuring these rules can be found in Chapter 20, “Configuring Group and VLAN Policies.” Please refer to that chapter for instructions on configuring each policy type. After you configure AutoTracker policies, you are done configuring this mobile group and a prompt similar to the following displays:

VLAN 9: 1 created successfully

You can configure rules for this group later through the **modatvl** command. This command also works with mobile groups as long as you indicate you want to alter VLAN 1 in the mobile group (i.e., the command line would read **modatvl 3:1** to modify mobile group 3).

◆ Note ◆

If the mobile group is initially created without rules, the **modatvl** command cannot be used to add them later. You must turn off group mobility and then reinstate it to add the rules.

Creating a WAN Routing Group

After entering basic Group information as described in *Step 1. Entering Basic Group Information* on page 19-19, you should have answered Yes to the following prompt:

Enable WAN Routing? (n):

if you want to enable WAN Routing. WAN Routing Groups are treated differently than other Groups, as described earlier. The following steps complete the configuration of the WAN Routing Group.

- a. After answering **y** to the **Enable WAN Routing?** prompt, the following prompt displays:

Enable IP (y):

Press **<Enter>** if you want to enable IP Routing on the virtual router port for this Group. If you do not enable IP, then this WAN Group will not be able to route IP data. If you don't want to set up IP routing, enter **n**, press **<Enter>** and skip to Step g.

◆ **Note** ◆

You may enable routing of both IP and IPX traffic over a WAN connection. If you set up dual-protocol routing, you must fill out information for both IP and IPX parameters.

- b. The following prompt displays:

IP Address:

Enter the IP address for this virtual router port in dotted decimal notation or hexadecimal notation (e.g., 198.206.181.10). This IP address is assigned to the virtual router port of the default VLAN within this Group. After you enter the address, press **<Enter>**.

- c. The following prompt displays:

IP Subnet Mask (0xfffff00):

The default IP subnet mask (in parentheses) is automatically derived from the default VLAN IP address class. Press **<Enter>** to select the default subnet mask or enter a new subnet mask in dotted decimal notation or hexadecimal notation and press **<Enter>**.

- d. The following prompt displays:

IP Broadcast Address (198.200.10.255):

The default IP broadcast address (in parentheses) is automatically derived from the default VLAN IP address class. Press **<Enter>** to select the default IP broadcast address or enter a new broadcast address in dotted decimal notation or hexadecimal notation and press **<Enter>**.

- e. The following prompt displays:

Description (30 chars max):

Enter a useful description for this virtual IP router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

- f. The following prompt displays:

```
IP RIP Mode {Deaf (d),  
Silent (s),  
Active (a),  
Inactive (i)} (s):
```

Define the RIP mode in which the virtual router port will operate. RIP (Router Information Protocol) is a network-layer protocol that enables the default VLAN in this Group to learn and advertise routes. The RIP mode can be set to one of the following:

Silent. The default setting shown in parentheses. RIP is active and receives routing information from other VLANs, but does not send out RIP updates. Other VLANs will not receive routing information concerning the default VLAN in this Group and will not include the VLAN in their routing tables. Simply press **<Enter>** to select Silent mode.

Deaf. RIP is active and sends routing information to other VLANs, but does not receive RIP updates from other VLANs. The default VLAN in this Group will not receive routing information from other VLANs and will not include other VLANs in its routing table. Enter **d** and press **<Enter>** to select Deaf mode.

Active. RIP is active and both sends and receives RIP updates. The default VLAN in this Group will receive routing information from other VLANs and will be included in the routing tables of other VLANs. Enter **a** and press **<Enter>** to select Active mode.

Inactive. RIP is inactive and neither sends nor receives RIP updates. The default VLAN in this Group will neither send nor receive routing information to/from other VLANs. Enter **i** and press **<Enter>** to select Inactive mode.

- g. You can now configure IPX routing on this port. The following message displays:

```
Enable IPX? (y) :
```

Press **<Enter>** if you want to enable IPX Routing on this virtual router port. If you do not enable IPX, then the default VLAN in this WAN Group will not be able to route IPX data. You can set up a virtual router port to route both IP and IPX traffic.

If you don't want to enable IPX routing for the default VLAN in this Group, enter **n** and press **<Enter>**. You can always set up IPX routing for other VLANs within this Group.

You are done configuring this WAN Routing Group. See the appropriate WAN interface chapter for further information on configuring this Routing service.

- h. After selecting to enable IPX, the following prompt displays:

```
IPX Network:
```

Enter the IPX network address. IPX addresses consist of eight hex digits and you can enter a minimum of one hex digits in this field. If you enter less than eight hex digits, the system prefixes your entry with zeros to create eight digits.

- i. The following prompt displays:

```
Description (30 chars max):
```

Enter a useful description for this virtual IPX router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

- j. The following prompt displays:

```
IPX Delay in ticks (0):
```

Enter the number of ticks you want for the IPX network. A tick is about 1/18th of a second. The default is 0.

- k. After entering a description, the following prompt displays:

```
IPX RIP and SAP mode {RIP and SAP active (a)
RIP only active (r)
RIP and SAP inactive (i)}
RIP and SAP triggered (t)}          (a):
```

Select how you want the IPX protocols, RIP (router internet protocol) and SAP (service access protocol), to be configured for the default VLAN in this Group. RIP is a network-layer protocol that enables this VLAN to learn routes. SAP is also a network-layer protocol that allows network services, such as print and files services, to advertise themselves. The choices are:

RIP and SAP active. The default setting. The default VLAN to which this IPX router port is attached participates in both RIP and SAP updates. RIP and SAP updates are sent and received through this router port. Simply press **<Enter>** to select RIP and SAP active.

RIP only active. The default VLAN to which this IPX router port is attached participates in RIP updates only. RIP updates are sent and received through this router port. Enter an **r** and press **<Enter>** to select RIP only active.

RIP and SAP inactive. The IPX router port is active, but the default VLAN to which it is attached does not participate in either RIP nor SAP updates. Enter an **i** and press **<Enter>** to select RIP and SAP inactive.

RIP and SAP triggered. The IPX router port is active, but RIP and SAP information will be sent out on the port only when a network change has occurred. This option is more cost effective for WAN links and is best suited for smaller network environments that don't change often. Enter a **t** and press **<Enter>** to select RIP and SAP triggered.

When you are done entering Router parameters, a message similar to the following displays:

```
GROUP 5 has been added to the system
```

You should now follow the instructions for configuring a WAN Routing Service described in the appropriate WAN interface chapter.

Viewing Current Groups

The **gp** command provides information on all currently defined Groups in a switch including Group number, network address, protocol type, and encapsulation type. You can obtain information on all groups in a switch by entering:

```
gp
```

A screen similar to the following displays:

Group ID (:VLAN ID)	Group Description	Network Address (IP Subnet Mask) or (IPX Node Addr)	Proto/ Encaps
1	Default GROUP (#1)	198.206.182.115 (ff.ff.ff.00)	IP / ETH2
2	New GROUP (#2)	198.206.101.12 (ff.ff.ff.00)	IP / SNAP
3	New GROUP (#3)	198.206.181.10 (ff.ff.ff.00)	IP/ 1490
4	New Group (#4)	198.206.183.44 (ff.ff.ff.00) 12314526 (0020da:020484)	IP / ETH2 IPX / 8023
5	New GROUP	198.206.143.11 (ff.ff.ff.00)	CIP / 1483

You can also get information on a specific Group by entering **gp** followed by the Group number. For example,

```
gp 3
```

displays information just on Group 3:

Group ID (:VLAN ID)	Group Description	Network Address (IP Subnet Mask) or (IPX Node Addr)	Proto/ Encaps
3	New GROUP (#3)	198.206.181.10 (ff.ff.ff.00)	IP / 1490

The following sections describe the columns in this table:

Group ID (:VLAN ID). The identification number assigned to this Group when it was created through the **crgp** command. The Group identifier is typically consistent network-wide (i.e., Group 3 in this switch should be the same Group as Group 3 configured in all other Omni Switch/Routers in the network). If this Group contains any VLANs, then they will be listed below the Group number. If the default VLAN in the Group supports both IP and IPX routing, then information on both (network address, etc) will display. Group 4 in the screen sample above shows a case where both IP and IPX routing are supported.

Group Description. The textual description of this Group that was entered when the Group was created or modified. This description is limited to 30 characters.

Network Address (IP Subnet Mask) or (IPX Node Addr). For each virtual router port configured, two addresses are listed. Both of these addresses were configured when the Group was created or modified through **crgp** or **modvl**. The first address is the Network Address, which is the address of the virtual router port for the default VLAN (VLAN #1) in this Group. For an IP virtual router port, this address is the IP address, which is shown in dotted decimal format. For an IPX virtual router port, this address is the IPX network address, which is shown as eight hex characters.

A second address is displayed below the Network address. For IP, this address is the IP Subnet Mask, which is normally derived from the default VLAN IP address class. For IPX, this address is the IPX Node Address.

Proto/Encaps. For each Group or VLAN listed, the top field is the Protocol supported by this virtual router port. Possible values in the field are: **IP** (IP router), **IPX** (IPX router), and **CIP** (Classical IP Group with CIP router). If you configured an IP and an IPX router port, then two router entries will be listed—one with a Protocol of IP and the other with a Protocol of IPX.

The bottom field is the encapsulation used for outgoing frames on the router port. This encapsulation was configured when the router port was configured. Possible values for this field depend on the Protocol and type of Group.

Frame Relay WAN Groups will always display **1490** to indicate RFC 1490 encapsulation is performed on frames.

IP and IPX routers have additional possible encapsulation types. For IP virtual router ports, the possible encapsulation types are as follows:

- **ETH2** Ethernet II
- **SNAP** Ethernet 802.3 SNAP
- **FDDI** FDDI
- **8025** Token Ring 802.5
- **TSRS** Token Ring Source Routing

For IPX virtual router ports, the possible encapsulation types are as follows:

- **ETH2** Ethernet II
- **LLC** Ethernet 802.3 LLC
- **SNAP** Ethernet 802.3 SNAP
- **8023** Ethernet 802.3 (Novell raw)
- **FDDI** FDDI SNAP
- **FSRS** FDDI Source Routing SNAP
- **FLLC** FDDI LLC
- **FSRL** FDDI Source Routing LLC
- **8025** Token Ring SNAP
- **TSRS** Token Ring Source Routing SNAP
- **TLLC** Token Ring LLC
- **TSRL** Token Ring Source Routing LLC

Modifying a Group or VLAN

After creating a Group (through **crgrp**) or VLAN (through **cratvl**, see Chapters 20 and 22), you can change any of their parameters through the **modvl** command. In addition, if you did not set up a virtual router port (IP or IPX) during the initial Group or VLAN configuration, you can set one up with **modvl**. To use this command, enter **modvl** followed by the Group number and VLAN number to change. For example, to modify parameters in Group 2, VLAN 1, enter:

```
modvl 2
```

Note that you do not need to specify a VLAN number to modify the default VLAN within a Group. To modify parameters in Group 2, VLAN 2, you would enter:

```
modvl 2:2
```

A screen similar to the following displays.

Current values associated with GROUP 2.1 are as follows:

```
1) GROUP Number      - 2:1
2) Description       - New GROUP (#2)
IP Parameters:
3) IP enabled        - Y
4) IP Network Address - 198.206.101.12
5) IP Subnet Mask    - 255.255.255.0
6) IP Broadcast Address - 198.206.101.255
7) Router Description - Router Port #2
8) RIP Mode          - Silent
                      {Active (a), Inactive (i), Deaf (d), Silent (s)}
9) Routing disabled  - N
11) Default Framing  - Ethernet II
                      {Ethernet II(e), Ethernet 802.3 (8), fddi (f),
                      token ring (t), source route token ring (s)}
IPX parameters:
12) IPX enabled      - N

(save/quit/cancel)
:
```

The Group number at the top of this sample screen is followed by the number 1 (**GROUP 2.1**), meaning that the information applies to default VLAN #1 in this Group. If this screen displayed information on Group 2, VLAN 2, then this field would read **GROUP 2:2**.

The colon prompt (:) at the bottom of the screen is used to prompt for user input. To change a value, type the line number of the item you want to change, followed by an equal sign (=) and the new value. For example, to set a new description you could enter:

```
2=Engineering
```

All of the **modvl** parameters are described in the section for creating a new Group, *Creating a New Group* on page 19-18.

◆ Note ◆

Line numbering for the **modvl** command will vary depending on whether you have an IP or IPX router configured. Each type of router contains several parameters that require extra line numbers.

Viewing Your Changes

When you enter a change at the colon prompt, the **modvl** screen does not normally refresh. If you want to see the current Group or VLAN settings, including any changes you made, enter a question mark (?) at the colon prompt. The **modvl** screen will refresh.

Saving Your Changes

Once you have entered all your modifications and you want to save them, type **save** at the colon prompt. You will exit the **modvl** command and your changes will take effect.

Canceling Your Changes

You can also exit the **modvl** command without saving any changes you made in the current session. Simply enter **cancel** at the colon prompt or enter **<Ctrl>-d**. The **modvl** command will end and none of the changes you made will be saved.

Changing the IP Address

Changing the IP address can also affect the Subnet Mask and the Broadcast Address. The new IP address means that the Subnet Mask and Broadcast Address must be re-generated and the following message displays:

**New IP address generates new subnet and broadcast address
Enter '?' to view the changes**

The system automatically creates new Subnet Mask and Broadcast addresses based on the new IP address. If you enter a question mark (?) at this point you could view these changes.

If you remove the last IP address in the system, you will see a warning message that SNMP (and other applications) are now inoperational.

Changing the IP Subnet Mask

Changing the IP Subnet Mask can also affect the IP Broadcast Address. The new Subnet Mask means that the Broadcast Address must be re-generated and the following message displays:

New mask caused change in broadcast address

The system automatically created a new Broadcast address based on the new Subnet Mask. If you entered a question mark (?) at this point you could view these changes.

Enabling IP or IPX Routing

If you enable IP or IPX routing by setting the corresponding **modvl** lines from **N** to **Y**, then the screen automatically refreshes with additional lines for the new router port parameters. All lines are set to router defaults. The router defaults are as follows:

IP Router

IP Network Address	0.0.0.0
IP Subnet Mask	0.0.0.0
IP Broadcast Address	0.0.0.0
Router Description	(no description shown for default)
Routing Disabled	No
RIP Mode	Silent
Default Framing Type	Ethernet II

IPX Router

IPX Network Address	0x0
Router Description	(no description shown for default)
Delay in Ticks	0
RIP/SAP Mode	RIP and SAP are active
Default Framing Type	Ethernet II

You can change any of these defaults as you would any other **modvl** parameters: enter the line number, followed by an equal sign (=) and the new parameter.

◆ Note ◆

You must at least enter a Network Address for a new router or you will not be able to save the configuration.

Deleting a Group

You can delete a Group as long as it does not contain any virtual ports. The default Group, Group #1, cannot be deleted. To delete a Group, enter **rmgp** followed by the Group number you want to delete. For example, if you wanted to delete Group 5, you would enter:

```
rmgp 5
```

If the Group does not contain any virtual ports, then a confirmation message displays:

```
GROUP 5 removed.
```

If the Group still contains virtual ports, then a message similar to the following displays:

```
GROUP 5 has active entries, you must remove  
these prior to removing the GROUP (use rmvp for this).
```

You must first remove the Group's virtual ports before the Group can be removed. The **rmvp** command allows you to remove virtual ports. See *Deleting a Virtual Port* on page 19-46 for information on using this command.

◆ Note ◆

Some commands in the Bridge Management menu (described in Chapter 17, "Configuring Bridging Parameters") require you to select a Group before making configuration changes. If you delete the currently selected Group with **rmgp**, then the new currently selected Group reverts to the default Group, Group #1.

Adding Virtual Ports

You can add virtual ports to a Group at any time after the Group is created. The **addvp** command allows you to add one or more ports to a Group you specify. If you have used the **crgp** command to add virtual ports, then you will find the **addvp** command fields very familiar.

To use **addvp**, enter the command followed by the Group number to which you want to add the port. Next, specify the port or ports you want to add.

addvp <Group Number for port> <Module Slot>/<Port Number>

For example, if you wanted to add ports 4 through 6 on the module in slot 4 to Group #5, then you would specify:

addvp 5 4/4-6

The procedure for using **addvp** is as follows:

1. Enter **addvp** followed by the Group number where you want this port to reside, followed by the physical slot and port numbers you want to configure.
2. If you enter a port that is already assigned to another Group, then you will be prompted on whether or not you want to change its assignment. A message similar to the following displays for each port that you enter:

**4/4 - This interface has already been assigned to GROUP 1 -
(Default GROUP #1).
Do you wish to remove it from that GROUP and assign it (with
new configuration values) to this GROUP (n)?**

Simply enter a **y** at each port prompt to change its Group assignment and begin setting port parameters. You could also enter a **c** at this prompt to accept all default port parameters and skip port configuration questions. If you enter a **c**, *all* remaining ports are automatically added to the Group with default settings, and your work is complete.

3. The virtual port configuration menu displays:

Modify Ether/8 Vport 4/4 Configuration

```

1) Vport                : 9
2) Description          :
3) Bridge Mode          : Auto-Switched
   31) Switch Timer      : 60
4) Flood Limit          : 192000
5) Output Format Type   : Default (IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status : Enabled, inactive
8) Mirrored Port Status : Disabled, available
9) MAC Address          : 000000:000000
    
```

Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Redraw) :

Descriptions for each of the fields in this display begin on page 19-29. To change any default value, enter the line number for the item, an equal sign (=), and then the value for the parameter. When you have completed the configuration for this port, enter **save** to save all configured settings.

Modifying a Virtual Port

You can modify a virtual port through the **modvp** command. The **modvp** command is very similar to the **addvp** command and the port configuration phase of the **crpp** command. To use **modvp**, enter the command, followed by the Group number for the port, and the physical slot and port number for the port:

```
modvp <Group Number for port> <Module Slot>/<Port Number>
```

You can specify only one port at a time. For example, if you wanted to modify the parameters for Port 7 on the module in Slot 4, and the Port currently resides in Group 6, then you would enter:

```
modvp 6 4/7
```

The procedure for using **modvp** is as follows:

1. Enter **modvp** followed by the Group number where the port currently resides, the physical slot and port number.
2. A prompt displays requesting your confirmation:

```
Modify local port 7 (Virtual port (#14)) ? (y) :
```

Simply press **<Enter>** if this is the correct virtual port. The Virtual Port number in parentheses (**Virtual Port #14** in this case) is the virtual port number within this entire Omni Switch/Router. Virtual ports are numbered sequentially within the switch, not within a Group or VLAN.

3. The virtual port configuration menu displays:

Modify Ether/8 Vport 4/7 Configuration

```

1) Vport                : 9
2) Description          :
3) Bridge Mode          : Auto-Switched
   31) Switch Timer     : 60
4) Flood Limit          : 192000
5) Output Format Type    : Default (IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status : Enabled, inactive
8) Mirrored Port Status : Disabled, available
9) MAC Address          : 000000:000000

```

```
Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Redraw) :
```

Descriptions for each of the fields in this display begin on page 19-29. To change any default value, enter the line number for the item, an equal sign (=), and then the value for the parameter. When you have completed the configuration for this port, enter **save** to save all configured settings.

Deleting a Virtual Port

You can delete a virtual port from its existing Group by using the **rmvp** command. When you remove a virtual port, the port is moved to the default switch Group, Group #1, and all port parameters are reset to defaults except for the port name. For example, if you configured a port with a special flood limit and customized translation settings and you then removed the port, you would lose those port settings.

To remove a port, enter the **rmvp** command, followed by the Group number where the port currently resides and the physical slot and port number for the port:

```
rmvp <Group number> <Module Slot>/<Port Number>
```

For example, to delete Port 7 on the module in Slot 4, and the Port currently resides in Group 6, you would enter:

```
rmvp 6 4/7
```

A prompt displays requesting that you confirm the deletion:

```
Local port 7 (Virtual po...) is attached to this slot/interface - remove? (n):
```

Enter a **y** and press **<Enter>** to remove the port. Another message displays confirming the deletion:

```
BRIDGE port on 4/7 moved to GROUP 1.
```

If the port you specified did not exist in the Group you specified in the **rmvp** command, then a message similar to the following would display:

```
Specified port(s) not found on GROUP 6.
```

Viewing Information on Ports in a Group

The **via** command allows you to view port attachments associated with a specified Group or all Groups in a switch. Entering

```
via
```

displays summary information for all virtual ports in the switch. You can also display virtual interface attachments for a specific Group by specifying the Group ID after the **via** command. For example, to view ports for Group 2, you would enter

```
via 2
```

The same type of information is displayed for a single Group as is displayed for all Groups. The following screen shows a sample from the **via** command when specified without a Group ID.

GROUP Interface Attachments For All Interfaces

GROUP: Slot/Intf	Description	Service/ Instance	Protocol	Admin Status
1.1 : *	GROUP #1.0 IP router vport	Rtr / 1	IP	Enabled
2.1 : *	for group 2	Rtr / 2	IP	Enabled
1:2/1	Virtual port (#2)	Brg / 1	Tns	Enabled
1:2/2	Virtual port (#3)	Brg / 1	Tns	Enabled
1:2/3	Virtual port (#4)	Brg / 1	Tns	Enabled
2:2/4	finance server	Brg / 1	Tns	Enabled
1:2/5	Virtual port (#6)	Brg / 1	Tns	Enabled
1:2/6	Virtual port (#7)	Brg / 1	Tns	Enabled
1:2/7	Virtual port (#8)	Brg / 1	Tns	Enabled
1:2/8	Virtual port (#9)	Brg / 1	Tns	Enabled
1:3/1	Virtual port (#1)	Brg / 1	Tns	Enabled
1:4/1	Virtual port (#10)	Brg / 1	Tns	Enabled
1:4/2	Virtual port (#11)	Brg / 1	Tns	Enabled
1:4/3	Virtual port (#12)	Brg / 1	Tns	Enabled
1:4/4	Virtual port (#13)	Brg / 1	Tns	Enabled
1:4/5	Virtual port (#14)	Brg / 1	Tns	Enabled
1:4/6	Virtual port (#15)	Brg / 1	Tns	Enabled

GROUP: Slot/Intf. **GROUP** is the group number to which this port is assigned. When the Group displays as a Group number followed by a decimal and a 1 (1.1 and 2.1 in the above sample), it represents the router port on the default VLAN within that Group. **Slot** is the position in the chassis of the switching module where this port is located. **Intf** (Interface) is the physical port on the switching module. When the Slot and Interface are shown as an asterisk (*)—as the top two entries in the above table display—it represents as virtual router port that does not have a corresponding physical interface.

Description. The textual description entered for either the virtual router port or the virtual switch port. This description was entered through **crgp** or **modvl** for virtual router ports, or through **crgp**, **addvp**, or **modvp** for virtual switch ports.

Service/Instance. **Service** is the service type configured for this port. **Instance** is an identifier of this service type within the switch. For example, multiple virtual router ports within the switch will be labelled consecutively (1, 2, 3, etc.), and will each have a different **Instance** number.

Values for the service type are as follows:

Viewing Information on Ports in a Group

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **CIP** Classical IP service port
- **Vlc** VLAN Clusters (X-LANE) service port

Protocol. The bridging protocol for virtual ports and services or the routing protocol for virtual router ports. Possible values are:

- **Tns** Transparent bridge. Bridges maintain a dynamic table of known MAC addresses on connected segments. The table is used to make forwarding decisions. When a frame is received that contains a destination address that matches an address in the table, it is forwarded to designated bridge ports that are in forwarding state.
- **SR** Source Routing Bridge. Normally used in Token Ring environments. Routing information is determined by looking at the Routing Information Field (RIF) in a frame. The RIF contains the segment and bridge numbers that create the path to the destination.
- **SRT** Source Routing Transparent. Normally used in Token Ring environments. Allows Source Routing and Transparent bridges to coexist. The Source Routing Transparent Bridge will form a Spanning Tree with other Transparent Bridges and Source Routing Transparent Bridges and will forward frames that do not contain a Routing Information Field (RIF) to destinations reachable by the Spanning Tree. If the bridge detects routing information in the RIF, it will forward it the same way Source Routing bridges do.”
- **IP** IP Routing Protocol. Routing Information Protocol (RIP) used to learn routes from neighboring routers. You configure an IP router through the **crgp** or **modvl** commands. Other IP routing parameters can be set through the Networking menu commands, which are described in Chapter 25, “IP Routing.”
- **IPX** IPX Routing Protocol. Uses RIP to learn routes from neighboring routers and the Service Advertising Protocol (SAP) to maintain a database of network services for requesting workstations. Other IPX routing parameters can be set through the Networking menu commands, which are described in Chapter 27, “IPX Routing.”
- **FR** Frame Relay IP Routing. WAN Routing Groups are configured slightly different from other Groups. Frame Relay IP Routing is IP Routing with some enhancements to account for the Frame Relay network.

Admin Status. Indicates whether the port is administratively **Enabled** or **Disabled**. When **Enabled**, the port can transmit and receive data as long as a cable is connected and no physical or operational problems exist. When **Disabled**, the port will not transmit or receive data even if a cable is connected and the physical connection is operational. You can set the Admin Status during port configuration phase of the **crgp**, **addvp**, or **modvp** commands.

Viewing Detailed Information on Ports

The **vi** command displays detailed information about virtual ports. Entering

```
vi
```

displays information for all virtual ports in the switch. You can also display information for only ports in a specific Group by specifying the Group ID after the **vi** command. For example, to view information only for ports in Group 6, you would enter

```
vi 6
```

The same type of information is displayed for a single Group as is displayed for all Groups. The following screen shows a sample from the **vi** command when specified without a Group ID.

Virtual Interface Summary Information- For All Interfaces										
									Status	
Group	Slot/ Intf	Type/ Inst/Srvc	MAC Address	Prt	Encp	Admin	Oper	Spn	Tr	Mode
1	All	Rtr/ 1	0020da:020d40	IP	ETH2	Enabl	Active	N/A	N/A	
2	All	Rtr/ 2	0020da:020d43	IP	ETH2	Enabl	Active	N/A	N/A	
2	All	Rtr/ 3	0020da:020d44	IP	ETH2	Enabl	Active	N/A	N/A	
1	3/1	Brg/ 1/ 1	0020da:048730	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
1	4/1	Brg/ 1/ na	0020da:030990	Tns	DFLT	Enabl	Active	Fwdng	Bridged	
1	4/2	Brg/ 1/ na	0020da:030991	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
1	4/3	Brg/ 1/ na	0020da:030992	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
1	4/4	Brg/ 1/ na	0020da:030993	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
1	4/5	Brg/ 1/ na	0020da:030994	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
1	4/6	Brg/ 1/ na	0020da:030995	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
1	4/7	Brg/ 1/ na	0020da:030996	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
2	4/8	Brg/ 1/ na	0020da:030997	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	
1	5/1	Brg/ 1/ na	0020da:022860	Tns	DFLT	Enabl	Inactv	Disabl	Bridged	

Group. The Group number to which this port is currently assigned.

Slot/Intf. The slot (**Slot**) is the position in the chassis of the switching module where this port is located. The interface (**Intf**) is the physical port on the switching module. If this column reads **All**, then this port is a router port that supports all virtual ports in the Group.

Type/Inst/Srvc. The Service Type (**Type**), Instance (**Inst**) of this Service Type in the switch, and service number (**Srvc**) for this virtual port. Service Type values are as follows:

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **Vlc** VLAN clusters (X-LANE) service port
- **CIP** Classical IP service port

The Instance (**Inst**) is an identifier of this type of service within the switch. For example, if more than one virtual router port is configured in the switch, then each “instance” of a router will be given a different number. The service number (**Srv**) is port-specific. If a port has more than one service configured on it, then each service will be identified by a different service number.

MAC Address. The MAC address for this virtual port. Each virtual port is allocated a MAC address.

Prt. The bridging or routing protocol supported by this virtual port. Descriptions of these protocol types are provided on page 19-48. Possible values are:

- **Tns** Transparent Bridge
- **SR** Source Routing Bridge
- **SRT** Source Routing Transparent Bridge
- **IP** IP Routing Protocol
- **IPX** IPX Routing Protocol
- **CIP** Classical IP Routing (RFC 1577)
- **FR** Frame Relay IP Routing

Encp. Encapsulation used for outgoing packets on this virtual router or switch port. Possible encapsulation values are:

- **DFLT** Default format for this switch port (differs for each interface type)
- **SWCH** Frame translations have been customized through the Switch menu
- **ETH2** Ethernet II
- **ESNP** Ethernet 802.3 SNAP (virtual router ports)
- **ELLC** Ethernet 802.3 LLC (IPX router ports only)
- **8023** Ethernet 802.3, Novell Raw (IPX router ports only)
- **8025** Token Ring 802.5 SNAP (virtual router ports)
- **TSRS** Token Ring Source Routing SNAP (virtual router ports)
- **TLLC** Token Ring LLC (IPX router ports only)
- **TSRL** Token Ring Source Routing LLC (IPX router ports only)
- **FDDI** FDDI SNAP (virtual router ports)
- **FSRS** FDDI Source Routing SNAP (IPX router ports only)
- **FLLC** FDDI LLC (IPX router ports only)
- **FSRL** FDDI Source Routing LLC (IPX router ports only)
- **1490** Frame Relay Routing (RFC 1490)
- **1483** Classical IP Routing (RFC 1483)
- **SNAP** SNAP (switch ports only)
- **LLC** LLC (switch ports only)

Admin. Indicates whether the port is administratively Enabled or Disabled. When **Enabld**, the port can transmit and receive data as long as a cable is connected and no physical or operational problems exist. When **Disabld**, the port will not transmit or receive data even if a cable is connected and the physical connection is operational. You can set the Administrative Status during the port configuration phase of the **crjgp** command, the **addvp** command, or the **modvp** command. A port can have an Administrative Status of Enabled, but still be operationally Inactive. See the description of the **Oper** column below.

Oper. Indicates the current Operational Status of the port. The port will be Active (**Active**) or Inactive (**Inactv**). If the port is Active, then the port can pass data and has a good physical connection. If it is Inactive, then it may not have a good physical connection and it is not capable of passing data at this time.

Spn Tr. The port's current state as defined by the Spanning Tree Protocol. The possible Spanning Tree States are: Disabled, Blocking, Listening, Learning, and Forwarding. This state controls the action a port takes when it receives and transmits a frame. For ports which are Administratively disabled or Operationally Inactive, this state will be Disabled (**Disabl**), meaning the Spanning Tree algorithm is not active on this port. If the state is **Blocking**, then only BPDUs will be transmitted and received. If the state is **Forwarding**, then both data and BPDU frames will be transmitted and received. This Spanning Tree Protocol state is not applicable to virtual router ports and will read **N/A** for those ports.

Mode. The Bridge Mode currently in use on this port. This mode is chosen during the port configuration phase of the **crgp** command, through the **addvp** command, or through the **modvp** command. It is not applicable to virtual router ports and will read **N/A** for those ports. Possible values are:

- **Bridged** Spanning Tree Bridge.
- **AutoSw** Auto Switch.
- **Optimzd** Optimized Device Switching.

See page 19-29 for a description of these bridge modes.

Viewing Port Statistics

The **vs** command displays transmit and receive statistics for ports in the switch. Entering

```
vs
```

displays statistics for all virtual ports in the switch. You can also display statistics for only ports in a specific Group by specifying the Group ID after the **vs** command. For example, to view statistics only for ports in Group 6, you would enter

```
vs 6
```

You can also display statistics for a specific port by entering the slot and port number after the **vs** command. For example, to view statistics only for Port 1 on the module in Slot 4, you would enter

```
vs 4/1
```

The same type of information is displayed for a single Group or port as is displayed for all ports in a switch. The following screen shows a sample from the **vs** command when specified without any Group or port parameters.

Virtual Interface Statistical Information- For All Interfaces						
Slot/ Group	Intf	Service/ Instance	Frames In Out	Octets In Out	UcastPkts In Out	NUcastPkts In Out
1	All	Rtr/ 1				
2	All	Rtr/ 2				
3	All	Rtr/ 3				
1	3/1	Tnk/ 1	0	0	0	0
			0	0	0	0
1	4/1	Brg/ 1	17774	1739560	1707	16067
			684	103048	681	3
1	4/2	Brg/ 1	0	0	0	0
			0	0	0	0
1	4/3	Brg/ 1	0	0	0	0
			0	0	0	0
1	4/4	Brg/ 1	0	0	0	0
			0	0	0	0
1	4/5	Brg/ 1	0	0	0	0
			0	0	0	0
1	4/6	Brg/ 1	0	0	0	0
			0	0	0	0
1	4/7	Brg/ 1	0	0	0	0
			0	0	0	0
1	4/8	Brg/ 1	0	0	0	0
			0	0	0	0
1	5/1	Brg/ 1	0	0	0	0
			0	0	0	0

Group, Slot/Intf. These columns are described for the **vi** command on page 19-50.

Service/Instance. The Service Type (**Service**) and Instance (**Instance**) of this Service Type in the switch.

Service Type values are as follows:

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **Vlc** VLAN clusters (X-LANE) service port
- **CIP** Classical IP service port

The Instance (**Inst**) is an identifier of this type of service within the switch. For example, if more than one virtual router port is configured in the switch, then each “instance” of a router will be given a different number.

Frames In/Out. The number of frames received or sent from this port. The top number for each port row is the number of frames received, and the bottom number is the number of frames sent. Statistics are not provided for virtual router ports in this display, but they are provided through Networking menu commands. See Chapters 25 and 27 for further information on router port statistics.

Octets In/Out. The number of octets, or bytes, received or sent from this port. The top number for each port row is the number of octets received, and the bottom number is the number of octets sent. Statistics are not provided for virtual router ports, but they are provided through Networking menu commands. See Chapters 25 and 27 for further information on router port statistics.

Ucast Pkts In/Out. The total number of unicast packets received or sent from this port. The top number for each port row is the number of unicast packets received, and the bottom number is the number of unicast packets sent. Statistics are not provided for virtual router ports, but they are provided through Networking menu commands. See Chapters 25 and 27 for further information on router port statistics.

Non Ucast Pkts In/Out. The total number of non-unicast packets received or sent from this port. Non-unicast frames include multicast and broadcast frames. The top number for each port row is the number of non-unicast packets received, and the bottom number is the number of non-unicast packets sent. Statistics are not provided for virtual router ports, but they are provided through Networking menu commands. See Chapters 25 and 27 for further information on router port statistics.

Viewing Port Errors

The **ve** command displays port error statistics for ports in the switch. Entering

```
ve
```

displays error statistics for all virtual ports in the switch. You can also display errors statistics for only ports in a specific Group by specifying the Group ID after the **ve** command. For example, to view errors only for ports in Group 6, you would enter

```
ve 6
```

You can also display error statistics for a specific port by entering the slot and port number after the **ve** command. For example, to view errors only for Port 1 on the module in Slot 4, you would enter

```
ve 4/1
```

The same type of information is displayed for a single Group or port as is displayed for all ports in a switch. The following screen shows a sample from the **ve** command when specified without any Group or port parameters.

Virtual Interface Error Information- For All Interfaces

Group	Slot/ Intf	Service/ Instance	Buffer Discards In	Out	Error Discards In	Out
2	All	Rtr/ 1				
3	All	Rtr/ 2				
1	All	Rtr/ 1				
1	3/1	Tnk/ 1	0	0	0	0
1	4/1	Brg/ 1	0	0	0	0
1	4/2	Brg/ 1	0	0	0	0
1	4/3	Brg/ 1	0	0	0	0
1	4/4	Brg/ 1	0	0	0	0
1	4/5	Brg/ 1	0	0	0	0
1	4/6	Brg/ 1	0	0	0	0
1	4/7	Brg/ 1	0	0	0	0
1	4/8	Brg/ 1	0	0	0	0
1	5/1	Brg/ 1	0	0	0	0

Group, Slot/Intf. These columns are described for the **vi** command on page 19-50.

Service/Instance. The Service Type (**Service**) and Instance (**Instance**) of this Service Type in the switch. Service Type values are as follows:

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **Vlc** VLAN clusters (X-LANE) service port
- **CIP** Classical IP service port

Viewing Port Errors

The Instance (**Inst**) is an identifier of this type of service within the switch. For example, if more than one virtual router port is configured in the switch, then each “instance” of a router will be given a different number.

Buffer Discards In/Out. For transmit (**Out**) and receive (**In**), the number of frames discarded due to a lack of buffer space. Buffer discard information is not provided for virtual router ports.

Error Discards In/Out. For transmit (**Out**) and receive (**In**), the number of frames discarded due to errors. Error discard information is not provided for virtual router ports.

Port Mirroring


You can set up Port Mirroring for any pair of Ethernet (10 or 10/100 Mbps) within the same switch chassis. Ethernet ports supporting port mirroring include 10BaseT (RJ-45), 10BaseFL (fiber), 10Base2 (BNC), and 10Base5 (AUI) connectors. When you enable port mirroring, the active, or “mirrored,” port transmits and receives network traffic normally, and the “mirroring” port receives a copy of all transmit and receive traffic to the active port. You can connect an RMON probe or network analysis device to the mirroring port to see an exact duplication of traffic on the mirrored port without disrupting network traffic to and from the mirrored port.

Port mirroring is supported on Omni Switch/Router chassis for Ethernet (10 or 10/100 Mbps) ports only. An Ethernet port can only be mirrored by one other Ethernet port. A mirroring port can only mirror one port at a time. Up to five (5) mirroring sessions (mirrored-mirroring port pairs) are supported in a single switch chassis. The mirrored and mirroring ports can be in different Groups and different VLANs.

How Port Mirroring Works

When a frame is received on a Mirrored Port it is copied and sent to the Mirroring Port. The received frame is actually transmitted twice across the switch backplane—once for normal bridging and then again to the Mirroring Port.

When a frame is transmitted by the mirrored port, a copy of the frame is made, tagged with the mirroring port as the destination, and sent back over the switch backplane to the mirroring port. The following diagram illustrates the data flow for a Mirrored-Mirroring port pair.



Relationship Between Mirrored and Mirroring Port

When port mirroring is enabled, there may be some performance degradation since all frames received and transmitted by the Mirrored port need to be copied and sent to the Mirroring port.

What Happens to the Mirroring Port

Once you set up port mirroring and attach cables to the Mirrored and Mirroring ports, the Mirroring port is administratively disabled and no longer a part of the Bridging Spanning Tree. The Mirroring port does not transmit or receive any traffic on its own. In addition, the Admin Status of the mirroring port displays in switch software commands, such as `vi`, as

```
M <slot> <port>
```

where **<slot>** is the slot number of the module containing the mirrored port, and **<port>** is the port number of the mirrored port. For example, if the Admin Status of a port displayed as

M 3 02

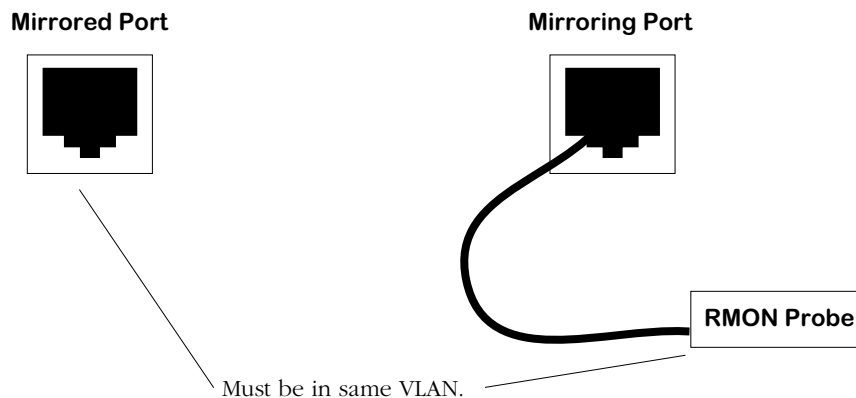
then you would know this port is mirroring traffic for Port 2 on the module in Slot 3.

If a cable is not attached to the Mirrored port, port mirroring will not take place. In this case, the Mirroring Port reverts back to its normally operational state and will bridge frames as if port mirroring were disabled.

Using Port Mirroring With External RMON Probes

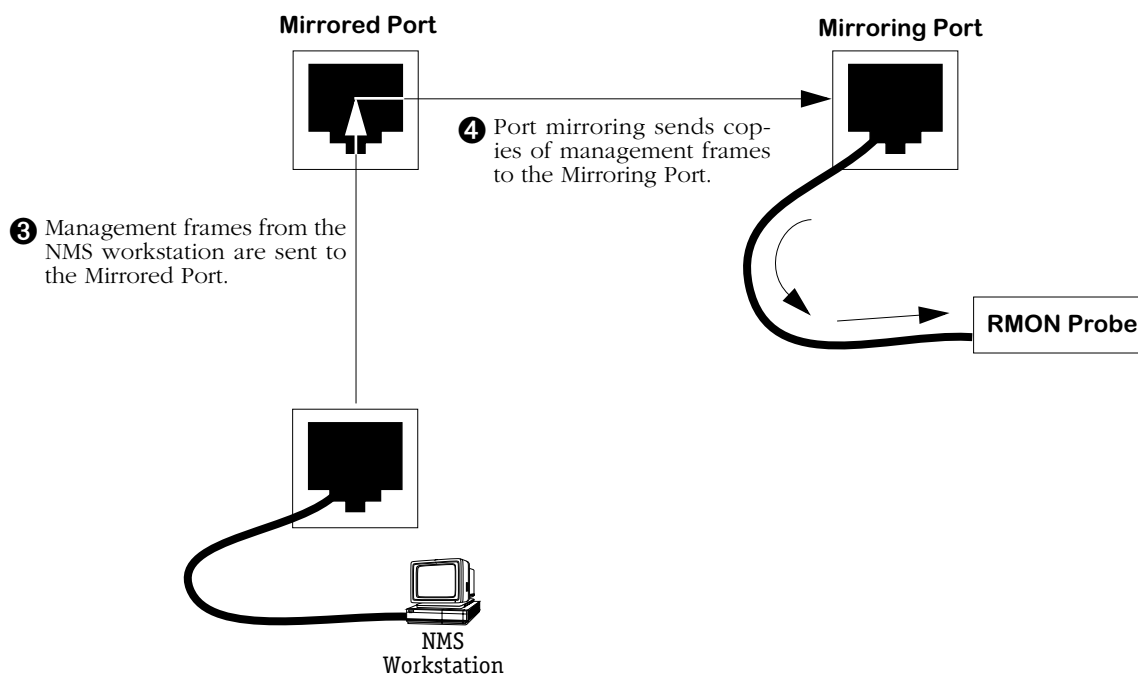
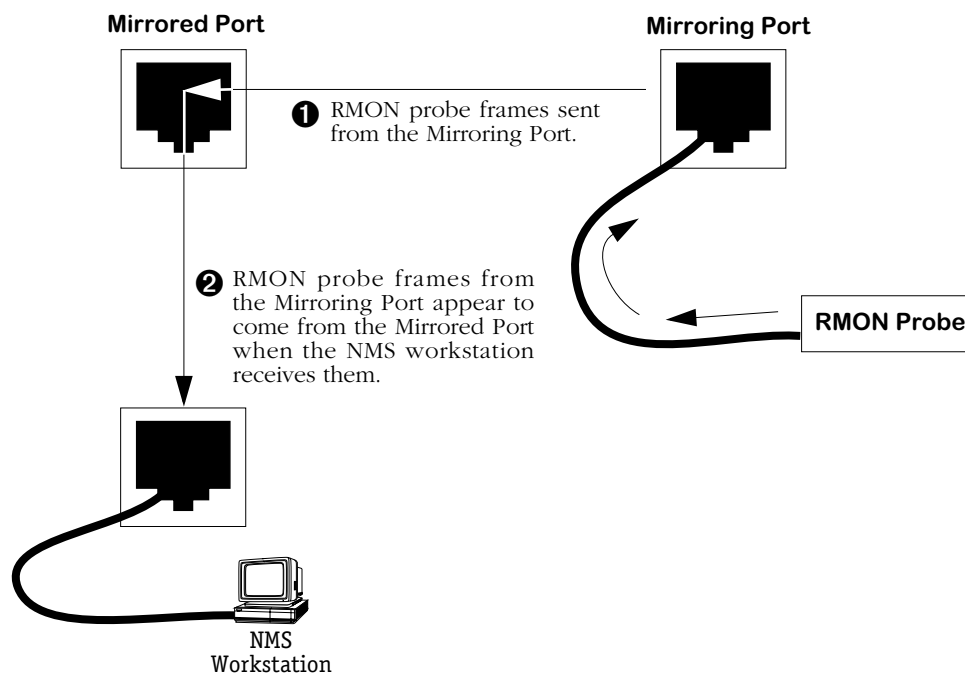
Port mirroring is a helpful monitoring tool when used in conjunction with an external RMON probe. Once you set up port mirroring, the probe can collect all relevant RMON statistics for traffic on the mirrored port. You can also move the Mirrored Port so that the Mirroring Port receives data from different ports. In this way, you can roam the switch and monitor traffic at various ports.

If you attach an external RMON probe to a mirroring port, that probe must have an IP address that places it in the same VLAN as the mirrored port. In addition if you change the mirrored port, then you must again make sure that the RMON probe is in the same VLAN as that new mirrored port.



Mirrored and Mirroring Ports in Same VLAN

Frames received from an RMON probe attached to the Mirroring Port can be seen as being received by the Mirrored Port. These frames from the Mirroring Port are marked *as if they are received on the Mirrored Port* before being sent over the switch backplane to an NMS station. Therefore, management frames from an NMS station that are destined for the RMON probe are first forwarded out the Mirrored Port. After being received on the Mirrored Port, copies of the frames are mirrored out the Mirroring Port—the probe attached to the Mirroring Port receives the management frames. The illustration on the following page shows this data flow.



Port Mirroring Using an External RMON Probe

◆ Important Note ◆

The Mirroring Port is not accessible from the NMS device. From the NMS station, the Mirroring Port will appear disabled or down.

Setting Up Port Mirroring

You set up port mirroring when you add or modify a port through the **addvp** or **modvp** commands. The switch software senses the type of port you are configuring, so it will only prompt you for port mirroring when configuring an Ethernet port. Follow the steps below to set up port mirroring.

1. Start the **addvp** or **modvp** command for the virtual port that you want to mirror.
2. At the **Command** prompt enter **8=e**, press **<Enter>** and you will be prompted for the slot and port number of the “mirroring” port (i.e., the port that can “see” all traffic for this port):

Mirroring vport slot/port ? () :

3. Enter the mirroring port’s slot, a slash (/), the port number, and then press **<Enter>**. The port that you indicate here will be disabled and only capable of receiving duplicate traffic from the mirrored port. If port mirroring is not supported on this port, then the following prompt will display:

mirroring not supported on this port type

After entering the Mirroring slot and port number, the **addvp** or **modvp** screen of options re-displays with the changes you entered. If you are done modifying or adding the port, enter **save** at the **Command** prompt. If using the **addvp** command a message indicating that you have successfully set up the port displays. Port mirroring takes place immediately, so you could now connect a probe or network analyzer to the Mirroring port.

Disabling Port Mirroring

You can disable port mirroring through the **modvp** command. Follow these steps to disable port mirroring.

1. Start the **modvp** command for the virtual port on which you want to disable port mirroring.
2. At the **Command** prompt enter **8=d**, press **<Enter>**. The **modvp** screen re-displays. The **Mirrored Port Status** field should read **Disabled, available**.

Port Monitoring

An essential tool of the network engineer is a network packet capture device. A packet capture device is usually a PC-based computer, such as the Sniffer®, that provides a means for understanding and measuring data traffic of a network. Understanding data flow in a VLAN-based switch presents unique challenges primarily because traffic takes place *inside* the switch, especially on dedicated devices.

The port monitoring feature built into OmniS/R software allows the network engineer to examine packets to and from a specific Ethernet 10BaseT port. Port monitoring has the following features:

- Software commands to enable and display captured port data.
- Captures data in Network General® file format.
- Limited protocol parsing (basic IP protocols and IPX) in console dump display.
- Data packets time stamped.
- One port monitored at a time.
- RAM-based file system.
- Memory buffer space from 1 MB to 8 MB.
- Statistics gathering and display
- Monitors only Ethernet 10BaseT ports
- Filtering limited to basic packet type—broadcast, multicast or unicast.

You can select to dump real-time packets to the terminal screen, or send captured data to a file. Once a file is captured, you can FTP it to a Sniffer for viewing.

Port Mirroring

An alternate method of monitoring ports is Port Mirroring, which allows a network engineer to attach a Sniffer to one Ethernet port and mirror traffic to and from any other Ethernet port. Port mirroring is described in *Port Mirroring* on page 19-57.

Port Monitoring Menu

The port monitoring commands are contained on the port monitoring menu, which is a sub-menu of the Networking menu. The port monitoring menu displays as follows:

<u>Command</u>	<u>Port Monitoring Menu</u>
pmon	Port monitor utility
pmcfg	Configure port monitor parameters
pmstat	View port monitor statistics
pmd	Port monitor disable
pmp	Port monitor pause
	Main File Summary VLAN Networking
	Interface Security System Services Help
	/Networking/Monitor %

The commands in this menu are described in the following sections.

RAM Disk System for Data Capture Files

Port monitoring uses a RAM disk for fast temporary storage of data capture files. The RAM disk has a separate directory designation of **/ram**. RAM-based files are created in DOS-FAT format and they are displayed in UPPERCASE.

You can copy files between the **/ram** disk system and the standard **/flash** file system. In addition, files in the RAM disk system are retrievable via FTP. Both the **/ram** file system and the **/flash** file system are accessible by using the UNIX/DOS-style change directory (**cd**) command.

◆ Note ◆

The RAM drive is part of DRAM memory. If you power off or reboot the switch, any files saved in the RAM drive will be lost.

Configuring RAM Drive Resources (pmcfg)

The **pmcfg** command allows you to select the size of the RAM disk file system or to delete the RAM disk. In addition, it allows you to configure the amount of data collected for each packet capture. To begin configuring RAM drive resources, enter

```
pmcfg
```

A screen similar to the following displays:

```
RAM disk size : 1000 Kilobytes
Lines displayed: 1
Change any of the above (y/n)? (n)
```

To change one of the settings, enter a **Y** and press **<enter>**. You will be prompted for a new RAM drive size. Select a size in kilobytes between 1000 and 8000. You can also delete the RAM drive by entering a size of zero (0). Changing the RAM disk size also requires that you reboot the system.

The **Lines displayed** controls the amount of data displayed to the terminal when you choose to dump session data to the computer screen. You can specify the number of lines to display while viewing port monitor data on the screen.

Changing the Default System Directory (cd)

After a port monitoring session is enabled the default directory is the RAM disk system (**/ram**). To switch back to the standard default flash file system (**/flash**) use the **cd** command. To switch back to the default directory, enter

```
cd /flash
```

To switch back to the RAM disk directory, enter

```
cd /ram
```

Starting a Port Monitoring Session (pmon)

You enable a port monitoring session through the **pmon** command. To start a session, enter **pmon** followed by the slot and port number that you want to monitor. For example, to monitor a port that is the first port in the fourth slot of the switch, you would enter

```
pmon 4/1
```

You can only monitor Ethernet 10BaseT ports. If a port is already being mirrored (enabled through the **addvp** or **modvp** command) you cannot monitor it. Also, you cannot set up more than one monitoring session on the same port.

If the port is currently being monitored, or mirrored, the following message displays:

```
Port 4/1 is being monitored.
Disable monitoring? (y)
```

If the port is not being monitored, or mirrored, the following message displays:

```
Port 4/1 is not being monitored, or mirrored.
Enable monitoring? (y)
```

Enter a **Y** and press **<enter>** at this prompt. The following screen of options displays:

```
Slot/Port           : 5/1
RAM disk size       : 1000 Kilobytes
Capture to filename : y
Capture filename    : PMONITOR.ENC
Dump to screen      : y
Broadcast frames    : y
Multicast frames    : y
Unicast frames      : y
Change any of the above (y/n)? (n) :
```

If you want to change any of the values, enter a **Y** and press **<enter>**. You will be prompted for all of the values in the screen except the **RAM disk size**, which you must change through the **pmcfg** command before starting the session. The information selected in this screen will be saved in flash configuration memory.

Enter any new values as prompted. The above screen re-displays to show the new values. Press **<enter>** to accept the updated values. Messages similar to the following display:

```
1048576 byte RAM drive /ram already initialized.
Bytes remaining on RAM disk = 1040384
```

The port monitoring session has begun. What happens at this point depends on whether you chose the **Dump to screen** option. The sections below describe what happens in each case.

◆ Note ◆

If you change the capture filename from the default, you must specify **/ram**. Otherwise, the file will be saved in the flash directory.

If You Chose *Dump to Screen*

If you selected the **Dump to screen** option, then a real-time synopsis of the session displays on your terminal screen. The following shows an example of this data

```
Enter 'p' to pause, 'q' to quit.
Destination | Source | Type | Data
-----|-----|-----|-----
00:20:DA:04:01:02 | 00:20:DA:04:01:01 | ICMP | 01:02:03:04:05:06:07:08
00:20:DA:04:01:02 | 00:20:DA:04:01:01 | ICMP | 01:02:03:04:05:06:07:08
FF:FF:FF:FF:FF:FF | 00:20:DA:02:10:E3 | ARP-C | 08:06:00:01:08:00:06:04
FF:FF:FF:FF:FF:FF | 00:20:DA:6F:97:A3 | RIP | 08:00:45:00:00:34:22:30
```

Each line in the display represents a packet. The destination MAC address, source MAC address, protocol type and actual packet data are shown. The amount of data shown is configured through the **pmcfg** command. The above sample shows 16 bytes of data per packet. You can stop the data dump to the screen at anytime by pressing **q** to quit. You can also pause the data dump by pressing **p** to pause.

If You Did Not Choose *Dump to Screen*

If you did not select the **Dump to screen** option, then the system prompt will return and port monitoring occurs in the background. You can continue using other UI commands. The port monitoring session data is saved in the file you indicated through the **pmmon** screen. You can monitor the session at anytime by using the **pmstats** command. You can also end or pause an in-progress session using the **pmdelete** or **pmpause** commands, respectively. The following sections describes **pmdelete** and **pmpause**.

Ending a Port Monitoring Session (**pmdelete**)

The **pmdelete** command ends a port monitoring data capture session that is being saved to file but not being dumped to the console screen. To end the session, enter:

```
pmd
```

A message similar to the following displays:

```
Port monitoring session terminated, data file is xxxxx.ENC.
```

If a port monitoring session was not in progress then the following message displays:

```
No ports being monitored.
```

Pausing a Port Monitoring Session (**pmpause**)

The **pmpause** command pauses a port monitoring data capture session that is being saved to file but not being dumped to the console screen. To pause the session, enter:

```
pmmp
```

The following message displays

```
Pausing monitor data capture/display.
```

To resume the port monitoring session, enter **pmmp** again. The following message displays:

```
Resuming monitor data capture.
```

If a port monitoring session was not in progress, then the following message would display:

```
No ports being monitored.
```

Ending a Port Monitoring Session

After you quit a port monitoring session, the default directory changes to **/ram** and the current files on the RAM drive are listed. The screen below shows an example of the display at the completion of a monitoring session.

```

Port monitoring capture done. Current capture files listed:
Current working directory '/ram'.

PM0302.ENC    65536  10/20/96 12:12
PM0303.ENC    32768  10/20/96 11:15

950272 bytes free

```

Viewing Port Monitoring Statistics (pmstat)

The **pmstat** command displays the statistics gathered for the current or most recent port monitoring session. If a port monitoring session is currently in progress, then it displays the results of the in-progress session. If a port monitoring session is not in progress, then it displays results of the most recently completed session. To view session statistics, enter

```
pmstat
```

A screen similar to the following displays:

```

Viewing capture statistics:
Percent RAM available: 96%
Frame type           #Frames
-----
Broadcast            108
Multicast             253
Unicast               301

```

The **Percent RAM available** indicates how much of the configured RAM disk has been used by this port monitoring session. You can configure the size of the RAM disk through the **pmcfg** command; the default size is 1 MB. The remaining items in the display show the number of packets passed on the port broken down into broadcast, multicast, and unicast frames.

Port Mapping

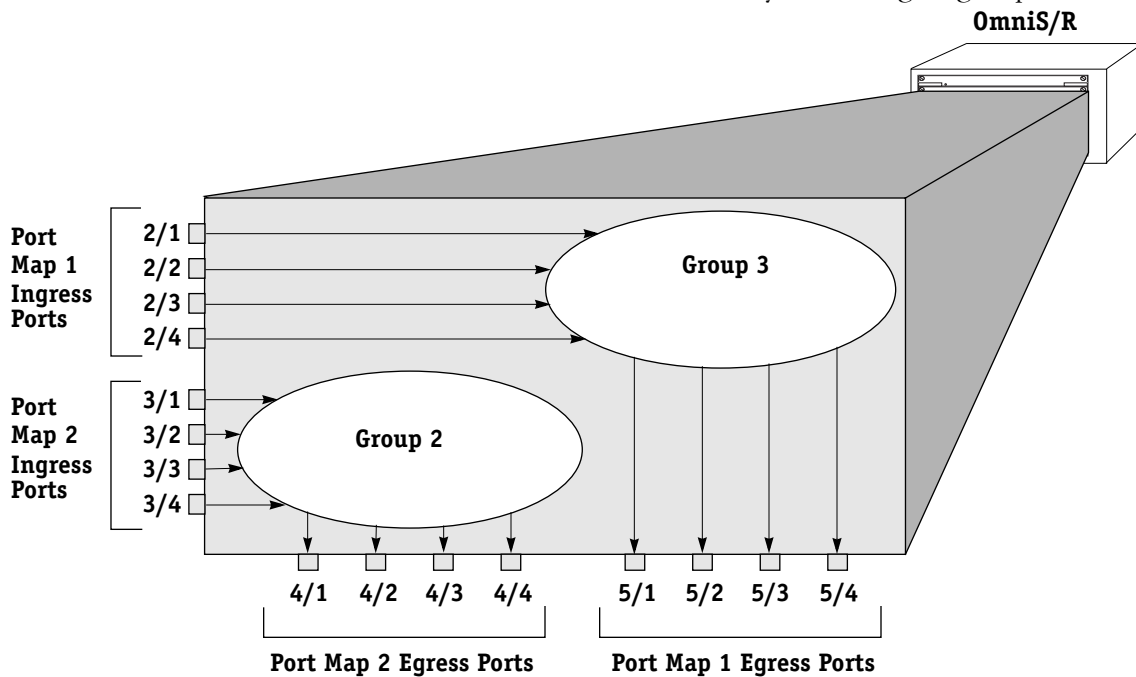
The OmniS/R began as an any-to-any switching device, connecting different LAN interfaces, such as Ethernet As networks grew and the traffic on them increased, a need arose for controlling some traffic, such as broadcasts. Virtual LANs, or VLANs, were introduced to segment traffic such that devices could only engage in switched communication with other devices in the same VLAN.

Some applications today require a further degree of traffic segmentation than that provided by VLANs. The port mapping feature allows you to further segment traffic *within* a VLAN or group by isolating a set of ports.

Groups/VLANs and Port Mapping

Port mapping does *not* affect existing group or AutoTracker VLAN operations in a switch. Group and VLAN membership are checked and applied before port mapping constraints are applied. Therefore, any constraints applied by port mapping only limit traffic flow *within* a group or VLAN; port mapping parameters do not provide any additional connectivity to a port. So if you add a port to a port mapping set, that port will be first subject to the constraints of its Group/VLAN and then the restrictions imposed by port mapping. Up to 128 port mapping sets can be configured per switch.

The illustration below helps show how group and port mapping constraints interact. The ports in slot 2 and 5 (2/1—2/4 and 5/1—5/4) are part of group 3. By group membership, all of these ports have switched communication with each other. Likewise, the ports in slot 3 and slot 4 have switched communication with each other as they all belong to group 2.



Groups and Port Mapping

Once a port mapping set is constructed, communication within each of the groups becomes more restricted. A port mapping set consists of *ingress* and *egress* ports; ingress ports can only send traffic to egress ports. In the above figure, all ports on slots 2 and 3 are ingress ports and ports on slots 4 and 5 are egress ports.

Port communication is uni-directional. A mapping between an ingress port and an egress port can only pass data from the ingress port to the egress port. To allow traffic to flow from the egress port to the ingress port, it is necessary to create a new mapping.

This configuration restricts each port to communication *only with the other four ports in the opposite port mapping subset within the same group*. For example, port 2/1 can only send traffic to ports 5/1, 5/2, 5/3, and 5/4. It can no longer communicate with ports 2/2, 2/3, and 2/4 even though they are part of the same group. Port mapping restricts ports from communicating with other ports within the same subset.

Port mapping does not affect other ports in the group that are not part of the port mapping set.

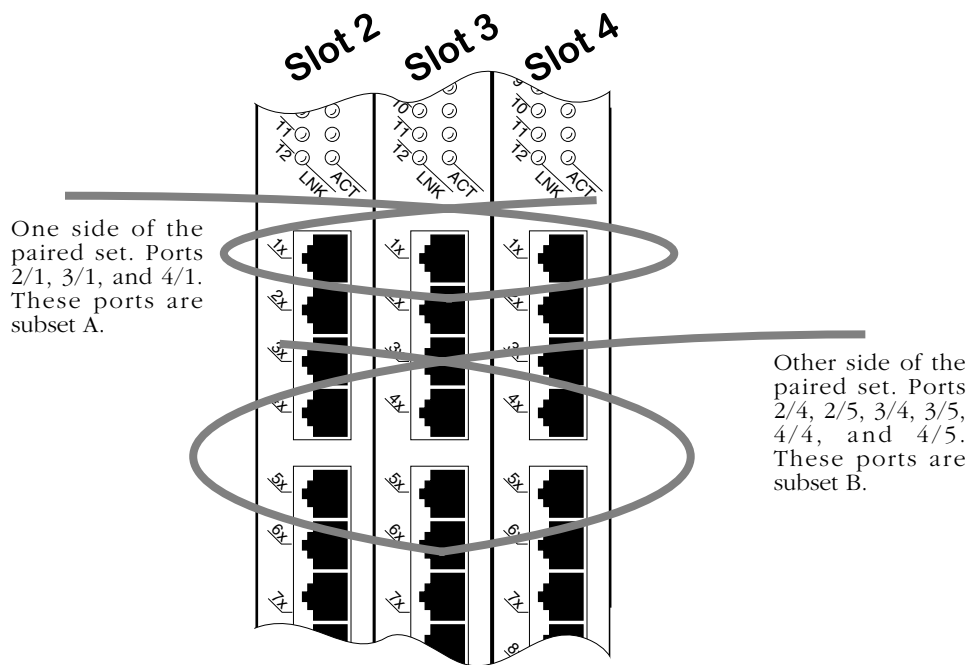
The Details of Port Mapping

Port mapping can be thought of as special rule that is applied after standard group and VLAN rules are applied. This rule statically assigns a port as either an ingress or egress port. Ingress ports can only communicate with egress ports. In this sense, one subset of ports is “mapped” to another subset of ports. Ports within the same subset can not communicate with each other or with another switch port that is not a member of the opposite port mapping subset.

◆ Note ◆

Port mapping restrictions are only applied to ports on 10/100 Ethernet modules (e.g., ESM-100F-8, ESM-C-32, ESM-FM-16W, ESM-100C-12).

As an illustration, see the diagram of three Ethernet modules below. The modules are in slots 2, 3, and 4. The ports that are circled are included in a port mapping subset. The three ports at the top—2/1, 3/1, and 4/1—are ingress ports. The six ports below—2/4, 2/5, 3/4, 3/5, 4/4, and 4/5—are egress ports in the port mapping set.



Port Subsets in the Port Mapping Set

Who Can Talk to Whom?

The following matrix outlines which ports can communicate with each other in the example shown on the previous page *assuming all ports are part of the same group or VLAN*. A port can only communicate with ports in the opposite subset within the port mapping set.

Switch Ports That May Communicate*

	2/1	2/4	2/5	3/1	3/4	3/5	4/1	4/4	4/5
2/1	N/A	Yes	Yes	No	Yes	Yes	No	Yes	Yes
2/4	No	N/A	No	No	No	No	No	No	No
2/5	No	No	N/A	No	No	No	Yes	No	No
3/1	No	Yes	Yes	N/A	Yes	Yes	No	Yes	Yes
3/4	No	No	No	No	N/A	No	No	No	No
3/5	No	No	No	No	No	N/A	No	No	No
4/1	No	Yes	Yes	No	Yes	Yes	N/A	Yes	Yes
4/4	Yes	No	No	Yes	No	No	Yes	N/A	No
4/5	Yes	No	No	Yes	No	No	Yes	No	N/A

***Read table from right (ingress ports) to left only.**

Port communication is uni-directional. A mapping between an ingress port and an egress port can only pass data from the ingress port to the egress port. To allow traffic to flow from the egress port to the ingress port, it is necessary to create a new mapping.

It's important to remember that the port mapping configuration is affected by existing group/VLAN rules. If the ports in the above example belonged to three groups based on IP network rules, then they would be restricted by group membership and port mapping.

Port mappings can be created between switch ports and uplink ports, but not between uplink ports. For example, you could map ethernet ports 3/1-12 to an WAN uplink port 4/1. This is useful when there is no traffic between ethernet ports, but all ports are to be forced to the uplink module. You *cannot*, however, map uplink port 4/1 to uplink port 4/2.

Port Mapping Limitations

The following are restrictions to the use of the port mapping feature:

- Port mapping cannot be used with ports assigned to an 802.1Q group.
- Port mapping cannot be used with an OmniChannel unless all ports in the OmniChannel are included in the port mapping (on either the ingress or egress list). For example, if ports 3/1-3/4 are an OmniChannel, all four ports must be in the ingress or egress list. You could not just map port 3/1.

Creating a Port Mapping Set

Use the **pmapcr** command to create a port mapping set. Follow these steps:

1. Enter **pmapcr** at a system prompt.
2. The following screen displays:

Port Map Configuration

```
1. Ingress List  :
2. Egress List   :
```

Enter the ingress ports and egress ports for this map set. This is done by entering the line number, an equal sign, and the port (or ports) to be added. For example, if you want to create a map set with an ingress port of 3/6 and an egress port of 4/6, you would enter the following at the prompt:

```
1=3/6
2=4/6
```

This must be done in two separate operations, one for the ingress and one for the egress lists. You can add more than one port to a list by using a comma (,) between slot/port designations, or a dash (-) between port numbers. For example, if you wanted to make ports 4/1, 4/6, 4/7, 4/8, and 4/9 egress ports for this map set, you would enter the following:

```
2=4/1, 4/6-9
```

A switch port in the ingress list can only communicate with switch ports in the egress list. Switch ports in the same list cannot communicate with each other or any other ports in the switch. For example, if you enter:

```
1=2/1, 3/1
2=2/2, 3/2
```

then you are creating a paired set of four ports. Port 2/1 can only communicate with ports 2/2 or 3/2. It cannot communicate with any other ports in the switch, including port 3/1. Port 3/1 also can only communicate with ports 2/2 and 3/2, but no others.

Any port type may be added to a port mapping set. However, only Mammoth-generation Ethernet ports will be restricted by port mapping limitations. For example, you could add a non-Ethernet port to the set, but traffic from that port would not be restricted.

3. You will want to save your configuration, so enter an **s** at the **port-mapping** prompt. Your configuration will be saved. A prompt similar to the following appears to confirm the creation of the port map:

```
Port Map 7 created.
```

The port map number is used when modifying the map set.

It is important to remember that port communication is uni-directional. A mapping between an ingress port and an egress port can only pass data from the ingress port to the egress port. To allow traffic to flow from the egress port to the ingress port, it is necessary to create a new mapping.

Adding Ports to a Port Mapping Set

You can add ports to a port map set once it has been created using the **pmapmod** command. Follow these steps:

1. Enter the **pmapmod** command at a system prompt, as shown:

```
pmapmod <pmap id>
```

where **<pmap id>** is the map set number shown when the map set was created. (To view a list of all existing map sets, see *Viewing a Port Mapping Set* on page 19-72.) For example, to modify map set 5, you would enter the following:

```
pmapmod 5
```

2. The following screen displays:

```

Port Mapping Configuration
=====
Port Map Id      Ingress Ports      Egress Ports
-----
5                3/1, 3/2, 3/3     4/1, 4/2, 4/3

Modify Port Map 5

1. Add Ports to Ingress List      :
2. Add Ports to Egress List      :
3. Delete Ports from Ingress List :
4. Delete Ports from Egress List  :
5. View Port Map Configuration   :
```

Note that the current ports in the port mapping set are displayed. Use this information to make decisions on the ports you want to add or remove from the set.

Enter the line number for the operation you want to perform (a **1** for the ingress list or a **2** for the egress list), an equal sign (=), and the ports to be added. For example, add port 3/2 to the ingress list and the egress list, enter the following (in two separate operations):

```
1=3/2
2=3/2
```

You can add more than one port to a list by using a comma (,) between slot/port designations, or a dash (-) between port numbers. For example, if you wanted to make ports 4/1, 4/6, 4/7, 4/8, and 4/9 egress ports for this map set, you would enter the following:

```
2=4/1, 4/6-9
```

3. To view the changes, enter a **5 (View Port Map Configuration)**, and equal sign (=), and a **y**, as shown:

```
5=y
```

This will refresh the Port Mapping Configuration screen and display any changes you have made.

4. Quit the session by entering a **q** at the prompt.

Removing Ports from a Port Mapping Set

You can remove ports to a port map set once it has been created using the **pmapmod** command. Follow these steps:

1. Enter the **modpmap** command at a system prompt, as shown:

```
pmapmod <pmap id>
```

where **<pmap id>** is the map set number shown when the map set was created. (To view a list of all existing map sets, see *Viewing a Port Mapping Set* on page 19-72.) For example, to modify map set 5, you would enter the following:

```
pmapmod 5
```

2. The Port Mapping Configuration screen displays (as shown above in *Adding Ports to a Port Mapping Set* on page 19-70).

Enter the line number for the operation you want to perform (a **3** for the ingress list or a **4** for the egress list), an equal sign (=), and the ports to be added. For example, remove port 3/2 to the ingress list and the egress list, enter the following (in two separate operations):

```
3=3/2  
4=3/2
```

You can remove more than one port to a list by using a comma (,) between slot/port designations, or a dash (-) between port numbers. For example, if you wanted to remove ports 4/1, 4/6, 4/7, 4/8, and 4/9 from the egress list of this map set, you would enter the following:

```
4=4/1, 4/6-9
```

3. To view the changes, enter a **5** (view port map configuration), an equal sign (=), and a **y**, as shown:

```
5=y
```

This will refresh the Port Mapping Configuration screen and display any changes you have made.

4. Quit the session by entering a **q** at the prompt.

Viewing a Port Mapping Set

You can view a port mapping set using the **vpmap** command. Enter the **pmapv** command as shown:

```
pmapv <pmap id>
```

where **<pmap id>** is the map set number shown when the map set was created. For example, to modify map set 5, you would enter the following:

```
pmapv 5
```

The following screen is shown:

```
Port Mapping Configuration
=====
Port Map Id      Ingress Ports      Egress Ports
-----
5                3/1, 3/2, 3/3     4/1, 4/2, 4/3
```

As a variation of this command, enter the **vpmap** command with no port map identification. This will display all port mapping sets configured for this switch.

Port Map Id. An identification number for the port map set, generated when the set is created.

Ingress Ports. The switch ports designated as ingress ports for this port map set. Ingress ports can only communicate with egress ports.

Egress Ports. The switch ports designated as egress ports for this port map set. Egress ports can only communicate with ingress ports.

Deleting a Port Mapping Set

You can delete a port mapping set after it is created. Enter **pmapdel** at a prompt as shown:

```
pmapdel <pmap id>
```

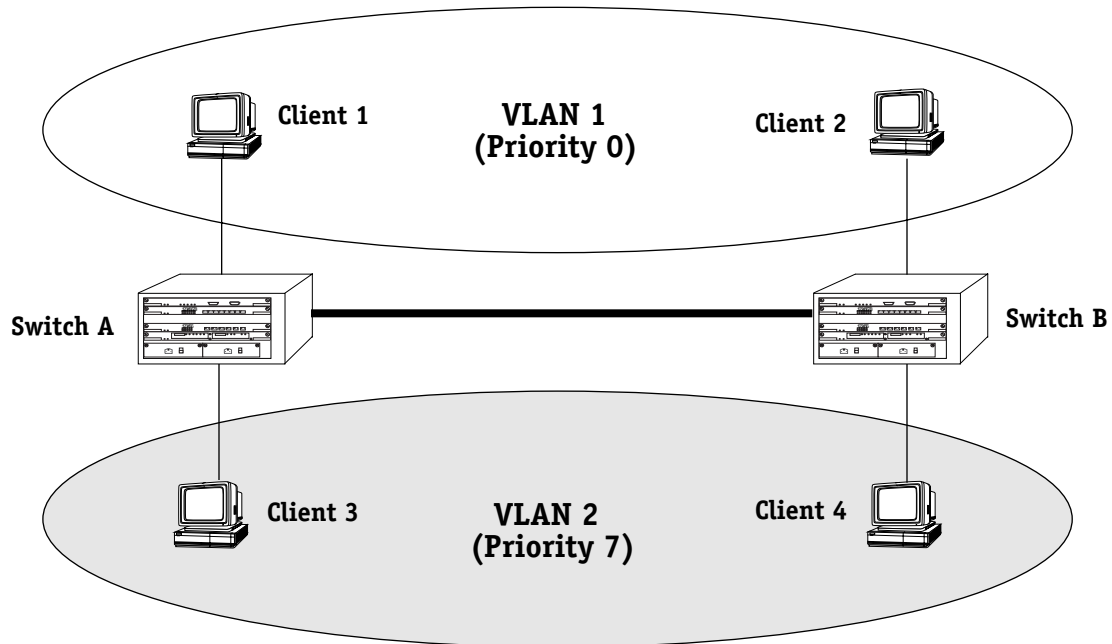
where **<pmap id>** is the map set number shown when the map set was created. (To view a list of all existing map sets, see *Viewing a Port Mapping Set* on page 19-72.) For example, to modify map set 5, you would enter the following:

```
pmapdel 5
```

Priority VLANs

Prioritizing VLANs allows you to set a value for traffic based on the destination VLAN of packets. Traffic with the higher priority destination will be delivered first. VLAN priority can be set from 0 to 7, with 7 being the level with the most priority.

The following diagram illustrates this idea:



In the above diagram, traffic from Client 3 in VLAN 2 (with a priority of 7) to Client 2 takes precedence over traffic from Client 1 in VLAN 1 (with a priority of 0) to Client 4.

Group priority can be set when creating a group using the **crgp** command. For more information on the **crgp** command, see *Creating a New Group* on page 19-18.

Group priority can be modified or viewed using the **prty_mod** and **prty_disp** commands, detailed below.

Mammoth vs. Kodiak Priority VLANs

Although the range of VLAN priority is 0-7, the Mammoth based modules only supports two levels of priority. In other words, 0-3 is one level and 4-7 is another. Future releases will expand the number of priority levels.

Kodiak based modules support up to 4 levels of priority (0-1, 2-3, 4-5, 6-7). **These two different implementations of the VLAN priority are not compatible.** Kodiak based priority VLANs can be used with other Kodiak based priority VLANs. This is true for Mammoth based VLANs as well.

Configuring VLAN Priority

To configure the priority of a VLAN:

1. Enter the **prty_mod** command at the system prompt, as shown:

```
prty_mod <groupid>
```

where **<groupid>** is the group number associated with the VLAN whose priority is being set. For example, to modify the priority of the VLAN for Group 2, you would enter the following:

```
prty_mod 2
```

The following prompt is shown:

```
Enter a priority value which is between 0 and 7: 0
```

2. Enter the number value that is to be the new priority level for this VLAN. The highest (most important) value is 7.
3. Press **<enter>**. A message similar to the following is displayed:

```
Priority for VLAN 2 has been set as 7
```

Viewing VLAN Priority

The priority level for all configured VLANs can be viewed by using the **prty_disp** command. Enter the **prty_disp** at the system prompt, as shown:

```
prty_disp <groupid>
```

where **<groupid>** is the group number associated with the VLAN whose priority is being viewed. For example, to view the priority of the VLAN for Group 2, you would enter the following:

```
prty_disp 2
```

A display similar to the following is shown:

```
The priority of group 2 is 7
```

As a variation of this command, you can enter **prty_disp** at the system prompt without a group number. This will display the priority of all VLANs.

20 Configuring Group and VLAN Policies

AutoTracker policies subdivide network traffic based on specific criteria. AutoTracker policies can be defined by port, MAC address, protocol, network address, user-defined, port binding, DHCP port, or DHCP MAC address policy. You can define multiple policies—also referred to as “rules”—for a mobile Group or an AutoTracker VLAN. A port or device is included in a mobile Group or AutoTracker VLAN if it matches any one AutoTracker rule. For example, you can define rules based on MAC address and rules based on protocol in the same mobile group or AutoTracker VLAN.

This chapter provides an overview of AutoTracker Policies as well as instructions for configuring these policies. AutoTracker policies may be applied to mobile groups (including authenticated groups) and to VLANs within standard groups. All policy types may be used with mobile groups and VLANs within standard Groups. However, only the Binding Rule may be used with authenticated groups.

◆ Note ◆

This chapter contains instructions for configuring AutoTracker policies for mobile groups or AutoTracker VLANs. Instructions for configuring groups (mobile and non-mobile) can be found in Chapter 19. More detailed overview and instructions for AutoTracker VLANs (created within non-mobile groups) can be found in Chapter 22.

AutoTracker policies enable you to control communications between end stations in your network. You define policies that determine membership in the mobile group or AutoTracker VLAN and AutoTracker automatically locates ports or devices that fit the policies and places them into the mobile group or AutoTracker VLAN.

You can define physical policies or logical policies (or combinations thereof) to determine membership. Physical policies consist of port rules: you define the members as one or more specific ports and membership is limited to the ports defined and the MAC addresses of devices connected to those ports.

Logical VLAN policies can consist of MAC address rules, protocol rules, network address rules, user-defined, or port binding rules. Ports are assigned to mobile groups or AutoTracker VLANs that have logical rules when the MPX module examines frames that originate from devices connected to the switch. If a frame is received that matches a logical rule, the source device's MAC address and the port to which the source device is connected are both made members.

The members of a mobile group or AutoTracker VLAN thus consist of source devices originating frames that fit the AutoTracker policies and the ports to which those source devices are connected.

AutoTracker Policy Types

You can define a maximum of 32 AutoTracker policies of each type per Group. There is no restriction on the number of rules you can define per AutoTracker VLAN, as long as the maximum number of policies for the Group is not exceeded. A port or device is included in a mobile group or AutoTracker VLAN if it matches any one rule.

You can define the following types of rules:

Port Policies. Port policies enable you to define membership on the basis of ports. Members of the mobile group or AutoTracker VLAN will consist of devices connected to specific ports on one switch or on multiple switches in the Group.

MAC Address Policies. MAC address policies enable you to define membership on the basis of devices' MAC addresses. This is the simplest type of rule and provides the maximum degree of control and security. Members of the mobile group or AutoTracker VLAN will consist of devices with specific MAC addresses. These devices may all be connected to one switch or they may be connected to different switches in the Group. A maximum of 1024 MAC addresses are supported per MAC address policy.

Protocol Policies. Protocol policies enable you to define membership on the basis of the protocol that devices use to communicate. All devices that communicate with the specified protocol become members of the mobile group or AutoTracker VLAN.

You can specify membership according to the following protocols: IP, IPX, AppleTalk, or DECNet. In addition, you can specify membership according to Ethernet type, source and destination SAP (service access protocol) header values, or SNAP (sub-network access protocol) type.

Network Address Policies. Network address policies enable you to define membership on the basis of network address criteria.

For example, you can specify that all IP users with a specific subnet mask be included in the mobile group or AutoTracker VLAN. Or, you can specify that all IPX users in a specific network address area using a certain encapsulation type be included.

If you define network address and port or protocol rules in the same VLAN, the network address rules will take precedence over the port and protocol rules should any conflict arise. To reverse this precedence (i.e., port and protocol rules take precedence over network address rules) you must add the following line to the switch's **mpx.cmd** file:

Precedence=0

User-Defined Policies. User-defined policies enable you to define membership on the basis of a specific pattern within a frame. All devices that originate frames containing this pattern are assigned to the mobile group or AutoTracker VLAN. The pattern is specified by defining an offset, a value, and a mask.

Port Binding Policies. A port binding policy specifies a particular device to be included in the mobile group or AutoTracker VLAN. There are six types of Port Binding Rules that can be created:

- Bind IP Address to a Port and a MAC address
- Bind MAC Address to a Protocol and a Port
- Bind Port to a Protocol
- Bind IP Address to a MAC Address
- Bind IP Address to a Port
- Bind MAC Address to a Port

You must specify a separate binding policy for each device, but you can specify an unlimited number of such policies. Binding policies take precedence over all other AutoTracker policies.

DHCP Port Policies. These policies are similar to standard port policies, but apply to switch ports to which DHCP client workstations are attached.

DHCP MAC Address Policies. These policies are similar to standard MAC address policies, but apply to the MAC addresses of DHCP client workstations only.

Defining and Configuring AutoTracker Policies

You can define AutoTracker policies by port, MAC address, protocol, network address, user definition, or port binding. You can define multiple policies for a mobile group or AutoTracker VLAN if you wish. A port or device is included in a mobile group or AutoTracker VLAN if it matches any one rule. For example, you can define rules based on ports, rules based on MAC address, and rules based on protocol in the same mobile group or AutoTracker VLAN. However, defining multiple rules is not trivial – exercise extreme care when you do so and make sure that you understand the consequences of your definitions. In most situations, it is advisable to use one of AutoTracker’s predefined rules.

The sections below provide directions for setting up each type of AutoTracker policy. Follow the directions for the policy you wish to set up.

Port Policy	See <i>Defining a Port Policy</i> on page 20-5.
MAC Address Policy	See <i>Defining a MAC Address Policy</i> on page 20-6.
Protocol Policy	See <i>Defining a Protocol Policy</i> on page 20-8.
Network Address Policy	See <i>Defining a Network Address Policy</i> on page 20-11.
User-defined Policy	See <i>Defining Your Own Rules</i> on page 20-13.
Binding Policy	See <i>Defining a Port Binding Policy</i> on page 20-15.
DHCP Port Policy	See <i>Defining a DHCP Port Policy</i> on page 20-20.
DHCP MAC Address Policy	See <i>Defining a DHCP MAC Address Policy</i> on page 20-21.

Where These Procedures Start

These policy configuration sections start in the middle of a sequence of steps with the **crgp** or **modatvl** commands. During the **crgp** command prompt sequence you can configure policies for mobile groups or for VLANs within non-mobile groups. The **modatvl** command contains an option for adding policies (option #3). The procedures in these sections pick up at the point after you choose to either to configure AutoTracker rules (**crgp**) or add more rules (**modatvl**).

Defining a Port Policy

After you enter the Administrative Status, the following menu displays:

- Select rule type:
1. Port Rule
 2. MAC Address Rule
 - 21) MAC Address Range Rule
 3. Protocol Rule
 4. Network Address Rule
 5. User Defined Rule
 6. Binding Rule
 7. DHCP PORT Rule
 8. DHCP MAC Rule
 - 81) DHCP MAC Range Rule

Enter rule type (1):

1. Press **<Return>**. If this is a VLAN in a non-mobile Group refer to Chapter 19 for a detailed explanation of the two ways port policies may be configured.

◆ **Note** ◆

As of the current release, the MAC Address Range Rule and DHCP MAC Range are not supported.

2. The following prompt displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the VLAN as it controls only to this specific rule within this specific VLAN. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

Enter the list of ports in Slot/Int/Service/Instance format:

Enter the physical ports that you want included in this VLAN. You may enter multiple ports at a time. Use the **<slot>/<port>** format. For example, to include port 7 from the module in slot 2, you would enter **2/7**. (The service and instance numbers are not necessary for specifying physical LAN ports. They are only necessary when specifying logical ports used over ATM, FDDI, and Frame Relay.)

4. The following prompt displays:

Configure more rules for this vlan (y/n):

You can set up multiple rules for the same VLAN. Enter a **Y** here if you want to set up more rules in addition to the port rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this VLAN. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

VLAN 1:2 created successfully

You are done setting up rules.

Defining a MAC Address Policy

After you enter the Administrative Status, the following menu displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule
```

Enter rule type (1):

1. Enter **2** and press **<Enter>**.
2. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for this mobile group or AutoTracker VLAN as it controls only to this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

```
Enter the list of MAC addresses (Enter save to end):
```

Enter the MAC addresses that you want to include in this VLAN. Separate addresses by a space. When you have entered the final MAC address, leave a space and type **save**.

4. The following prompt displays:

```
Configure more rules for this vlan (y/n):
```

You can set up multiple rules. Enter a **Y** here if you want to set up more rules in addition to the MAC Address rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the mobile group or AutoTracker VLAN was set up.

```
VLAN 1:2 created successfully
```

You are done setting up rules.

Defining a MAC Address Range Policy

After you enter the Administrative Status, the following menu displays:

```

Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule
  
```

Enter rule type (1):

1. Enter **21** and press **<Enter>**.
2. The following prompt displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for this mobile group or AutoTracker VLAN as it controls only to this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

Enter the lower end MAC addresses (AABBCC:DDEEFF) in canonical form followed by the higher end:

Enter the low end MAC address followed by the high end MAC address. Separate addresses by a space. The range is specified using the last two bytes of the MAC address.

When you have entered the high end MAC address press **<enter>**.

4. The following prompt displays:

Configure more rules for this vlan (y/n):

You can set up multiple rules. Enter a **Y** here if you want to set up more rules in addition to the MAC Address rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the mobile group or AutoTracker VLAN was set up.

VLAN 1:2 created successfully

You are done setting up rules.

◆ Note ◆

MAC range rules only apply to mobile groups. They cannot be configured for AutoTracker VLANs.

Defining a Protocol Policy

After you enter the Administrative Status for this mobile group or AutoTracker VLAN, the following menu displays:

- Select rule type:**
1. Port Rule
 2. MAC Address Rule
 - 21) MAC Address Range Rule
 3. Protocol Rule
 4. Network Address Rule
 5. User Defined Rule
 6. Binding Rule
 7. DHCP PORT Rule
 8. DHCP MAC Rule
 - 81) DHCP MAC Range Rule

Enter rule type (1):

1. Press **3** and press **<Enter>**.
2. The following prompt displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Indicate whether or not you want to enable this rule. Type **e** to enable or **d** to disable. If you enable the rule, the mobile group or AutoTracker VLAN will use it to determine membership of devices. If you disable the rule, then this rule will not be used in assigning devices, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the mobile group or AutoTracker VLAN as it controls only to this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

- Select Protocol:**
1. IP
 2. IPX
 3. DECNET
 4. APPLETALK
 5. Protocol specified by ether-type
 6. Protocol specified by DSAP and SSAP
 7. Protocol specified by SNAP

Enter protocol type (1):

Enter the number for the protocol that will be used to define this mobile group or AutoTracker VLAN. Numbers are listed next to the protocol names. By selecting a specific protocol, you are indicating that all traffic originating from network devices using that protocol will be assigned to this mobile group or AutoTracker VLAN. You can select the IP, IPX, DECNET, and APPLETALK protocols by entering 1, 2, 3, or 4, respectively.

◆ Please Take Note ◆

ARP (address resolution protocol) is included as IP. DDP (datagram delivery protocol) and AARP (AppleTalk ARP) are included as AppleTalk. DECNET is DECNET Phase IV traffic only.

If you want to define a protocol other than IP, IPX, AppleTalk, or DECNet, you can do so by specifying an Ethernet type, or by specifying source and destination SAP (service access protocol) header values, or by specifying a SNAP (sub-network access protocol) type. The following three sections describe how to specify these protocol types. If you are not specifying one of these special protocol types, continue with Step 4 below.

Protocol Specified by Ether-Type

- a. To specify a protocol by Ethernet type, enter **5** at the **Select Protocol:** menu. The following prompt displays:

Enter the Ether-type value in hex:

- b. Enter the desired Ethernet type in hex. You must enter two bytes of data. For example, enter 0800 to specify IP or enter 0806 to specify ARP. All devices that use the specified Ethernet type will be members of the mobile group or AutoTracker VLAN.
- c. Go on to Step 4 below.

Protocol Specified by DSAP and SSAP

- a. To specify a protocol by SAP (service access protocol) header, enter **6** at the **Select Protocol:** menu. The following prompt displays:

Enter the DSAP value in hex:

- b. Enter the destination service access protocol (DSAP) value in hex and press **<Enter>**. The following prompt displays:

Enter the SSAP value in hex:

- c. Enter the source service access protocol (SSAP) value in hex. Each entry must consist of one byte of data. All devices that use the specified source and destination SAP types will be members of the mobile group or AutoTracker VLAN.
- d. Go on to Step 4 below.

Protocol Specified by SNAP

- a. To specify a protocol by SNAP (sub-network access protocol) type, enter **7** at the **Select Protocol:** menu. The following prompt displays:

Enter the SNAP value in hex

- b. Enter the desired SNAP value in hex. You must enter five bytes of data. For example, enter 0000008137 to specify IPX SNAP or enter 00000080F3 to specify AppleTalk ARP SNAP. All devices that use the specified SNAP type will be members of the mobile group or AutoTracker VLAN.
- c. Go on to Step 4 below.

4. The following prompt displays:

Configure more rules for this vlan (y/n):

You can set up multiple rules for the same mobile group or AutoTracker VLAN. Enter a **Y** here if you want to set up more rules in addition to the protocol rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

VLAN 1:2 created successfully

You are done setting up rules for this mobile group or AutoTracker VLAN.

Defining a Network Address Policy

After you enter the Administrative Status for this mobile group or AutoTracker VLAN, the following menu displays:

```

Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule
  
```

Enter rule type (1):

1. Press **4** and press **<Enter>**.
2. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the mobile group or AutoTrackerVLAN as it controls only to this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

```

Select the Network Protocol:
1. IP
2. IPX
  
```

```
Enter the protocol type:
```

Enter the protocol for which you want to define this network address rule. Enter a **1** for IP and a **2** for IPX. The prompts that follow are different for IP and IPX. These differences are due to the different conventions used by the protocols for network address formats. Follow the procedure below the network protocol you are setting up.

Set Up an IP Address

- a. To specify an IP address, enter a **1** at the **Select the Network Protocol:** prompt.
- b. The following prompt displays:

```
Enter the IP address:
```

Enter the IP address that you want to include in this mobile group or AutoTracker VLAN. Enter the address in dotted decimal notation or hexadecimal notation (e.g., 198.206.181.10).

- c. The following prompt displays:

```
Enter the IP Mask (0xfffff00):
```

Enter the IP Subnet mask for this address. The default subnet mask is shown in parentheses and is automatically derived from the IP address class entered in Step b.

- d. Go on to Step 4 below.

Set Up an IPX Address

- a. To specify an IPX address, enter a **2** at the **Select the Network Protocol:** prompt.
- b. The following prompt displays:

Enter the IPX Network Number:

Enter an IPX network number to define the network devices you want included in the mobile group or AutoTracker VLAN. IPX addresses consist of eight hex digits and you can enter a minimum of one hex digit in this field. If you enter less than eight hex digits, the system prefixes your entry with zeros to create eight digits. All devices with the specified network number will be included in the mobile group or AutoTracker VLAN.

- c. The following prompt displays:

Select the IPX Network Encapsulation

1. Ethernet-II
2. IEEE 802.2 LLC
3. IEEE SNAP
4. IPX Proprietary

Enter the IPX Network Encapsulation (1):

Select the encapsulation type from the list. IPX devices do not know their network number at bootup. Typically, IPX servers assign different network numbers to devices using different encapsulation types within the same physical network. When an encapsulation type is specified here, an IPX device that does not know its network number at bootup will be assigned to the mobile group or AutoTracker VLAN as long as the device uses the encapsulation type you specify here.

- d. Go on to Step 4 below.
4. The following prompt displays:

Configure more rules for this vlan (y/n):

You can set up multiple rules for the same mobile group or AutoTracker VLAN. Enter a **Y** here if you want to set up more rules in addition to the Network Address rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this mobile group or AutoTracker VLAN. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

VLAN 1:2 created successfully

You are done setting up rules for this mobile group or AutoTracker VLAN.

Defining Your Own Rules

A user-defined rule enables you to include all devices in the mobile group or AutoTracker VLAN that originate frames containing a specified pattern at a specified location. Each user-defined rule requires an Offset, a Value, and a Mask; you will be prompted for each of these values. The Offset specifies the location of the pattern within the frame. The Value specifies the pattern. The Mask specifies the bits that you care about within the Value pattern.

After you enter the Administrative Status for this mobile group or AutoTracker VLAN, the following menu displays:

```

Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule
  
```

Enter rule type (1):

1. Enter **5** and press **<Return>**.
2. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the mobile group or AutoTracker VLAN as it controls only to this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

```
Enter the Offset into the frame ( < 64 ) :
```

Enter an **Offset** value, in number of bytes, to define the location where the **Value** – or pattern – is found. The offset value can be any number from 0 – 63. The first byte of the frame's MAC header is considered byte 1. An offset of 0 specifies that the pattern begins in byte 1 of the frame.

As an example, enter an offset value of **14** if you want to specify the pattern that defines NETBIOS, because that pattern begins in the 21st byte of the frame.

4. The following prompt displays:

```
Enter the value of the pattern to match:
```

Enter a **Value**, in hex, to specify the pattern itself. The value can be a maximum of eight bytes. For example, enter **FOFO** to specify the pattern that identifies NETBIOS.

5. The following prompt displays:

```
Enter the mask for the pattern to match:
```

Enter a **Mask** value, in hex, to specify the bits within the **Value** that you care about. The mask can be a maximum of eight bytes, but must be the same length as the **Value** you entered. The mask value is ANDed with the **Value** and frames are searched for the result.

Defining and Configuring AutoTracker Policies

For example, if you enter **FFEF** as the value and **FFFF** as the mask:

	<u>Hex</u>		<u>Binary</u>
Value=	FFEF	=	1111 1111 1110 1111
Mask=	FFFF	=	1111 1111 1111 1111

When a bit in the mask is set to 1, the corresponding bit of the value must be literal. When a bit in the mask is set to 0, the corresponding bit in the value is ignored and can be either a 0 or a 1. In the example above, since the mask is FFFF, all bits in the value must be literal and the actual pattern searched for is the binary value 1111 1111 1110 1111. Only devices that originate frames containing this binary value beginning at the 21st byte will be included in the mobile group or AutoTracker VLAN.

As a second example, if you enter FFEF as the pattern and FFF7 as the mask:

	<u>Hex</u>		<u>Binary</u>
Value=	FFEF	=	1111 1111 1110 1111
Mask=	FFF7	=	1111 1111 1111 0111

In this example, bits 0–2 and bits 4–15 of the value must be literal, since the corresponding bits in the mask are 1s. However, since bit 3 of the mask is a 0, bit 3 of the value can be either a 0 or a 1. Therefore, in this example, two actual binary patterns are searched for:

1111 1111 1110 1111 **or** 1111 1111 1110 0111

Devices originating frames containing either one of these binary values beginning at the 21st byte of the frame will be included in the mobile group or AutoTracker VLAN.

6. The following prompt displays:

Configure more rules for this vlan (y/n):

You can set up multiple rules for the same mobile group or AutoTracker VLAN. Enter a **Y** here if you want to set up more rules in addition to the Network Address rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this mobile group or AutoTracker VLAN. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

VLAN 1:2 created successfully

You are done setting up rules for this mobile group or AutoTracker VLAN.

Defining a Port Binding Policy

Port binding policies require devices to match two or three criteria. The criteria can be one of six combinations:

1. The device can attach to a specific switch port *and* use a specific MAC address *and* use a specific protocol (IP or IPX).
2. The device can attach to a specific switch port *and* use a specific MAC address *and* use a specific IP network address
3. The device can attach to a specific switch port and use a specific protocol (IP or IPX)
4. The device can use a specific IP address *and* use a specific MAC address
5. The device can use a specific port *and* a specific IP address
6. The device can use a specific port *and* a specific MAC address.

A device must match all values in the criteria set.

Port binding policies have two additional features. First, if a policy violation is detected, an SNMP trap is generated to alert the network manager which rule was violated. Secondly, if you attempt to configure a port binding rule that creates a conflict with another binding rule, an error message is generated to alert the user of the problem.

For example, if a port binding rule is created with a policy that links IP address 1.1.1.1 and MAC address aabbcc:ddeeff, and you attempt to create a port binding rule for the same IP address with a policy that links it to port 3/1, an error message will appear as shown:

This IP address has already been assigned to a different rule

In this example the second port binding rule is not created because the purpose of the first rule is to provide mobility for the IP address 1.1.1.1 (i.e., it is not restricted to a port), while the second rule specifically limits the mobility of IP address 1.1.1.1 to port 3/1.

A general rule for port binding policies is that once an address has been assigned (MAC or IP), it cannot be assigned to another policy until it is removed from the first policy. The following table is a reference for policy conflicts:

Limitations for Port Policies

	IP Address	MAC Address	Port	Protocol
IP Address	N/A	IP and MAC address cannot be used again	IP address cannot be used again	N/A
MAC Address	IP and MAC address cannot be used again	N/A	MAC address cannot be used again	MAC address cannot be used again
Port	IP address cannot be used again	MAC address cannot be used again	N/A	None
Protocol	N/A	MAC address cannot be used again	None	N/A

Defining and Configuring AutoTracker Policies

After you indicate you want to set up rules for this mobile Group or AutoTracker VLAN (using the **cratvl** command), the following menu displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule
```

Enter rule type (1):

1. Enter a 6 and press <Return>.
2. The following prompt displays:

```
Set Rule Admin Status to [(e)nable/(d)isable] (d) :
```

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the mobile group or AutoTracker VLAN as it applies only to this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

```
Please select one of the following bindings:
1. Bind IP Address to a Port and a MAC Address.
2. Bind MAC Address to a Protocol and a Port
3. Bind Port to a Protocol
4. Bind IP Address to a MAC Address
5. Bind IP Address to a Port
6. Bind MAC Address to a Port
Enter the type of binding (1) :
```

Enter the type of binding you want to use for this policy. Each binding policy specifies a particular device to be included in the mobile group or AutoTracker VLAN. Therefore, you must set up a separate binding policy for each device you want included in this mobile Group or AutoTracker VLAN.

You can bind a device's IP address to a switch port and a MAC address (select option 1), bind a device's MAC address to a protocol and a switch port (select option 2), bind a switch port to a specific protocol (select option 3), bind an IP address to a MAC address (select option 4), bind an IP address to a switch port (select option 5), or bind a MAC address to a switch port (select option 6).

◆ Note ◆

It is important to remember the line number of the binding policy you chose in order to follow the correct sequence for the remainder of these steps.

If you select option 1, 2, 3, 5, or 6, go to step 4. If you select option 4, go to step 5.

- The following prompt displays:

Enter the port in the form of slot/interface:

Enter the switch port to which this device must be attached. If the device is not attached to this port, it will not be included in this mobile Group or AutoTracker VLAN. You should first enter the slot for the module, then a slash (/), then the port number.

If you selected binding policy 1 or 5, then continue with step 5. If you selected binding policy 2 or 6, then continue with step 6. If you selected binding policy 3, then continue on with step 7.

- The following prompt displays:

Enter the IP address:

Enter the IP address for the device. If the device does not have this IP address, it will not be included in this mobile Group or AutoTracker VLAN.

If you selected binding policy 1 or 4, continue with step 6. If you selected binding policy 5, continue with step 8.

- The following prompt displays:

Enter the Canonical MAC address in AABBC:DDEEFF format:

Enter the MAC address for the device. If the device does not have this MAC address, it will not be included in this mobile Group or AutoTracker VLAN.

If you selected binding policy 1, 4, or 6, then continue with step 8. If you selected binding policy 2, then continue with step 7.

- The following prompt displays:

Select Protocol:

1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type
6. Protocol specified by DSAP and SSAP
7. Protocol specified by SNAP

Enter protocol type (1):

Enter the number for the protocol that will be used to define this binding policy. Numbers are listed next to the protocol names. By selecting a specific protocol, you are indicating that a device with the MAC address you specified previously that are attached to the switch port you specified previously, and with traffic using this protocol will be assigned to this mobile group or AutoTracker VLAN. You can select the IP, IPX, DECNET, and APPLE-TALK protocols by entering 1, 2, 3, or 4, respectively.

◆ **Note** ◆

ARP (address resolution protocol) is included as IP. DDP (datagram delivery protocol) and AARP (Apple-Talk ARP) are included as AppleTalk. DECNET is DECNET Phase IV traffic only.

If you want to define a protocol other than IP, IPX, AppleTalk, or DECNet, you can do so by specifying an Ethernet type, or by specifying source and destination SAP (service access protocol) header values, or by specifying a SNAP (sub-network access protocol) type. The following three sections describe how to specify these protocol types. If you are not specifying one of these special protocol types, continue with Step 8 below.

Protocol Specified by Ether-Type

- a. To specify a protocol by Ethernet type, enter **5** at the **Select Protocol:** menu. The following prompt displays:

Enter the Ether-type value in hex:

- b. Enter the desired Ethernet type in hex. You must enter two bytes of data. For example, enter 0800 to specify IP or enter 0806 to specify ARP. All devices that use the specified Ethernet type will be members of the mobile group or AutoTracker VLAN.
- c. Go on to Step 8 below.

Protocol Specified by DSAP and SSAP

- a. To specify a protocol by SAP (service access protocol) header, enter **6** at the **Select Protocol:** menu. The following prompt displays:

Enter the DSAP value in hex:

- b. Enter the destination service access protocol (DSAP) value in hex and press **<Enter>**. The following prompt displays:

Enter the SSAP value in hex:

- c. Enter the source service access protocol (SSAP) value in hex. Each entry must consist of one byte of data. All devices that use the specified source and destination SAP types will be members of the mobile group or AutoTracker VLAN.
- d. Go on to Step 8 below.

Protocol Specified by SNAP

- a. To specify a protocol by SNAP (sub-network access protocol) type, enter **7** at the **Select Protocol:** menu. The following prompt displays:

Enter the SNAP value in hex

- b. Enter the desired SNAP value in hex. You must enter five bytes of data. For example, enter 0000008137 to specify IPX SNAP or enter 00000080F3 to specify AppleTalk ARP SNAP. All devices that use the specified SNAP type will be members of the mobile group or AutoTracker VLAN.
- c. Go on to Step 8 below.

8. The following prompt displays:

Configure more rules for this vlan (y/n):

You can set up more devices for this binding policy group. Enter a **Y** here if you want to set up more devices. If you enter **Y**, you will be prompted for the next rule that you want to set up. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

VLAN 2:1 created successfully

You are done setting up rules for this mobile group or AutoTracker VLAN.

Defining a DHCP Port Policy

DHCP port polices simplify network configurations requiring DHCP clients and servers to be in the same mobile group or AutoTracker VLAN. You can see how DHCP port policies were used in an application example on page 20-27.

DHCP port policies differ fundamentally from standard port policies. In a standard port policy, the port is placed in the mobile group or AutoTracker VLAN as soon as the port rule is configured; no traffic on the port is required. A DHCP port rule *requires* traffic on the port in the form of a DHCP request packet before the port gains membership.

After you indicate you want to set up rules for this mobile Group or AutoTracker VLAN, the following menu displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule
```

Enter rule type (1):

1. Enter a 7 and press <Enter>.
2. The following prompt displays:

```
Set Rule Admin Status to [(e)nable/(d)isable] (d) :
```

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the mobile group or AutoTracker VLAN as it applies only to this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

```
Enter the list of ports in Slot/Int/Service/Instance format:
```

Enter the physical switch ports that you want included in this mobile Group or AutoTracker VLAN. You may enter multiple ports at a time. Use the <slot>/<port> format. For example, to include port 7 from the module in slot 2, you would enter **2/7**. (The service and instance numbers are not necessary for specifying physical LAN ports. They are only necessary when specifying logical ports used over ATM, FDDI, and Frame Relay.)

4. The following prompt displays:

```
Configure more rules for this vlan [y/n] (n) :
```

You can set up multiple rules for the same mobile group or AutoTracker VLAN. Enter a **Y** here if you want to set up more rules in addition to the port rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up. Follow the directions in the appropriate section to configure that rule. If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

```
VLAN 1:2 created successfully
```

You are done setting up rules for this mobile group or AutoTracker VLAN.

Defining a DHCP MAC Address Policy

You can see how DHCP MAC address policies were used in an application example on page 20-27.

After you enter the Administrative Status for this mobile group or AutoTracker VLAN, the following menu displays:

```

Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule

```

Enter rule type (1):

1. Enter **8** and press **<Enter>**.
2. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the mobile group or AutoTracker VLAN as it applies only this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

```
Enter the list of MAC addresses (AABBCC:DDEEFF) in Canonical format
(Enter save to end):
```

Enter the MAC addresses that you want to include in this mobile group or AutoTracker VLAN. Separate addresses by a space. When you have entered the final MAC address, leave a space and type **save**.

4. The following prompt displays:

```
Configure more rules for this vlan (y/n):
```

You can set up multiple rules for the same mobile group or AutoTracker VLAN. Enter a **Y** here if you want to set up more rules in addition to the port rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

```
VLAN 1:2 created successfully
```

You are done setting up rules for this mobile group or AutoTracker VLAN.

Defining a DHCP MAC Address Range Policy

You can see how DHCP MAC address policies were used in an application example on page 20-27.

After you enter the Administrative Status for this mobile group or AutoTracker VLAN, the following menu displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
   21) MAC Address Range Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP MAC Rule
   81) DHCP MAC Range Rule
```

Enter rule type (1):

1. Enter **81** and press **<Enter>**.
2. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Indicate whether or not you want to enable the Administrative Status for this rule. Type **e** to enable or **d** to disable. If you enable the rule, the switch will use it to determine membership of devices. If you disable the rule, then the switch will not use this rule, but the parameters you set up will be saved. The Admin Status for a Policy is different from the Admin Status for the mobile group or AutoTracker VLAN as it applies only this specific rule. You can enable or disable the rule at a later time using the **modatvl** command.

3. The following prompt displays:

```
Enter the lower end DHCP MAC addresses (AABBCC:DDEEFF) in canonical form
followed by the higher end:
```

Enter the low end DHCP MAC address followed by the high end DHCP MAC address. Separate addresses by a space. The range is specified using the last two bytes of the MAC address.

When you have entered the high end MAC address press **<Enter>**.

4. The following prompt displays:

```
Configure more rules for this vlan (y/n):
```

You can set up multiple rules for the same mobile group or AutoTracker VLAN. Enter a **Y** here if you want to set up more rules in addition to the port rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

```
VLAN 1:2 created successfully
```

You are done setting up rules for this mobile group or AutoTracker VLAN.

◆ Note ◆

MAC range rules only apply to mobile groups. They cannot be configured for AutoTracker VLANs.

Viewing Mobile Groups and AutoTracker VLANs

You can view the current status of all mobile groups or AutoTracker VLANs in the switch using the **atvl** command. Enter **atvl** and a table similar to the following displays.

VLAN Group :	VLAN Id	VLAN Description	Admin Status	Operational Status
	6	New Mobile Group 6	Enabled	Active
	8	New Mobile Group 8	Enabled	Active

VLAN Group. The Group to which this AutoTracker VLAN is assigned. The Group is specified when first creating an AutoTracker VLAN.

VLAN ID. An identification number that you assigned when you created this VLAN. A value will not display in this column for mobile groups.

VLAN Description. A textual description that you entered to describe a VLAN when you created or modified it through **cratvl** or **modatvl**. This description is limited to 30 characters.

Admin Status. The Administrative Status for the VLAN may be enabled or disabled. You enable or disable the Administrative Status for a VLAN when you create or modify it. If the VLAN is enabled, the switch will use the policies you configured to filter traffic to the devices in this VLAN. If you disable the rule, then policies will not be used, but the parameters you set up for the VLAN will be saved.

Oper Status. The VLAN is shown as **Active** or **Inactive**. In order for an enabled VLAN to become “active” it must be able to assign a switch port to the VLAN. If the port rule is used for a VLAN, then the VLAN automatically becomes active. If any other rule is used (MAC address, protocol, etc.), then a frame matching the VLAN rule must first be received by a switch port before the VLAN is active. So, an Active VLAN requires the following:

- Admin Status must be enabled.
- A port must be assigned to the VLAN through either a port-based rule or by a device transmitting data that matches the VLAN policy.

Viewing Policy Configurations

Typing **viatr1** brings up the Policy Configuration Table, which shows the policies defined for the mobile Group or VLAN specified.

VLAN Group :	VLAN Id	Rule Num	Rule Type	Rule Status	Rule Definition
3:	5	1	PORT RULE	Disabled	2/7/Brg/1
3:	11	1	NET ADDR RULE	Enabled	IPX Addr = 11223344 IPX Encapsulation = Ethernet
3:	12	1	NET ADDR RULE	Enabled	DECNET Area = 13579
3:	22	1	PORT RULE	Enabled	2/7/Brg/1
3:	23	1	PORT RULE	Enabled	2/7/Brg/1
3:	24	1	MAC RULE	Enabled	082008:003002 082009:803728
3:	25	1	PROTOCOL RULE	Enabled	Protocol = IP
3:	26	1	NET ADDR RULE	Enabled	IP Addr = 131.1.2.3 IP Mask = 255.255.0.0
3:	27	1	USER RULE	Enabled	Offset = 64 Length = 2 Value = FFFF Mask = FFFF
3:	31	1	PROTOCOL RULE	Enabled	Protocol = IP
3:	32	1	NET ADDR RULE	Enabled	IPX Addr = 00000001 IPX Encapsulation = Ethernet

VLAN Group. The Group to which this AutoTracker VLAN is assigned. The Group number is specified when first creating the VLAN.

VLAN ID. An identification number that you assigned when you created this virtual LAN. A value will not display in this column for mobile groups.

Rule Num. The number of the policy within the VLAN definition. Each rule defined for a VLAN is numbered sequentially in the order of creation. The rule number is needed when you want to modify or delete a rule definition.

Rule Type. The type of VLAN policy. The Rule Type can be a port policy (PORT RULE), MAC Address policy (MAC RULE), network address policy (NET ADDR RULE), Protocol policy (PROTOCOL RULE), a user-defined policy (USER RULE), port-binding policy (BIND RULE), DHCP Port policy (DHCP PORT RULE), or a DHCP MAC address policy (DHCP MAC RULE). You set up VLAN policies when you create or modify the VLAN.

Rule Status. Indicates whether the rule for this row is Enabled or Disabled. If the rule is enabled, then the VLAN is using the rule definition to determine VLAN membership. If Disabled, then the VLAN is not using this rule to determine membership. Note that this Rule Status is different from the Admin Status for the VLAN since it controls only this specific rule within this specific VLAN. You can enable or disable the rule using the **modatvl** command.

Rule Definition. Details of this rule. For a Port Rule, this column lists the virtual interface for the Port included in the VLAN as

<slot>/<port>/<service>/<instance>

For example, the port defined for the first row in the table applies to the first bridge instance on port 7 on the module in slot 2 of the switch. For a MAC address rule, this column lists the MAC address for the device in the VLAN. For a Network Address Rule, the column will list the address (IP or IPX) and the IP Mask (IP) or the Encapsulation type (IPX). For a Protocol policy, the column list the protocol used to determine membership. And in a User-Defined rule, the offset, length, value, and mask are listed.

Viewing Virtual Ports' Group/VLAN Membership

You can view the VLAN membership of each virtual interface in the switch. For physical LAN ports, the virtual interface is the same as a virtual port. However, when multiple services are set up for a physical port, then each service has a virtual port.

Type **vi** and a Virtual Interface Table displays similar to the one that follows. You can also specify just the slot and port number to narrow the range of ports displayed.

Virtual Interface VLAN Membership

Slot/Intf/Service/Instance	Group	Member of VLAN#
1 /1 /Rtr /1	1	1
1 /1 /Rtr /2	3	1
1 /1 /Rtr /3	3	23
1 /1 /Rtr /4	3	24
1 /1 /Rtr /5	3	25
1 /1 /Rtr /6	3	5
2 /1 /Brg /1	1	1
2 /2 /Brg /1	1	1
2 /3 /Brg /1	1	1
2 /4 /Brg /1	1	1
2 /5 /Brg /1	1	1
2 /6 /Brg /1	1	1
2 /7 /Brg /1	1	1 22
2 /8 /Brg /1	1	1
3 /1 /Brg /1	1	1
4 /1 /Brg /1	1	1
4 /2 /Brg /1	1	1
4 /3 /Brg /1	1	1
4 /4 /Brg /1	1	1
4 /5 /Brg /1	1	1
4 /6 /Brg /1	1	1
5 /1 /Brg /1	1	1

Slot/Intf/Service/Instance. Specifies the virtual interface for which AutoTracker VLAN information will be displayed. The **Slot** is the physical slot location to which the virtual interface maps. The **Intf** is the physical port to which the virtual interface maps. The **Service** is the service type for this interface. The service type may be a Router (**Rtr**), Bridge (**Brg**), Classical IP (**CIP**), FDDI Trunk (**Trk**), or an 802.10 Trunk (**T10**). **Instance** is the specific instance of this service type. These different instances are identified numerically. The first instance of a service type belonging to a physical port is identified as 1, the second instance is identified as 2, etc.

Group. The Group to which this virtual interface is assigned. The Group is specified when first creating an AutoTracker VLAN.

Member of VLAN #. The AutoTracker VLANs to which this virtual interface belongs. An interface may belong to more than one VLAN. For example, a port may contain devices using the IP Protocol and could match the Port policy of one AutoTracker VLAN and the Protocol policy of another AutoTracker VLAN. Also, physical ports always remain members of the default VLAN #1.

View VLAN Membership of MAC Devices

The **fwtl** command displays a table of learned MAC addresses and the VLAN membership of those MAC addresses. Follow these steps to view this table.

1. Enter **fwtl**.
2. The following prompt displays:

Enter Slot/Interface (return for all ports) :

Enter the slot and port for which you want to view MAC Address/VLAN information. You can also press **<Enter>** to view information on all ports in the switch.

3. The following message and prompt displays:

Total number of MAC addresses learned for Group 1: 4
Maximum number of entries to display [20] :

The top line displays the number of MAC addresses learned on this switch. This number indicates the potential number of entries you can display in the Learned MAC Address Table. The second line allows you to indicate how many of these MAC addresses you want to display. Enter the number of MAC entries you want to display or press **<Enter>** to select the default in brackets [20].

4. The Learned MAC Address/VLAN Membership Table displays as follows:

MAC Address	Slot/Intf/Service/Instance	AT VLAN Membership
0020DA:05F623	4/ /1 /Brg 1	1
0020DA:021533	4/ /1 /Brg 1	1
0020DA:0205B3	4/ /1 /Brg 1	1
0020DA:06BAD3	4/ /1 /Brg 1	1
0020DA:05F610	4/ /1 /Brg 1	1

MAC Address. The MAC address for which virtual interface and VLAN membership information will be displayed.

Slot/Intf/Service/Instance. Specifies the virtual port for which AutoTracker VLAN information will be displayed. The **Slot** is the physical slot location to which the MAC address maps. The **Intf** is the physical port to which the MAC address maps. The **Service** is the service type for this MAC address. The service type may be a Router (**Rtr**), Bridge (**Brg**), Classical IP (**CIP**), FDDI Trunk (**Trk**), or an 802.10 Trunk (**T10**). **Instance** is the specific instance of this service type. These different instances are identified numerically. The first instance of a service type belonging to a physical port is identified as 1, the second instance is identified as 2, etc.

AT VLAN Membership. The AutoTracker VLANs to which this MAC Address belongs. An MAC address may belong to more than one VLAN. For example, let's say a MAC device runs on an IPX network. It could be included in a MAC Address policy for one AutoTracker VLAN and the IPX Protocol Policy of another VLAN.

Application Example: DHCP Policies

This application example shows how Dynamic Host Configuration Protocol (DHCP) port and MAC address policies can be used in a DHCP-based network. DHCP is built on a client-server model in which a designated DHCP server allocates network addresses and delivers configuration parameters to dynamically configured clients.

Since DHCP clients initially have no IP address, placement of these clients in an AutoTracker VLAN presents a problem. AutoTracker determines VLAN membership by looking at traffic from source devices. Since the first traffic transmitted from a source DHCP client does not contain the actual address for the client (because the server has not allocated the address yet), the client may not be placed in the same VLAN as its server.

Before the introduction of DHCP port and MAC address rules, various strategies were deployed to use DHCP with Groups and VLANs. Typically these strategies involved IP protocol and network rules along with Bootp relay functionality. (See Chapter 24 for some application examples of these strategies.) These solutions required that all DHCP clients in a particular mobile group or VLAN be grouped together through a common IP policy.

DHCP port and MAC address rules simplify the configuration of DHCP networks. Instead of relying on IP-based policies to group all DHCP clients in the same network as a DHCP server, you can manually place each individual DHCP client in the VLAN or mobile group of your choice. DHCP port and MAC address policies operate the same way as standard port and MAC address policies except these new rules have been enhanced for use with DHCP clients.

The VLANs

This application example contains three (3) AutoTracker VLANs within a single non-mobile group. These VLANs are called Test, Production, and Branch.

The Test VLAN connects to the main network, the Production VLAN, through an external router. This VLAN is intended to be self-contained such that copies of it could be made and attached to the Production VLAN in the same way this VLAN does. The Test VLAN contains its own DHCP server and DHCP clients. The clients gain membership to the VLAN through DHCP port rules.

The Production VLAN carries most of the traffic in this network. It does not contain a DHCP server, but does contain DHCP clients that gain membership through DHCP port rules. Two external routers connect this VLAN to the Test VLAN and a Branch VLAN. One of the external routers—the one connected to the Branch VLAN—has Bootp relay functionality enabled. It is through this router that the DHCP clients in the Production VLAN access the DHCP server in the Branch VLAN.

The Branch VLAN contains a number of DHCP client stations and its own DHCP server. The DHCP clients gain membership to the VLAN through both DHCP port and MAC address rules. The DHCP server allocates IP addresses to all clients in this VLAN as well as the DHCP clients in the Production VLAN.

DHCP Servers and Clients

DHCP clients must be able to communicate with a DHCP server at initialization. The most reliable way to ensure this communication is for the server and its associated clients to share the same VLAN or mobile group. However, if the network configuration does not lend itself to this solution (as the Production VLAN does not in this application example), then the server and clients can communicate through a router with Bootp relay enabled.

The DHCP servers and clients in this example are either in the same VLAN or are connected through a router with Bootp relay. All clients in the Test VLAN receive IP addresses from the server in their VLAN (Server 1). Likewise, all clients in the Branch VLAN receive IP addresses from their local server (Server 2). The DHCP clients in the Production VLAN do not have a local DHCP server, so they must rely on the Bootp relay functionality in external Router 2 to obtain their IP addresses from the DHCP server in the Branch VLAN.


Both DHCP servers gain membership to their VLANs through IP network address policies.

The following table summarizes the VLAN architecture and policies for all devices in this network configuration. The diagram on the following page illustrates this network configuration.

Devices and VLAN Membership

Device	VLAN Membership	Policy Used/Router Role
DHCP Server 1	Test VLAN	IP subnetwork rule=10.15.X.X
DHCP Server 2	Branch VLAN	IP subnetwork rule=10.13.X.X
External Router 1	Test VLAN Production VLAN	Connects Test VLAN to Production VLAN
External Router 2	Production VLAN Branch VLAN	Bootp relay provides access to DHCP server in Branch VLAN for clients in Production VLAN.
DHCP Client 1	Test VLAN	DHCP Port Rule
DHCP Client 2	Test VLAN	DHCP Port Rule
DHCP Client 3	Production VLAN	DHCP Port Rule
DHCP Client 4	Production VLAN	DHCP Port Rule
DHCP Client 5	Branch VLAN	DHCP Port Rule
DHCP Client 6	Branch VLAN	DHCP Port Rule
DHCP Client 7	Branch VLAN	DHCP MAC Address Rule
DHCP Client 8	Branch VLAN	DHCP MAC Address Rule

DHCP Port and MAC Rules



DHCP Servers
 Both DHCP servers become members in their respective VLANs via IP subnet rules.

Routers

Router 1 provides connectivity between the Test VLAN and the Production VLAN. It does not have Bootp functionality enabled so it cannot connect DHCP servers and clients from different VLANs.

Router 2 connects the Production VLAN and the Branch VLAN. With Bootp relay enabled, this router can provide connectivity between the DHCP server in the Branch VLAN and the DHCP clients in the Production VLAN.

DHCP Clients

Clients 1 to 6 are assigned to their respective VLANs through DHCP port rules. Clients 3 and 4 are not in a VLAN with a DHCP server so they must rely on the server in the Branch VLAN for initial addressing information. Clients 7 and 8 share a port with other devices, so they are assigned to the Branch VLAN via DHCP MAC address rules.

21 Interswitch Protocols

This chapter describes Interswitch Protocols, which are used to discover adjacent switches, and track VLAN membership and retain mobile group information across switches. They include two new protocols and one existing protocol that is updated for release 4.0:

- Mapping Adjacency Protocol (XMAP), a new protocol used to discover the topology of Omni Switch/Routers (Omni S/Rs)
- Group Mobility Advertisement Protocol (GMAP), a new protocol used to retain learned mobile group and protocol information
- VLAN Advertisement Protocol (VAP), an existing interswitch protocol used to exchange VLAN information between switches

The protocols are independent of each other and perform separate functions. Each protocol is described in detail in separate sections of this chapter.

Interswitch Protocol Commands

There is an Interswitch Protocol (XIP) submenu. Select **XIP** from the AutoTracker submenu, and the submenu displays as follows:

<u>Command</u>	<u>XIP Menu</u>
gmapst	Turn Group Mobility Advertisement Protocol (GMAP) ON or OFF
gmapgaptime	Set GMAP inter-message gap time in milliseconds
gmapholdtime	Set GMAP hold time interval time in minutes
gmapupdtime	Set GMAP update interval time in seconds
vlap	Turn VLAN Advertisement Protocol (VAP) ON or OFF
xmapst	Turn the Xylan Mapping Adjacency Protocol (XMAP) ON or OFF
xmapls	List adjacent switches found using the XMAP protocol
xmapdisctime	Set XMAP message interval for discovery phase in seconds
xmapcmntime	Set XMAP message interval for common phase in seconds

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

These commands are described in this chapter.

XMAP

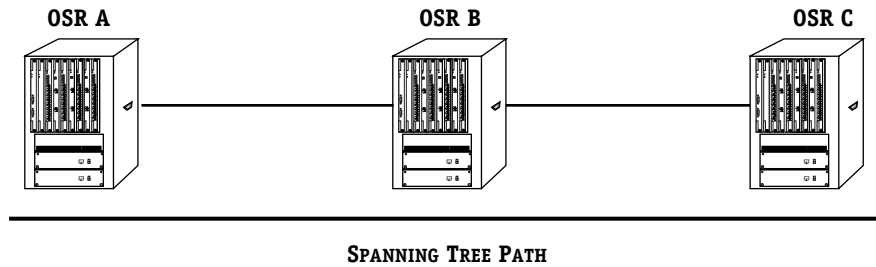
The Mapping Adjacency Protocol (XMAP) is used to discover the topology of OmniS/Rs in a particular installation. Using this protocol, each switch determines which OmniS/Rs are adjacent to it by sending and responding to Hello update packets. For the purposes of XMAP, *adjacent* switches are those that:

- have a Spanning Tree path between them
- do not have any switch between them on the Spanning Tree path that has XMAP enabled

◆ **Note** ◆

XMAP replaces the Adjacency Only mode of earlier versions of VAP.

In the illustration here, all switches are on the Spanning Tree path. Omni Switch/Router A and Omni Switch/Router C have XMAP enabled. Omni Switch/Router B does not. Omni Switch/Router A is adjacent to Omni Switch/Router C and vice versa. If Omni Switch/Router B enables XMAP, the adjacency changes. A would be adjacent to B, B would be adjacent to both A and C, and C would be adjacent to B.



XMAP Adjacency


XMAP Transmission States

XMAP switch ports are either in the *discovery transmission state*, *common transmission state*, or *passive reception state*. Ports transition to these states depending on whether or not they receive Hello responses from adjacent switches.

◆ Note ◆

All Hello packet transmissions are sent to a well-known MAC address (0020DA000003).

The transmission states are illustrated here.



XMAP State Transitions

Discovery Transmission State

When XMAP is active, at startup all active switch ports are in the discovery transmission state. In this state ports send out Hello packets and wait for Hello responses. Ports send out Hello packets at a configurable interval called the *discovery transmission time*. The discovery transmission time is configurable; 30 seconds is the default. The ports send out Hello packets up to *three* timeouts of this interval trying to discover adjacent switches.

Any switch ports that receive Hello packets before three discovery transmission times expire send a Hello reply and transition to the common transmission state. Any switch ports that do not receive a Hello response before three discovery transmission times have expired are placed in the passive reception state.

Common Transmission State

In the common transmission state, ports detect adjacent switch failures or disconnects by sending Hello packets and waiting for Hello responses. Ports in this state send out Hello packets at a configurable interval (the default is 5 minutes) called the *common transmission time*. To avoid synchronization with adjacent switches, the common transmission time is jittered randomly by plus or minus ten percent.

Ports wait for Hello responses using the *discovery transmission time* (the default is 30 seconds). If Hello responses are detected within one discovery transmission time, the port remains in the common transmission state. If Hello responses are not detected within one discovery transmission time, the port reverts to the discovery state.

Passive Reception State

In the passive reception state, switch ports are in receive-only mode. Hello packets are not sent out from these ports, and there is no timer on waiting for Hello responses. If the port receives a Hello packet at any time, it enters the common transmission state and transmits a Hello packet in reply.

If a port transitions to the passive reception state, any remote switch entries for that port are deleted.

Common Transmission and Remote Switches

If an XMAP switch is connected to multiple XMAP switches via a hub, the switch sends and receives Hello traffic to and from the remote switches through the same port. If one of the remote switches stops sending Hello packets and other remote switches continue to send Hello packets, the ports in the common transmission state will remain in the common transmission state.

The inactive switch will eventually be aged out of the switch's XMAP database because each remote switch entry has a "last seen" field that is updated when Hello packets are received. The switch checks the "last seen" field at least once every common transmission interval. Switch ports that are no longer "seen" may still retain an entry for up to three common transmission intervals. The slow aging out prevents the port from sending Hello packets right away to the inactive switch and creating additional unnecessary traffic.

Configuring XMAP

XMAP is active by default. In addition to disabling or enabling XMAP, you can view a list of adjacent switches or configure the timeout intervals for Hello packet transmission/reception.

Enabling or Disabling XMAP

To display whether or not XMAP is active or inactive, or to activate or deactivate XMAP, enter the following command:

```
xmapst
```

A screen displays similar to the following:

```
XMAP is currently ACTIVE. (a)ctivate, (d)e-activate : (a) :
```

Enter **a** or **d** to change the current state, or press **<Enter>** to keep the current value. A message similar to the following displays:

```
XMAP is ACTIVE.
```

To change the state of XMAP without displaying the current state first, enter the command with the desired value. For example:

```
xmapst d
```

A message similar to the following displays:

```
XMAP is INACTIVE.
```


Viewing a List of Adjacent Switches

Use the **xmapls** command to view a list of adjacent switches and their associated MAC addresses, ports, groups, and IP addresses. For remote switches that stop sending Hello packets *and* are connected via a hub, entries may take up to three times the common transmission interval to age out of this table.

The example display shows three virtual ports on a local XMAP switch connected to remote virtual ports on two switches. VPN 3 is connected to a remote switch through a hub.

VPN	Rem Switch ID	Rem VPN	Pri Group	IP Addresses
=====	=====	=====	=====	=====
1	0020da:032c40	1	2	18.1.1.1 27.0.0.2 192.168.10.1 198.206.184.40
3	0020da:999660	1	2	192.168.10.1
		3	7	198.206.184.177
4	0020da:999660	4	9	192.168.10.1 198.206.184.177

A visual illustration of these connections is shown here:



XMAP Network Example

The fields in `xmapls` table are defined as follows:

VPN. The local virtual port number which is connected to an adjacent switch.

Rem Switch ID. The MAC address of the MPX in the adjacent switch.

Rem VPN. The remote virtual port number in the adjacent switch.

Pri Group. The primary group associated with the remote port. The primary group is the group upon which Spanning Tree converges. For more information about primary groups, see Chapter 19, “Managing Groups and Ports.”

IP Addresses. All IP addresses associated with the adjacent switch.

Configuring the Discovery Transmission Time

The discovery transmission time is used in both the discovery transmission state *and* the common transmission state to determine how long the port will wait for Hello packets. For ports in the discovery transmission state, this timer is also used as the interval between sending out Hello packets.

◆ Note ◆

Ports in the common transmission state send out Hello packets based on the common transmission time as described in the next section.

Use the `xmapdisctime` command to view or update the discovery transmission time.

To view the current discovery transmission time, enter the following command:

```
xmapdisctime
```

A message similar to the following displays:

XMAP Discovery Phase Timeout Interval is 30 seconds.

To change the interval, enter the command with the desired value (any value between 1 and 65535). For example:

xmapdisctime 20

A message similar to the following displays:

XMAP Discovery Phase Timeout Interval is 20 seconds.

Configuring the Common Transmission Time

Use the **xmapcmntime** command to view or change the time between sending Hello update packets in the common transmission state. (This timer is only used in the common transmission state.) A switch sends an update for a port just before or after the common transmission time expires.

◆ Note ◆

The switches avoid synchronization by jittering the common transmission time by plus or minus ten percent of the configured value. For example, if the default common transmission time is used (300 seconds), the jitter is plus or minus 30 seconds.

When a Hello packet is received from an adjacent switch before the common transmission time expires, the switch sends a Hello reply and restarts the common transmission timer.

To view the current common transmission time, enter the following:

xmapcmntime

A message similar to the following displays:

XMAP Common Phase Timeout Interval is 300 seconds.

To change the interval, enter the command with the desired value (the value must be between 1 and 65535):

xmapcmntime 200

A message similar to the following displays:

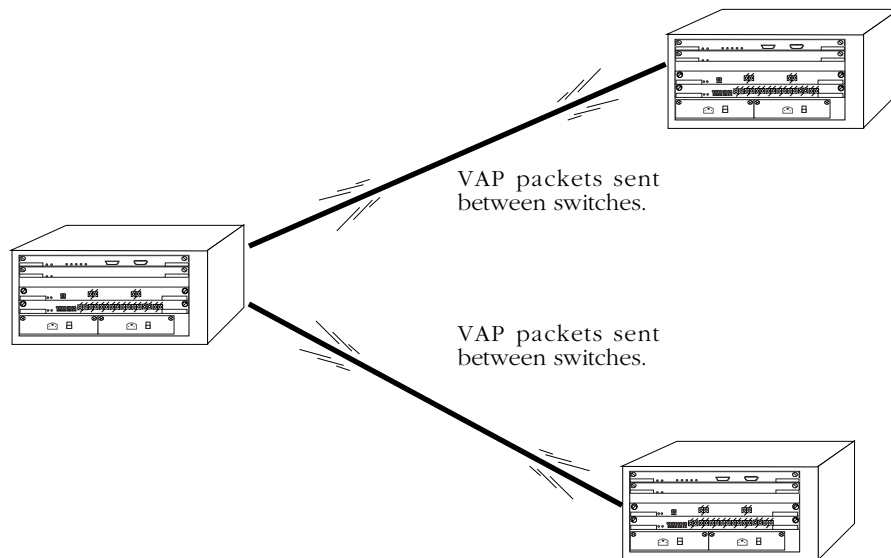
XMAP Common Phase Timeout Interval is 200 seconds.

VLAN Advertisement Protocol (VAP)

The VLAN Advertisement Protocol (VAP) is an interswitch protocol that keeps the VLAN membership databases stored on switches in sync and enables the auto-discovery of network nodes. VAP is useful when you want all VLANs to communicate over a backbone, but do not want locally connected devices to receive all backbone traffic.

In order for a switch to participate in VAP exchanges, VAP must be enabled through a software configuration command. The switch does not need to have attached devices that are a part of all groups and VLANs for which VAP information is exchanged; however, all groups and VLANs must be defined on each switch.

Each switch in a network maintains an AutoTracker database. This database is built by observing traffic that matches user-configured policies. The VAP protocol reads this database on all switches and then advertises entries in the database to all other switches in the network.



VAP Exchanges Between Switches

VAP updates nodes on any new entries in AutoTracker databases every 60 seconds.

VAP also stores information in its own database. Currently this information is used by SNMP-based network management software. The database contains information on VLAN membership; it maps each learned MAC address to a group and to any associated VLANs. This database can contain information on up to 40,000 MAC addresses.

VAP and Port Policies

One of the main purposes of VAP is to advertise the connectivity of devices attached to the switch via AutoTracker port policies. VAP eliminates the need to apply port policies to backbones to ensure that connectivity is established and maintained. When you use port policies, all devices heard through a port will become a member of the VLAN. Using port policies across backbones is not efficient because all devices learned over the backbone would be placed in the same VLAN since they would be attached to the same port.

For this reason, port policies should not be used to interconnect switches because these policies classify MAC addresses on VLANs. Backbone ports should be left in the default VLAN, and only learned devices should be segregated into VLANs by port policies.

There are two types of port policies (or rules), regular port rules and port forwarding rules. Only one can be active at a time. The type of port rule is determined by a command line in the `mpx.cmd` file. See *Port Policy Functionality* in Chapter 22, “Managing AutoTracker VLANs,” for a detailed explanation of the two port rules settings.

Regular port policy places frames received on a particular port into a VLAN; VLAN membership is based on the port. The current version of VAP supports regular port rules only.

If you set up VAP in its full mode (VLAN membership exchanges and auto-discovery), the switch will automatically set the port policy to *regular* mode.

◆ Note ◆

Earlier versions of VAP include an Adjacency Only mode. If an earlier version of VAP is running on the switch in Adjacency Only mode, when new code is loaded and the switch is rebooted, VAP will be set to off. If an earlier version of VAP is set to Full mode, VAP remains in Full mode when the new code is loaded and the switch is rebooted.

Configuring VAP

There are two settings for the VLAN advertisement protocol, off or full mode. These modes are defined as follows:

- *Full mode*—VLAN membership information exchanged between switches and auto-discovery of network nodes is enabled. This option automatically sets the port policy to regular mode.
- *Off*—Disables VAP exchanges. Nodes will not be auto-discovered and VLAN information will not be exchanged between switches.

To change the VAP mode, at a UI command prompt, enter `vlap` and select the mode in which you want VAP to run. A screen similar to the following displays:

```
The VLAN advertisement protocol is currently not running
To change the mode type: F - full mode, O - off : ( ) :
```

Or, enter the `vlap` command with the desired mode. For example:

```
vlap f
```

A message similar to the following displays:

```
The VLAN advertisement protocol is currently running.
```

The new mode takes effect immediately. You do not need to reboot the switch.

GMAP

The Group Mobility Advertisement Protocol (GMAP) enables workstation users to move from port to port among interconnected switches and still retain all learned mobile group and protocol information. Using GMAP the switch sends a complete list of learned MAC addresses and associated group/protocol information to all interconnected switches in the network. Update and retention times are configurable. A switch that receives a GMAP update packet updates its internal GMAP tables and queries the forwarding database to make any necessary updates.

At startup time and for three successive update intervals, GMAP sends update packets on all virtual ports that are active non-leaf ports (that is, ports that are running Spanning Tree). GMAP packets are sent using the VAP multicast address. After startup and three transmissions, interval packets will only be sent on virtual ports that are active and are known to have an OmniS/R running GMAP connected to them.

GMAP will send updates only for MAC addresses that are learned on leaf ports (ports that are not running Spanning Tree). It does not advertise MAC addresses for groups assigned by authentication, and it does not advertise group 1 entries or nonmobile group entries. If conflicting information is received for a MAC address, the last packet received for that address will take precedence.

When AutoTracker learns a new MAC address on a leaf port it attempts to assign it to a mobile group. It consults GMAP tables and any appropriate group membership entries are added to the forwarding database.

GMAP Updating Rules

Upon receiving a packet, GMAP updates its internal tables and queries the forwarding database. When GMAP reviews the forwarding database to update it with new information, it uses the following rules:

- GMAP will only update information for leaf ports.
- GMAP does not add a new MAC address to a port.
- GMAP will only overwrite group 1 entries. If there is no group 1 entry, it will add a new entry, provided that it will not create a conflict with existing entries in the forwarding database.
- GMAP will not add an entry for an authenticated group.
- GMAP will not add an entry that is in conflict or potential conflict with a binding rule. A potential conflict would be a binding rule that requires the IP address be known for the MAC address. GMAP does not have access to IP information.
- GMAP will not add an entry for a group/protocol pair when there is an existing entry for that protocol on the requested port.
- When GMAP finds an entry for the desired group already on the switch but not on the requested port, it will move it to the requested port.
- When GMAP finds an entry with the appropriate group but a protocol value of 0 (indicating all protocols), it will update the protocol value in that entry to that in its database.

Configuring GMAP

GMAP is inactive by default. In addition to enabling and disabling GMAP, you can configure the time between packet transmissions (when multiple packets are required for an update), the time between updates, and the length of time GMAP will retain its current information.

Enabling and Disabling GMAP

Use the **gmapst** command to display or change the state of GMAP. A prompt similar to the following displays:

```
GMAP is currently INACTIVE. (a)ctivate, (d)e-activate: (d)
```

Enter **a** or **d** or press **<Enter>** to keep the current value. A message similar to the following displays:

```
GMAP is ACTIVE.
```

To change the state of GMAP without displaying the current state first, enter the command with the desired value. For example:

```
gmapst d
```

The following message displays:

```
GMAP is INACTIVE.
```

Configuring the Gap Time

Use the **gmappgptime** command to display or change the interpacket gap time used when multiple packets are required for an update. When there are many MAC addresses on mobile ports, more than one GMAP packet is required for an update. Typically the gap time does not have to be changed, but you may want to modify it if traffic spikes are occurring in the network.

To view the current gap time, enter the following command:

```
gmappgptime
```

A message similar to the following displays:

```
GMAP Gap Time is 133 milliseconds.
```

To change the gap time, enter the command with the desired value (any value between 0 and 65535). For example:

```
gmappgptime 100
```

A message displays similar to the following:

```
GMAP Gap Time is 100 milliseconds.
```

The switch approximates the gap time because its internal clock does not use milliseconds. For any value shorter than one second, the switch uses 1/60 second increments called “ticks.” The default for gap time is 8 ticks or approximately 133 milliseconds. Any value you enter will be rounded to the nearest tick.

Configuring the Interpacket Update Time

Use the **gmapupdtime** command to display or change the time between sending updates.

◆ **Note** ◆

The switches avoid synchronization by jittering the update time by plus or minus one quarter of the configured interval. For example, if the default of 300 seconds is used, the jitter is plus or minus 75 seconds.

To view the current update time, enter the following:

```
gmapupdtime
```

A message similar to the following displays:

```
GMAP Update Time is 300 seconds.
```

To change the update time, enter the command and the desired time (any value between 1 and 65535). For example:

```
gmapupdtime 100
```

A message similar to the following displays:

```
GMAP Update Time is 100 seconds.
```

Configuring the Hold Time

Use the **gmapholdtime** command to display or change the length of time for which GMAP will retain information it has learned.

To view the current hold time, enter the following:

```
gmapholdtime
```

A message similar to the following displays:

```
GMAP Hold Time is 4320 minutes.
```

The default is 4320 minutes (72 hours). To change the current hold time, enter the command followed by the desired value (any value between 1 and 65535). For example:

```
gmapholdtime 2880
```

A message similar to the following displays:

```
GMAP Hold Time is 2880 minutes.
```

Displaying GMAP Statistics by MAC Address

To display GMAP statistics for all MAC addresses, use the **gmapls** command. The screen displays similar to the following:

GMAP Table						
MAC Address	Protocol	Group	Src Switch ID	Flags	Timeout(sec)	
000502:C07F11	1809B	12	0020DA:ECC770	00:00:00:00	3536	
	800	12	0020DA:ECC770	00:00:00:00	3536	
00105A:1873B9	1809B	12	0020DA:ECC770	00:00:00:00	3536	
	800	23	0020DA:ECC770	00:00:00:00	3536	

To limit the display, specify the MAC address. For example:

```
gmapls 00105A:C07F11
```

Fields in this table are defined as follows:

MAC Address. The MAC address of the local end station.

Group. The group(s) to which the MAC address belongs.

Protocol. The protocol associated with the group on the switch from which the information was received. Protocol values are defined as follows:

- e0e0 or ffff — IPX over 802.3
- 8137 — IPX over Ethernet II
- 18137 — IPX over SNAP
- 28137 — any IPX encapsulation
- 800 — IP
- 809b — AppleTalk
- 1809b — AppleTalk over SNAP
- 6003 — DECNET

Src Switch ID. The MAC address of the switch from which the entry was received.

Flags. The first two bytes are not used. The third byte displays the AutoTracker flags associated with the entry on the source switch. The last byte displays the router flags associated with this entry on the source switch.

Timeout (sec). The number of seconds remaining until this entry is deleted (unless another GMAP message is received and then the entry is refreshed).

22 Managing AutoTracker VLANs

In a large, flat, switched network, broadcast traffic can overload a network based primarily on port-based Groups. Through the use of AutoTracker VLANs, you can control broadcast traffic such that it is forwarded only to those VLANs where it needs to be sent.

VLANs are created within a Group to subdivide network traffic based on specific criteria. The criteria you use to define a VLAN are called AutoTracker policies. AutoTracker policies can be defined by port, MAC address, protocol, network address, a user-defined policy, or a multicast policy. You can also define multiple policies—also referred to as “rules”—for a VLAN if you wish. A port or device is included in a VLAN if it matches any one VLAN rule. For example, you can define rules based on MAC address and rules based on protocol in the same VLAN.

A Group defines a physical space within the network—a set of ports. The policies that you define for VLAN membership are applied to all traffic on those ports, but not to traffic on ports outside the Group.

You can create two types of policy-based VLANs: AutoTracker VLANs and multicast VLANs. You can create up to 31 AutoTracker VLANs and up to 32 multicast VLANs in any one Group. AutoTracker VLANs and multicast VLANs operate independently of one another: the policies you establish for AutoTracker VLANs neither conflict nor interfere with the policies you establish for multicast VLANs, even when those policies involve the same ports or MAC addresses.

This chapter provides an overview of AutoTracker VLANs and multicast VLANs as well as instructions for managing and monitoring each type of VLAN. Instructions for configuring AutoTracker policies can be found in Chapter 20, “Configuring Group and VLAN Policies.”

The AutoTracker Menu

All software commands for configuring AutoTracker policies and AutoTracker/multicast VLANs are in the AutoTracker menu. This menu is a submenu of the VLAN menu. You can access the AutoTracker menu by typing **at** any prompt. The menu displays as follows:

Command	Auto-Tracker Management Menu
cratvl	Create an Auto-Tracker VLAN
atvl	View definition of Auto-Tracker VLAN
viatrl	View Auto-Tracker Rule Configuration
rmatvl	Delete an Auto-Tracker VLAN
modatvl	Modify definition of an Auto-Tracker VLAN
vivl	View list of Active Auto-Tracker VLANs on an interface
fwtvvl	View VLAN assignment of learned MAC addresses
defvl	Enable or disable membership in default VLAN
crmcvl	Create a Multicast VLAN
mcvl	View definition of Multicast VLAN
vimcrl	View Multicast VLAN Rule Configuration
rmmcvl	Delete a Multicast VLAN
modmcvl	Modify definition of a Multicast VLAN
vimcvl	View list of Active Multicast VLANs on an interface
gmstat	Turn Group Mobility Status ON or OFF
vpl	View Virtual Ports in a Mobile Group
vigl	View Mobile Group List for a Virtual Port
cats	Create Auto-Activated Services
data	Delete Auto-Activated Services
vats	View Auto-Activated Services
vag	View Authenticated Groups
gmcfg	Configure Group Mobility Parameters
mag	Modify Authenticated Group
xip	Enter the Xylan Inter-switch Protocol (XIP) sub-menu

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The commands on the AutoTracker menu can be roughly divided into two halves. The first half of commands—listed from **cratvl** to **vimcvl**—apply mainly to AutoTracker VLANs (i.e., VLANs created inside non-mobile groups). An exception to this rule is the **modatvl** command, which can be used to modify AutoTracker policies for VLANs or mobile groups. In addition many of the informational commands apply to both VLANs and mobile groups. The commands that apply to AutoTracker VLANs are described in this chapter. Multicast VLANs are described in Chapter 23, “Multicast VLANs.” The mag command is described in the *Switched Network Services User Manual*. The XIP sub-menu is described in Chapter 21, “Inter-switch Protocols.”

The commands from **gmstat** to **gmcfg** apply strictly to mobile groups. All of the commands in this second set are described in 19 24, “Managing Groups and Ports.”

AutoTracker VLANs

AutoTracker VLANs enable you to control communications between end stations in your network. You define policies that determine membership in the VLAN and AutoTracker automatically locates ports or devices within the Group that fit the policies and places them into the VLAN.

You can define physical policies or logical policies (or combinations thereof) to determine membership in AutoTracker VLANs. Physical policies consist of port rules: you define the VLAN members as one or more specific ports and VLAN membership is limited to the ports defined and the MAC addresses of devices connected to those ports.

Logical VLAN policies can consist of MAC address rules, protocol rules, network address rules, or user-defined rules. Ports are assigned to VLANs that have logical rules when the MPX examines frames that originate from devices connected to the Group's set of ports. If a frame is received that matches a logical VLAN rule, the source device's MAC address and the port to which the source device is connected are both made VLAN members.

The members of an AutoTracker VLAN thus consist of source devices originating frames that fit the VLAN's policies and the ports to which those source devices are connected. Instructions for creating AutoTracker VLANs begin on page 22-16.

AutoTracker VLAN Policies

You can define a maximum of 32 AutoTracker policies of each type per Group. There is no restriction on the number of rules you can define per VLAN, as long as the maximum number of policies for the Group is not exceeded.

A switch port – or a device connected to a switch port – can belong to more than one VLAN simultaneously, as determined by the rules the port or device matches. A port or device is included in a VLAN if it matches any one rule.

You can define the following types of rules:

Port Policies. Port policies enable you to define membership in the VLAN on the basis of ports. Members of the VLAN will consist of devices connected to specific ports on one switch or on multiple switches in the Group.

MAC Address Policies. MAC address policies enable you to define membership in the VLAN on the basis of devices' MAC addresses. This is the simplest type of rule and provides the maximum degree of control and security. Members of the VLAN will consist of devices with specific MAC addresses. These devices may all be connected to one switch or they may be connected to different switches in the Group. A maximum of 10,240 MAC addresses are supported per policy.

Protocol Policies. Protocol policies enable you to define membership in the VLAN on the basis of the protocol that devices use to communicate. All devices that communicate with the specified protocol become members of the VLAN.

You can specify VLAN membership according to the following protocols: IP, IPX, AppleTalk, or DECNet. In addition, you can specify membership according to Ethernet type, source and destination SAP (service access protocol) header values, or SNAP (sub-network access protocol) type.

Network Address Policies. Network address policies enable you to define membership in the VLAN on the basis of network address criteria.

For example, you can specify that all IP users with a specific subnet mask be included in the VLAN. Or, you can specify that all IPX users in a specific network address area using a certain encapsulation type be included in the VLAN.

If you define network address and port or protocol rules in the same VLAN, the network address rules will take precedence over the port and protocol rules should any conflict arise. To reverse this precedence (i.e., port and protocol rules take precedence over network address rules) you must add the following line to the switch's **mpx.cmd** file:

Precedence=0

User-Defined Policies. User-defined policies enable you to define membership in the VLAN on the basis of a specific pattern within a frame. All devices that originate frames containing this pattern are assigned to the VLAN. The pattern is specified by defining an offset, a value, and a mask.

Port Binding Policies. A port binding policy specifies a particular device to be included in the mobile group or AutoTracker VLAN. You can bind a device's IP address to a switch port and a MAC address, or bind a device's MAC address to a protocol and a switch port.

DHCP Port Policies. These policies are similar to standard port policies, but apply to switch ports to which DHCP client workstations are attached.

DHCP MAC Address Policies. These policies are similar to standard MAC address policies, but apply to the MAC addresses of DHCP client workstations only.

The Default VLAN

The default AutoTracker VLAN, also referred to as VLAN #1, is different from other AutoTracker VLANs. The following list outlines some of these differences.

1. The default VLAN is automatically created when you create a new Group. Non-default VLANs must be created through the **cratvl** command.
2. The default VLAN cannot be removed. Other VLANs can be removed through the **rmatvl** command.
3. You cannot apply AutoTracker policies to the default VLAN. Other non-default AutoTracker VLANs allow you to apply any policy to them.

You can enable routing on the default VLAN. You enable the default VLAN virtual router through the **crgp** or **modvl** command. See Chapter 19, "Managing Groups and Ports," for further information on the virtual router port on the default VLAN.

All ports and devices in a Group initially belong to default VLAN #1. All physical switch ports always remain members of the default VLAN, but they can also become members of other VLANs. It is not possible to delete a physical switch port from VLAN #1. Individual network devices, however, can move out of VLAN #1. All MAC devices are also initially part of default VLAN #1. However, when a MAC device is removed from default VLAN #1 and moved into a non-default VLAN, it is deleted from default VLAN #1.

The default VLAN is explained further in other sections of this chapter. See *How Devices are Assigned to AutoTracker VLANs* on page 22-5 for a discussion of default VLAN membership issues and the **defvl** command. Also, see *Application Example 4* in Chapter 24, "AutoTracker VLAN Application Examples," for discussions of routing issues and the default VLAN.

How Devices are Assigned to AutoTracker VLANs

When a broadcast frame, a multicast frame, or a unicast frame from an unknown device is received at a switching module, the frame is forwarded to the MPX for processing. Source learning logic on the MPX module examines the entire frame to determine the VLAN or VLANs in which the originating device should be a member. If the frame matches any one policy defined for a VLAN, the originating device (and the port to which it is connected) are made members of that VLAN. If the frame does *not* match any VLAN policy, one of the following occurs:

- If the **defvl** command is on, the source device is made a member of Default VLAN #1 in the Group of which the source port is a member. The **defvl** command determines whether traffic from devices that do not match any policies is assigned to the default VLAN or dropped. (See “The defvl Command” below for more information on this command.)
- If the **defvl** command is off, all traffic from the source device is dropped.

Please Take Note

A broadcast or multicast frame is processed to determine the source device’s VLAN membership each time it is received. A unicast frame is processed to determine the source device’s VLAN membership only the first time it is received.

When the MPX module has determined the VLAN or VLANs in which the originating device belongs, it relays this information to the switching module. The switching module updates a VLAN membership flag attached to the frame’s source MAC address in the CAM (content-addressable memory). The frame is then switched based on this membership flag.

Refer to Chapter 24, “AutoTracker VLAN Application Examples,” for information on AutoTracker VLAN assignments in specific network situations.

The defvl Command

You can turn the **defvl** command on and off simply by entering **defvl on** or **defvl off**. If you enter the command without any parameters, it displays the current setting for the Default VLAN. For example, if source devices are automatically placed in the Default VLAN when they do not match any VLAN policy rule, the following message would display:

membership in default vlan is currently on

If source devices are automatically dropped when they do not match any VLAN policy, the following message would display:

membership in default vlan is currently off

The **defvl** command applies to all Groups in an Omni Switch/Router and it is only applicable if there is at least one AutoTracker rule configured.

Devices that Generate a Secondary Traffic Type

Source devices sometimes generate more than one traffic type; for example, a device could generate IP traffic primarily but also generate a secondary stream of AppleTalk. When a device generates secondary traffic that does not match any existing VLAN policy, that traffic is grouped into the primary VLAN of which the device is a member.

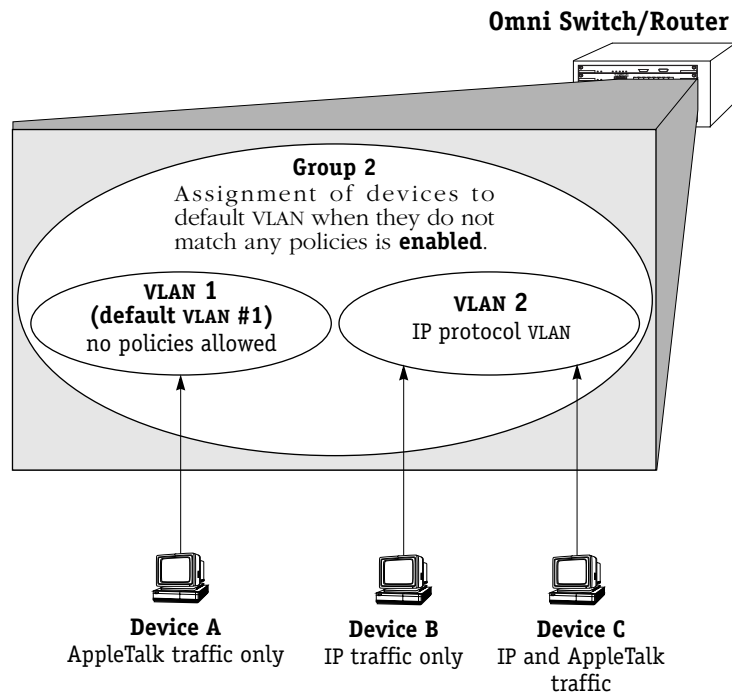
To continue the example, if a device generates both IP and AppleTalk, and both an IP VLAN and an AppleTalk VLAN exist, that device is made a member of both VLANs and no problem occurs. If, however, an AppleTalk VLAN does not exist, all traffic from that device is grouped into the existing VLAN of which the device is a member – in this example, the IP VLAN. This can cause communication problems, as explained below. **For this reason, it is advisable to create VLANs that accommodate all known network traffic.**

In this example Device A is assigned to default VLAN #1 because it does not match any existing VLAN policy.

Devices B and C are assigned to VLAN 2 because they generate IP traffic. The secondary AppleTalk traffic Device C generates is also grouped into VLAN 2, since the AppleTalk traffic does not match any existing VLAN policy.

The result is that Devices A and C are unable to communicate.

Creation of an AppleTalk protocol VLAN solves this problem. If an AppleTalk VLAN exists, Device A will be assigned to it and removed from Default VLAN #1. Device C will be assigned to both the IP VLAN and the AppleTalk VLAN. Devices A and C can then communicate.



How Devices are Assigned to AutoTracker VLANs (*continued*)

Router Traffic in IP and IPX Network Address VLANs

Prior to release 2.1, AutoTracker handled VLAN assignments for router traffic in IP and IPX network address VLANs in the same manner as normal traffic. In release 2.1 and later, AutoTracker differentiates router traffic from normal traffic and can distinguish traffic that is routed *through* a router from traffic that is generated *by* a router.

AutoTracker now determines VLAN assignments for router interfaces (that is, the MAC addresses of router interface ports) in IP and IPX network address VLANs based on router update messages generated by the router itself. This minimizes VLAN leakage and avoids the problem situation described on the facing page.

The Problem with Router Traffic

AutoTracker functions on the assumption that data in a frame can be associated with the frame's source MAC address. For example, if a frame has an IPX network number of 300, AutoTracker assumes that it has received the frame directly and that the source device is a member of IPX network 300. This is not true in the case of routed frames. Routers route frames from one network to another by changing the frame's MAC header but keeping the layer 3 content intact. This can lead to the problem situation described on the facing page.

In the network on the facing page, Device A gets correctly assigned to VLAN 2 and Device B gets correctly assigned to VLAN 3 without problem. The two router interfaces will be assigned to the correct VLANs *if AutoTracker learns the router interface MAC addresses from their RIP updates*. However, this may not happen. The problem situation on the facing page shows what can occur if AutoTracker learns the router interface MAC addresses from traffic routed through the router rather than from traffic generated by the router (such as a RIP update).

How AutoTracker Handles Router Traffic

To avoid the problem situation on the facing page, AutoTracker now determines if any IP or IPX device it has learned is a router. If it is, AutoTracker marks the device as a router, unlearns all previous VLAN assignments for that device, and reassigns the device based on a router-generated update packet (such as a RIP packet).

AutoTracker determines if a learned device is a router by searching further within the frame. For example, if AutoTracker receives an IP frame, it searches beyond the source IP address and also checks if the IP frame is a RIP, OSPF, BGP, DVMRP, or IGRP update. If it is, as explained, AutoTracker marks the device as a router, unlearns its previous VLAN assignments, and reassigns it using the router-generated update packet.

AutoTracker recognizes the following types of router-generated frames:

- IP protocol: RIP frames, OSPF frames, BGP4 frames, DVRMP frames, and IGRP frames
- IPX protocol: IPX RIP frames and SAP frames

AutoTracker maintains a record of the devices it has learned are routers. Each time a router-generated frame is received from a device marked as a router, AutoTracker updates that device's membership in IP or IPX network address VLANs. If a frame received from a device marked as a router is not IP or IPX, VLAN membership is updated normally.

Please Take Note

This special handling of router traffic occurs in IP and IPX network address VLANs only. Note that it does not alter normal VLAN assignment processes such as checking for VLAN policy matches other than IP or IPX network address.

How Routed Frames can Confuse VLAN Assignment

- 2 The router receives the frame on the interface for Network 2 and routes the frame to the interface for Network 3. To do this, the router strips the MAC header from the frame and inserts the MAC address of its interface for Network 3. The frame now specifies its source as Network 2, MAC address Y. **Frame** from Network 2, MAC address Y

- 1 Device A initiates a request to route a frame to Device B. The switch forwards the frame to the router interface for Network 2.

Frame
from Network 2, MAC address A


Please Note

This example assumes that AutoTracker has not yet learned the MAC addresses of the router interfaces for Networks 2 and 3.

- 3 The switch receives the frame from Network 2, MAC address Y, forwards it to its destination, Device B, and gives it to AutoTracker to be learned.

Per the policies, AutoTracker assigns the frame's source MAC, MAC address Y, to VLAN 2. The router interface for network 3 is now a member of network 2!

Frame
from Network 2, MAC address Y



- 4 Let's say that the next transmission is a RIP update from the router interface for network 3. The source of the RIP update is Network 3, MAC address Y. AutoTracker thus assigns MAC address Y to VLAN 3. MAC address Y is now assigned to both VLAN 2 and VLAN 3.

The same situation can occur with MAC address X on the router interface for network 2. Both router interfaces will be members of both VLANs and will transmit RIP updates to both.

If this is an IPX network and IPX servers are members of these VLANs, they will respond with router configuration errors. If this is an IP network and devices A and B are IP workstations listening to RIP, they will respond with invalid network address errors.

How Devices are Assigned to AutoTracker VLANs (*continued*)

Port Policy Functionality

In release 2.1 and later, AutoTracker's VLAN port policy can be set to operate in either of two distinct modes:

- In the original mode, wherein membership in all VLANs active on a port is inherited by all devices connected to that port. Original port policy functionality is explained on page 22-10.
- In a new mode, wherein membership in all VLANs active on a port **is not** inherited by all devices connected to that port. This is the current, default functionality with which the switch ships. Current port policy functionality is explained on page 22-11.

Port policy functionality is set on a switch-wide basis, via a flag in the switch's `mpx.cmd` file called `reg_port_rule`. The switch ships with port policy functionality set to operate in the new mode. You can revert the switch to original port policy functionality by editing the file and setting the `reg_port_rule` flag to 1. You must then restart the switch. (The file is accessed, and can be edited, via the switch User Interface. You can view the current setting of `reg_port_rule` with the `view mpx.cmd` command. See Chapter 7, "Managing Files," for information on editing the `mpx.cmd` file.)

Why the New Functionality?

Port policies can cause problems in a multi-switch environment. AutoTracker assumes that each switch in a multi-switch environment can independently arrive at identical VLAN assignments for all devices in the network. This is not true when port policies are in effect because of their very nature: port policies are switch-specific and not network wide. The figure on page 22-10, which explains original port policy functionality, provides an example of how port policies can result in inconsistent VLAN membership between two switches – notice the inconsistent VLAN membership in Omni Switch/Router 1 and in Omni Switch/Router 2.

The use of port policies in a multi-switch environment can result in connectivity problems if the source switch and the destination switch are separated by other switches. The switches along the path of the frame will not have identical VLAN memberships. At any particular switch along the path, frames could be lost because of inconsistencies in the VLAN membership of the frames' source and destination devices.

In addition, AutoTracker maintains devices in the same VLAN without regard to the devices' location – provided the devices match the same AutoTracker policies throughout the network. Multiple switches will assign a device to the same VLANs provided that device matches the same policies on each switch. This is not possible when port policies are in effect because, as stated, by their very nature port policies are switch-specific and not network-wide.

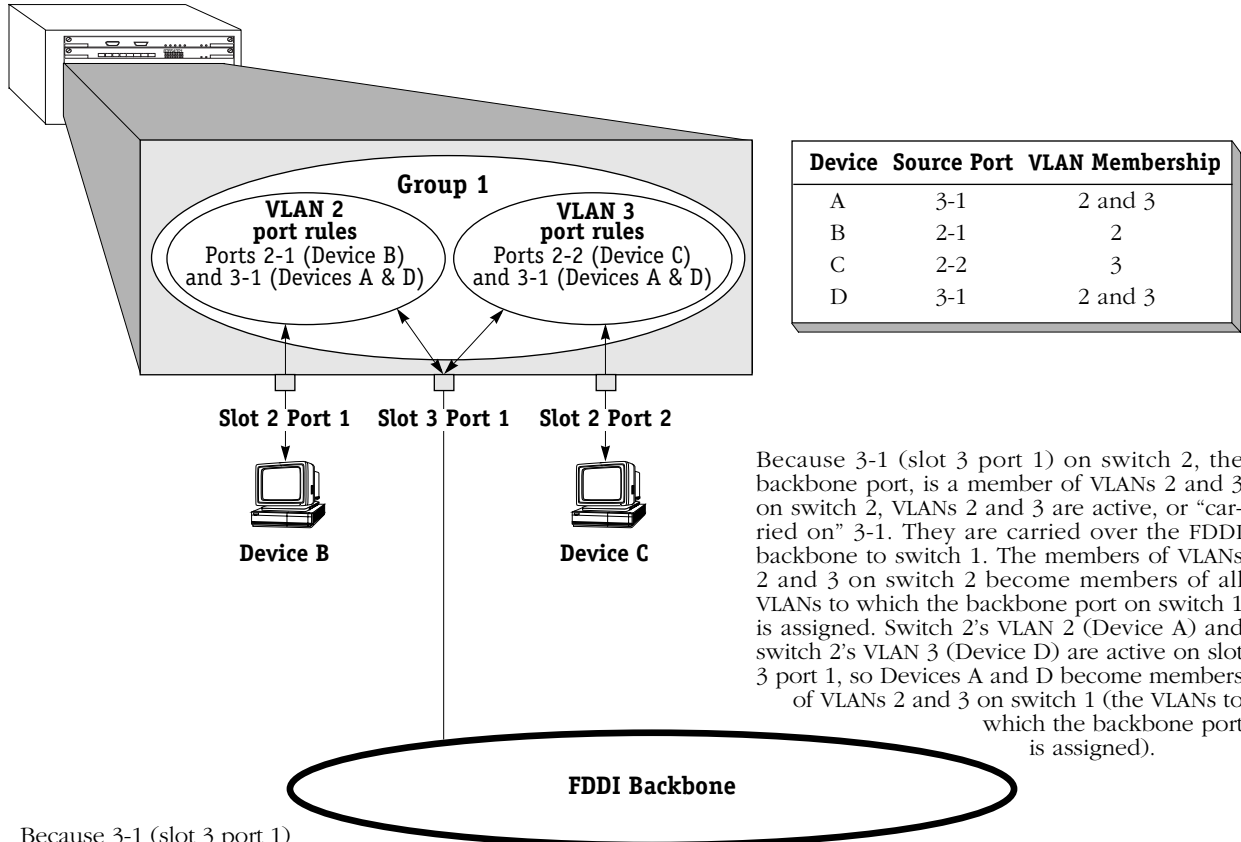
For these reasons, the Omni Switch/Router now ships with new port policy functionality (although, as explained, you can revert the switch to original port policy functionality if you wish). The new functionality still enables users to assign ports to VLANs and still enables those ports to carry traffic for those VLANs. However, with the new functionality, port policies are not used to learn VLAN assignments for traffic received on ports (as explained on page 22-11). In order for a device to be assigned to a VLAN, it must match an existing logical policy of the VLAN. This is explained on page 22-13.

The Following Examples

The following pages provide examples of original and current port policy functionality. The limitations of port policies become apparent if one tries to use port policies to create two VLANs in these sample networks, one for Devices A and B and one for Devices C and D.

Original Port Policy Functionality
(reg_port_rule = 1)

Omni Switch/Router 1

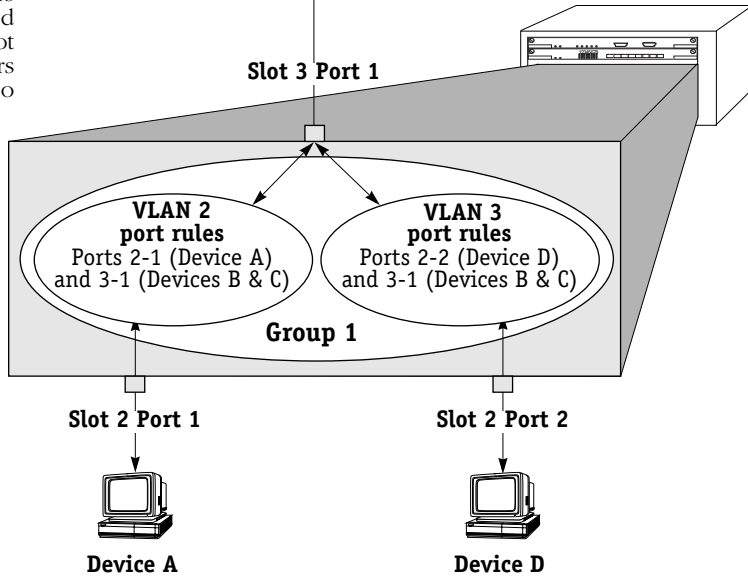


Because 3-1 (slot 3 port 1) on switch 2, the backbone port, is a member of VLANs 2 and 3 on switch 2, VLANs 2 and 3 are active, or “carried on” 3-1. They are carried over the FDDI backbone to switch 1. The members of VLANs 2 and 3 on switch 2 become members of all VLANs to which the backbone port on switch 1 is assigned. Switch 2’s VLAN 2 (Device A) and switch 2’s VLAN 3 (Device D) are active on slot 3 port 1, so Devices A and D become members of VLANs 2 and 3 on switch 1 (the VLANs to which the backbone port is assigned).

Because 3-1 (slot 3 port 1) on switch 1, the backbone port, is a member of VLANs 2 and 3 on switch 1, VLANs 2 and 3 are active, or “carried on” 3-1. They are carried over the FDDI backbone to switch 2. The members of VLANs 2 and 3 on switch 1 become members of all VLANs to which the backbone port on switch 2 is assigned. Switch 1’s VLAN 2 (Device B) and switch 1’s VLAN 3 (Device C) are active on slot 3 port 1, so Devices B and C become members of VLANs 2 and 3 on switch 2 (the VLANs to which the backbone port is assigned).

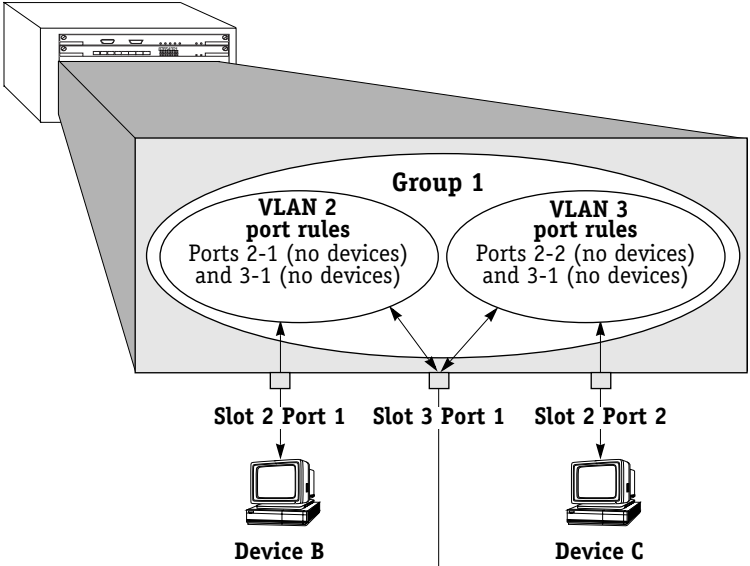
Device	Source Port	VLAN Membership
A	2-1	2
B	3-1	2 and 3
C	3-1	2 and 3
D	2-2	3

Omni Switch/Router 2



Current Port Policy Functionality
(reg_port_rule = 0)

Omni Switch/Router 1



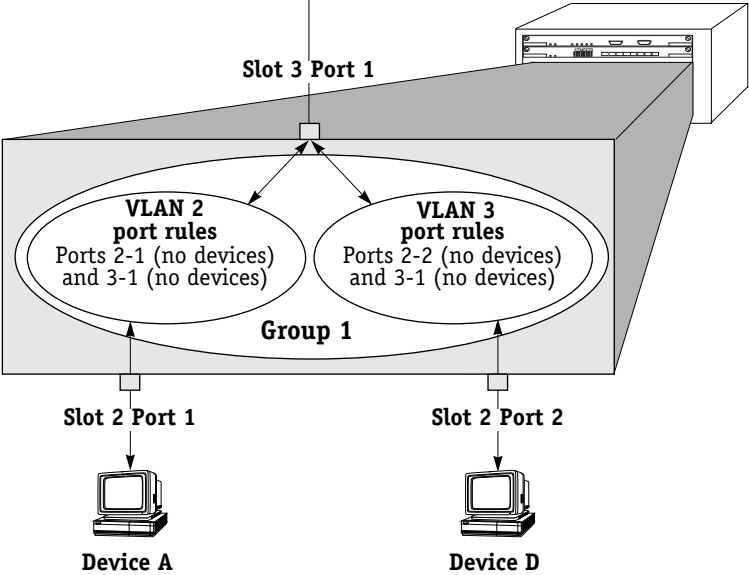
Device	Source Port	VLAN Membership
A	3-1	default VLAN #1
B	2-1	default VLAN #1
C	2-2	default VLAN #1
D	3-1	default VLAN #1

With current port policy functionality, VLANs are not active, or “carried on” ports. Port 3-1 (slot 3 port 1) on switch 2, the backbone port, is still a member of VLANs 2 and 3 on switch 2, but VLANs 2 and 3 **are not** carried over the FDDI backbone to switch 1. The members of VLANs 2 and 3 on switch 2 **do not** become members of all VLANs to which the backbone port on switch 1 is assigned. Rather, they become members of default VLAN #1 on switch 1 if they do not match any of switch 1’s existing VLAN policies and if **defvl** is on for the switch. If **defvl** is off, the traffic is dropped.

With current port policy functionality, VLANs are not active, or “carried on” ports. Port 3-1 (slot 3 port 1) on switch 1, the backbone port, is still a member of VLANs 2 and 3 on switch 1, but VLANs 2 and 3 **are not** carried over the FDDI backbone to switch 2. The members of VLANs 2 and 3 on switch 1 **do not** become members of all VLANs to which the backbone port on switch 2 is assigned. Rather, they become members of default VLAN #1 on switch 2 if they do not match any of switch 2’s existing VLAN policies and if **defvl** is on for the switch. If **defvl** is off, the traffic is dropped.



Omni Switch/Router 2



Device	Source Port	VLAN Membership
A	2-1	default VLAN #1
B	3-1	default VLAN #1
C	3-1	default VLAN #1
D	2-2	default VLAN #1

The Usefulness of Port Policies

As has been explained – and as illustrated on page 22-10 – original port policy functionality is not well-suited to the creation of consistent VLAN membership in a multi-switch environment. Current port policy functionality – as illustrated on page 22-11 – neither contributes to nor participates in VLAN assignments. Port policies, either original or current, are in fact not useful in the creation of consistent VLAN membership across multiple switches. Logical policies are of far greater use, as illustrated on page 22-13. So, why use port policies at all?

Port Policies are Useful in these Situations:

- **Silent stations.** If a device does not transmit traffic (such as a printer), the port to which the device is connected never gets assigned to VLANs. It is then impossible for other stations to communicate with that device. Creating a port policy that assigns the silent device's port to one or more VLANs will enable traffic to flow out that port to the silent device.
- **Inactive VLANs.** AutoTracker does not activate a VLAN – or its internal router – until a port is assigned to that VLAN. AutoTracker assigns ports to VLANs with port policies immediately. However, AutoTracker only assigns ports to VLANs with logical policies when a frame is received from a source device that matches the VLAN's policies. This means that, in some network situations, you may need to assign a port policy to a VLAN to force it active. *Application Example 5* in Chapter 24 provides an example of this.
- **Backbone connections.** A port policy that assigns the backbone port to a VLAN will enable traffic from that VLAN to flow out onto the backbone.

◆ Important Note ◆

If you are using port policies to extend VLANs across a backbone, you are strongly advised to use current (default) port policy functionality. If you use original port policy functionality, you are, in effect, placing all devices learned from the backbone port into the same VLAN. If the port policy is configured for all VLANs (so that all VLANs can communicate over the backbone), all devices learned from the backbone port are assigned to all VLANs. This is not desirable – it would subject locally-connected devices to all the backbone traffic.

So How Do I Get Devices Assigned to VLANs Over a Backbone?

The way to get devices assigned to VLANs over a backbone is to define logical VLAN policies that so assign them. An example is shown on the facing page utilizing IP and IPX protocol policies. The network on the facing page uses port policies (and current port policy functionality) to assign the backbone port to VLANs on each switch so that traffic can flow out onto the backbone from these VLANs.

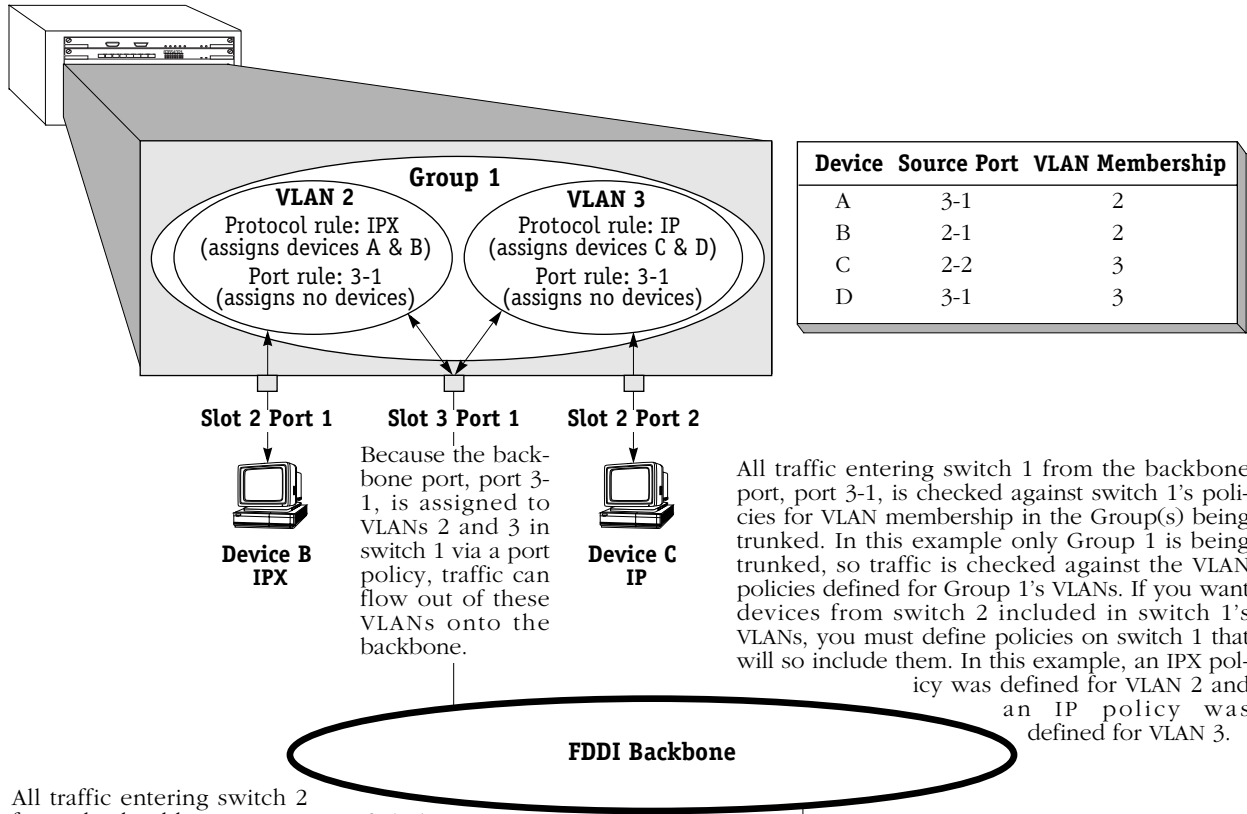
The problem of remote VLAN assignments is solved by the IP and IPX protocol policies. When a frame is received from a backbone port, the frame is examined to determine if it matches any VLAN membership rules. Let's say Device D on switch 2 transmits an IP frame. The frame travels the FDDI backbone and enters switch 1 on port 3-1. AutoTracker learns the frame and assigns it to VLAN 3, since VLAN 3 has an IP protocol policy and the frame is IP.

Notice that with this approach:

- VLAN membership is consistent between the two switches.
- In a multi-switch environment, no frames are lost in switches along the traffic path because of the inconsistent VLAN membership of a frame's source and destination devices.
- Devices can be moved from switch to switch and they will be assigned to the same VLAN – without reconfiguring AutoTracker or the device.
- As was the original intent, it is possible to create two VLANs in this sample network, one for Devices A and B and one for Devices C and D. As is apparent, this was impossible using port policies.

An Example of VLAN Assignment Using Logical Policies and Current Port Policy Functionality (reg_port_rule = 0)


Omni Switch/Router 1



All traffic entering switch 2 from the backbone port, port 3-1, is checked against switch 2's policies for VLAN membership in the Group(s) being trunked. In this example only Group 1 is being trunked, so traffic is checked against the VLAN policies defined for Group 1's VLANs. If you want devices from switch 1 included in switch 2's VLANs, you must define policies on switch 2 that will so include them. In this example, an IPX policy was defined for VLAN 2 and an IP policy was defined for VLAN 3.

Because the backbone port, port 3-1, is assigned to VLANs 2 and 3 in switch 2 via a port policy, traffic can flow out of these VLANs onto the backbone.

Omni Switch/Router 2



Frame Flooding in AutoTracker VLANs

Flooding occurs when a frame is received addressed to a device that is unknown to the switch or broadcast or multicast frames are received addressed to multiple users. In a typical bridged environment, the frame would be forwarded out all ports. However, this is not true with VLANs as VLANs segment the network into smaller broadcast domains. In this environment, flooding occurs as follows:

Unicast Traffic

- If the destination address of the frame is unknown but its source address is known and the source device is a member of one or more VLANs, the frame is flooded out all ports of all VLANs in which the source device is a member. Please note the following:
 - If the source device is a member of multiple VLANs, some leakage may occur during the flooding process. Leakage may occur only among VLANs in the same Group—frames do not leak between Groups.
 - If the source device is a member of multiple VLANs and some or all of those VLANs share the same physical port, only one copy of the frame is forwarded out that port.
 - If the source device is a member of multiple VLANs that use trunking, only one copy of the frame is sent to each trunk port.
- If both the source and destination addresses of the frame are unknown, the frame is forwarded to the MPX for processing (to determine the VLAN or VLANs in which the originating device should be a member) **and** the frame is flooded out all ports of all VLANs in which the source port is a member.

Broadcast and Multicast Traffic

Frames are forwarded out all ports that are members of the same VLANs as the source MAC address. If the source MAC address is unknown, it is forwarded out all ports that have VLANs active on the source ports.

Routing Between AutoTracker VLANs

Devices that do not share membership in a common VLAN must use routers to communicate with one another. You can configure a virtual router port that is capable of IP and/or IPX routing for each VLAN. By enabling a router port on a VLAN, you are creating a static route entry within the switch to that VLAN. If this router port is not configured for a VLAN, then that VLAN will not be able to communicate with other VLANs unless an external router is between those VLANs. You may configure up to 16 virtual router ports within a single Omni Switch/Router. Each VLAN may contain only one router port.

Routing and the Default VLAN. You can enable routing for the default VLAN when you initially create a Group, or when you modify the Group. There are several issues about which you should be aware when enabling routing on the Default VLAN. See *Application Example 4* in Chapter 24, “AutoTracker VLAN Application Examples,” for more information.

Creating AutoTracker VLANs

You create AutoTracker VLANs through the AutoTracker menu options. Creating an AutoTracker VLAN includes the following steps:

- A.** Enter basic information such as the name and number for the VLAN. See *Step A. Entering Basic VLAN Information* on page 22-16 for instructions on this step.
- B.** Define policies that define membership in the VLAN. See *Step B. Defining and Configuring VLAN Policies* on page 22-18 for instructions on this step.
- C.** Configure the type of routing used for communication between VLANs. In order for devices in a VLAN to communicate with devices in other VLANs, a virtual router must be configured or an external router must exist between those VLANs. See *Step C. Configuring the Virtual Router Port (Optional)* on page 22-19 for instructions on this step.

These steps are explained in detail in the sections that follow.

Step A. Entering Basic VLAN Information

1. To begin setting up the AutoTracker VLAN type **cratvl** at any prompt.
2. The following prompt displays:

Enter the VLAN Group id for this VLAN (1):

Enter the number for the Group to which this VLAN will belong. All VLANs belong to a Group. You can create up to 31 VLANs per Group (each Group already contains a default VLAN, VLAN #1).

3. The following prompt displays:

Enter the VLAN Id for this VLAN (2):

Enter the number that will identify this VLAN with the Group specified above. Up to 32 VLANs may belong to the same Group (including the default VLAN). By default the system displays the next available VLAN ID number. Press **<Enter>** to accept this default.

4. The following prompt displays:

Enter the new VLAN's description:

Enter a textual description that will help you identify the VLAN. For example, if you know this VLAN will be composed of only workstations using the IPX protocol, you might call the VLAN, "IPX VLAN." You may use up to 30 characters for this description.

5. The following prompt displays:

Enter the Admin Status for this vlan (Enable (e) / Disable (d):

Enter whether or not you want the Administrative Status for this VLAN to be enabled or disabled. Once enabled, the switch begins using the policies you defined. A disabled VLAN is still defined (name, number, policies intact), but the switch keeps the VLAN disabled. The enable/disable status may be changed at a later time using the **modatvl** command.

Note

A VLAN may not always be operational even when its **Admin** Status is enabled. The VLAN becomes operational as soon as a port is assigned to it. In addition, a VLAN's operation may be disabled by the switch because devices in the VLAN cease transmitting data, among other reasons.

After you enter the Administrative Status, additional prompts display that allow you to select the rules governing membership in this VLAN. Go on to the next section, *Step B. Defining and Configuring VLAN Policies* on page 22-18 to continue setting up this VLAN.

Step B. Defining and Configuring VLAN Policies

You can define AutoTracker policies by port, MAC address, protocol, network address, user definition, or port binding. You can define multiple policies for a AutoTracker VLAN if you wish. A port or device is included in a AutoTracker VLAN if it matches any one rule. For example, you can define rules based on ports, rules based on MAC address, and rules based on protocol in the same AutoTracker VLAN. However, defining multiple rules is not trivial – exercise extreme care when you do so and make sure that you understand the consequences of your definitions. In most situations, it is advisable to use one of AutoTracker’s predefined rules.

Instructions for defining each AutoTracker policy type are included in Chapter 20, “Configuring Group and VLAN Policies.” Follow the directions in that chapter for the policy you wish to set up.

The sections below provide directions for setting up each type of AutoTracker policy. Follow the directions for the policy you wish to set up.

1. When are done specifying AutoTracker policies the following prompt displays:

Configure more rules for this vlan (y/n):

You can set up multiple rules for the same VLAN. Enter a **Y** here if you want to set up more rules in addition to the Network Address rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this VLAN. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

VLAN 1:2 created successfully

You are done setting up rules for this VLAN, so you can start configuring the virtual router for this VLAN. See *Step C. Configuring the Virtual Router Port (Optional)* on page 22-19 for information on configuring a virtual router port.

Step C. Configuring the Virtual Router Port (Optional)

You can now optionally configure the virtual router port that this VLAN will use to communicate with other AutoTracker VLANs. A virtual router port for the VLAN is created within the switch. If you do not define a virtual router port for this VLAN, devices within the VLAN will only be able to communicate with devices in other VLANs through an external router.

You will have the choice of configuring IP, IPX, or both IP and IPX routing. Continue with the steps below:

1. After you finish configuring AutoTracker Policies for this VLAN, the following prompt displays:

Enable IP (y):

Press **<Enter>** if you want to enable IP Routing on this virtual router port. If you do not enable IP, then this VLAN will not be able to internally route IP data. If you don't want to set up the IP router port, enter **n**, press **<Enter>** and skip to Step 10.

Note

You may enable routing of both IP and IPX traffic on this router port. If you set up dual-protocol routing, you must fill out information for both IP and IPX parameters.

2. The following prompt displays:

IP Address:

Enter the IP address for this virtual router port in dotted decimal notation or hexadecimal notation (e.g., 198.206.181.10). This IP address is assigned to the virtual router port for this VLAN. After you enter the address, press **<Enter>**.

3. The following prompt displays:

IP Subnet Mask (0xfffff00):

The default IP subnet mask (in parentheses) is automatically derived from the VLAN's IP address class. Press **<Enter>** to select the default subnet mask or enter a new subnet mask in dotted decimal notation or hexadecimal notation and press **<Enter>**.

4. The following prompt displays:

IP Broadcast Address (198.200.10.255):

The default IP broadcast address (in parentheses) is automatically derived from the VLAN's IP address class. Press **<Enter>** to select the default address or enter a new IP broadcast address in dotted decimal notation or hexadecimal notation and press **<Enter>**.

5. The following prompt displays:

Description (30 chars max):

Enter a useful description for this virtual IP router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

6. The following prompt displays:

Disable routing? (n) :

Indicate whether you want to disable routing in the VLAN. You can enable routing later through the **modvl** command.

7. The following prompt displays:

Enable NHRP? (n) :

Indicate whether you want to enable NHRP.

8. The following prompt displays:

**IP RIP Mode {Deaf (d),
Silent (s),
Active (a),
Inactive (i)} (s):**

Define the RIP mode in which the virtual router port will operate. RIP (Router Information Protocol) is a network-layer protocol that enables this VLAN to learn and advertise routes. The RIP mode can be set to one of the following:

Silent. The default setting shown in parentheses. RIP is active and receives routing information from other VLANs, but does not send out RIP updates. Other VLANs will not receive routing information concerning this VLAN and will not include the VLAN in their routing tables. Simply press **<Enter>** to select Silent mode.

Deaf. RIP is active and sends routing information to other VLANs, but does not receive RIP updates from other VLANs. This VLAN will not receive routing information from other VLANs and will not include other VLANs in its routing table. Enter **d** and press **<Enter>** to select Deaf mode.

Active. RIP is active and both sends and receives RIP updates. This VLAN will receive routing information from other VLANs and will be included in the routing tables of other VLANs. Enter **a** and press **<Enter>** to select Active mode.

Inactive. RIP is inactive and neither sends nor receives RIP updates. This VLAN will neither send nor receive routing information to/from other VLANs. Enter **i** and press **<Enter>** to select Inactive mode.

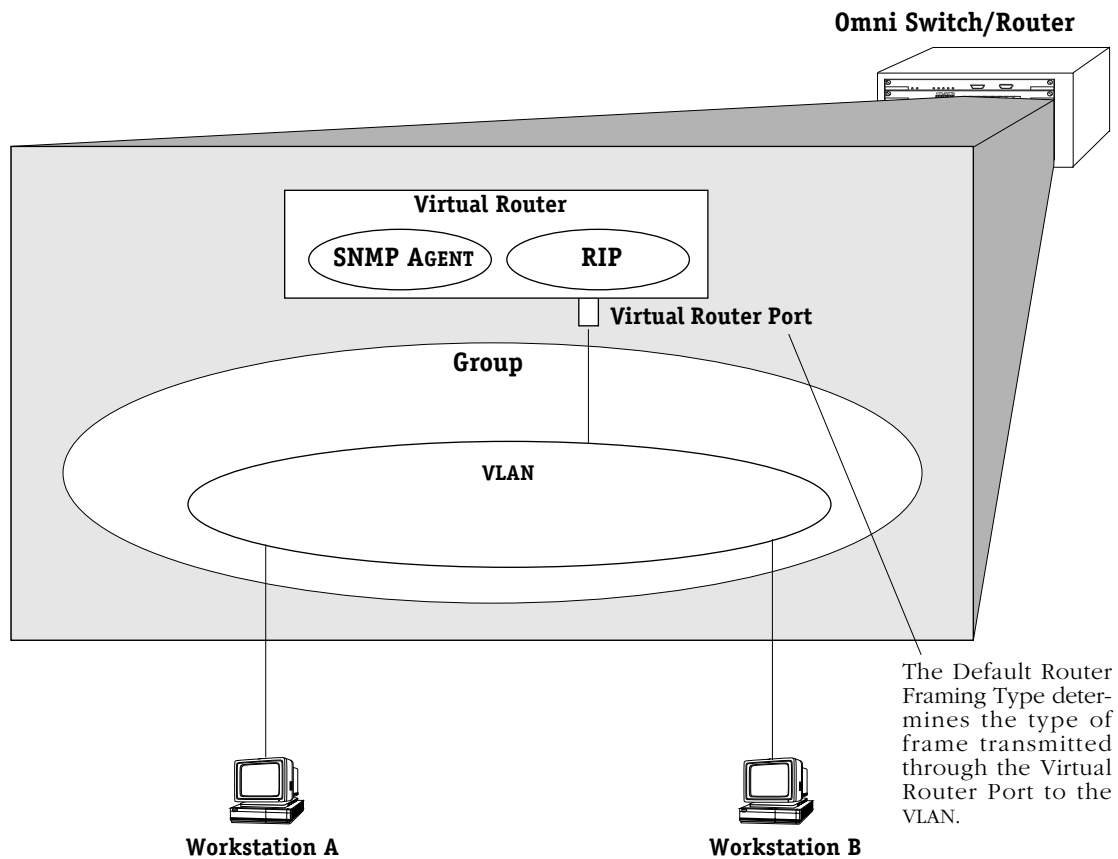
9. After you enter the RIP mode, the following prompt displays:

**Default framing type [Ethernet II(e),
fddi (f),
token ring (t),
Ethernet 802.3 SNAP (8),
source route token ring(s)} (e):**

Select the default framing type for the frames that will be generated by this router port and propagated over this VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in this VLAN. If this VLAN contains devices using encapsulation types other than those defined here, the MPX module must translate those frames, which slows throughput. The figure on the next page illustrates the Default Framing Type and its relation to Virtual Router Port communications.

After you enter the framing type a message displays indicating that this IP router port was created:

Created router port for vlan 1:3



Default Framing Type and the Virtual Router Port

10. You can now configure IPX routing on this port. The following message displays:

Enable IPX? (y) :

Press **<Enter>** if you want to enable IPX Routing on this virtual router port. If you do not enable IPX, then this VLAN will not be able to internally route IPX data. You can set up a virtual router port to route both IP and IPX traffic.

If you don't want to enable IPX routing, enter **n** and press **<Enter>**. You are now done configuring this VLAN. You can monitor activity on this VLAN through other AutoTracker commands. See later section in this chapter for more information on these commands.

11. After selecting to enable IPX, the following prompt displays:

IPX Network:

Enter the IPX network address. IPX addresses consist of eight hex digits and you can enter a minimum of one hex digit in this field. If you enter less than eight hex digits, the system prefixes your entry with zeros to create eight digits.

12. The following prompt displays:

Description (30 chars max):

Enter a useful description for this virtual IPX router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

13. After entering a description, the following prompt displays:

```
IPX RIP and SAP mode {RIP and SAP active (a)
RIP only active (r)
RIP and SAP inactive (i)}                (a):
```

Select how you want the IPX protocols, RIP (router internet protocol) and SAP (service access protocol), to be configured for this VLAN. RIP is a network-layer protocol that enables this VLAN to learn routes. SAP is also a network-layer protocol that allows network services, such as print and files services, to advertise themselves. The choices are:

RIP and SAP active. The default setting. The VLAN to which this IPX router port is attached participates in both RIP and SAP updates. RIP and SAP updates are sent and received through this router port. Simply press **<Enter>** to select RIP and SAP active.

RIP only active. The VLAN to which this IPX router port is attached participates in RIP updates only. RIP updates are sent and received through this router port. Enter an **r** and press **<Enter>** to select RIP only active.

RIP and SAP inactive. The IPX router port is active, but the VLAN to which it is attached does not participate in either RIP nor SAP updates. Enter an **i** and press **<Enter>** to select RIP only active.

14. After selecting the RIP and SAP configuration, the following prompt displays the default router framing type options:

```
Default router framing type for : {
Ethernet Media:
Ethernet II (0),
Ethernet 802.3 LLC (1),
Ethernet 802.3 SNAP (2),
Novell Ethernet 802.3 raw (3),
FDDI Media:
fdi SNAP (4),
source route fdi SNAP (5),
fdi LLC (6),
source route fdi LLC (7),
Token Ring Media:
token ring SNAP (8),
source route token ring SNAP (9),
token ring LLC (a),
source route token ring LLC (b) }      (0) :
```

Select the default framing type for the frames that will be generated by this router port and propagated over the VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in the VLAN. If the VLAN contains devices using encapsulation types other than those defined here, the MPX module must translate those frames, which slows throughput. See the figure, *Default Framing Type and the Virtual Router Port* on page 22-21 for an illustration of the Default Framing Type and its relation to Virtual Router Port communications.

15. If you chose a Source Routing frame format in the last step (options 5, 7, 9, or b), the an additional prompt displays:

**Default source routing broadcast type : {
ARE broadcasts(a), STE broadcasts(s)} (a) :**

Select how broadcasts will be handled for Source Routing. The choices are:

ARE broadcasts. All Routes Explorer, the default setting. Broadcasts are transmitted over every possible path on inter-connected source-routed rings. This setting maximizes the generality of the broadcast. Simply press **<Enter>** to select All Routes Explorer.

STE broadcasts. Spanning Tree Explorer. Broadcasts are transmitted only over Spanning Tree paths on inter-connected source-routed rings. This setting maximizes the efficiency of the broadcast. Enter an **s** and press **<Enter>** to select Spanning Tree Explorer.

After you enter framing type information a message displays indicating that this IPX router port was created:

Created router port for vlan 1:3

You have now completed the configuration of the virtual router port for this VLAN. You can monitor activity on this VLAN through other AutoTracker commands. See later section in this chapter for more information on these commands.

Modifying an AutoTracker VLAN

After you set up a VLAN you can modify its Admin Status, description, rules, and the Admin Status of each of the rules. You use the **modatvl** command to modify a VLAN as follows:

```
modatvl <Group Number>:<VLAN Number>
```

You must specify the Group and VLAN numbers and they must be separated by a colon. For example, to modify the VLAN 3 in Group 4, you would specify:

```
modatvl 4:3
```

After entering a valid **modatvl** command a screen similar to the sample below displays:

```
VLAN 4: 3 is defined as:
  1. Description = AT VLAN 3
  2. Admin Status = Enabled
  3. Rule Definition
      Rule Num Rule Type Rule Status
        1      Protocol Rule Disabled
Available options:
  1. Set VLAN Admin Status
  2. Set VLAN Description
  3. Add more rules
  4. Delete a rule
  5. Set rule Admin Status
  6. Quit
Option =
```

The first half of the display shows the current configuration of this VLAN. For example, this sample shows a VLAN 3 in Group 4 with a description, "AT VLAN 3." The VLAN is Enabled and a Protocol Rule has been set up, but this rule has not been enabled.

The second half of the displays a list of the VLAN attributes you can modify. You can modify basic information such as the Admin Status and Description. You can also add rules, delete rules, and enable or disable the rule. To modify an attribute, enter the number next to the option you want to modify and press **<Enter>**.

The following sections describe each of the six Available Options for the **modatvl** command.

Changing a VLAN's Admin Status

1. At the **Option=** prompt enter a **1** and press **<Enter>**.
2. The following prompt displays:

```
Set Admin Status to ((e)nable/(d)isable):
```

Type an **E** to enable the VLAN or a **D** to disable it. An enabled VLAN starts using policies to direct data flow. A disabled VLAN is saved, but cannot become active.

The system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Changing a VLAN's Description

1. At the **Option=** prompt enter a **2** and press **<Enter>**.
2. The following prompt displays:

Enter a new description:

Type in the revised description for this VLAN. The description can be up to 30 characters long. Press **<Enter>** when you have completed the new description.

The system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Adding More Policies for This VLAN

1. At the **Option=** prompt enter a **3** and press **<Enter>**.
2. The following menu displays:

Select rule type:

1. **Port Rule**
2. **MAC Address Rule**
3. **Protocol Rule**
4. **Network Address Rule**
5. **User Defined Rule**
6. **Binding Rule**
7. **DHCP PORT Rule**
8. **DHCP MAC Rule**

Enter rule type (1):

This is the same menu used by the **cratvl** command. This menu has eight options, some of which contain multiple branching options. This menu is documented fully in Chapter 20, "Configuring Group and VLAN Policies." Please consult that chapter for information on this menu.

When you have entered all new rule types, the system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Deleting A Policy for This VLAN

1. At the **Option=** prompt enter a **4** and press **<Enter>**.
2. The following menu displays:

Enter rule number to delete:

The rule number is listed with other information on the VLAN just after you entered the **modatvl** command. Find the number corresponding to the rule you want to delete and enter it at this prompt and press **<Enter>**. The rule is deleted and the system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Changing the Admin Status for a VLAN Policy

1. At the **Option=** prompt enter a **5** and press **<Enter>**.
2. The following menu displays:

Enter rule number:

The rule number is listed with other information on the VLAN just after you entered the **modatvl** command. Find the number corresponding to the rule you want to change and enter it at this prompt and press **<Enter>**.

3. The following menu displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Type an **E** to enable this rule or a **D** to disable it. If the rule is enabled, the VLAN will start using the rule criteria to segment data traffic.

The system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Deleting an AutoTracker VLAN

You can delete an AutoTracker VLAN. When you delete a VLAN, traffic is no longer filtered according to the VLAN's policies. Follow these steps to delete a VLAN.

1. Enter **rmatvl** followed by the Group number, a colon (:), and the VLAN number that you want to delete. For example to delete VLAN 2 in Group 3, you would enter:

rmmcvl 3:2

2. The following prompt displays:

Delete VLAN 3:2 ? (n):

Enter a **Y** and press **<Enter>** to complete the deletion of the VLAN. A message display confirming the deletion.

VLAN 3:2 deleted

Viewing AutoTracker VLANs

You can view the current status of all AutoTracker VLANs in the switch using the **atvl** command. Enter **atvl** and a table similar to the following displays.

VLAN Group :	VLAN Id	VLAN Description	Admin Status	Operational Status
3:	5	VLAN 5	Enabled	Active
3:	11	VLAN 11	Enabled	Inactive
3:	12	VLAN 12	Enabled	Inactive
3:	22	VLAN 22	Enabled	Active
3:	23	VLAN 23	Enabled	Active
3:	24	VLAN 24	Enabled	Inactive
3:	25	VLAN 25	Enabled	Inactive
3:	26	VLAN 26	Enabled	Inactive
3:	27	VLAN 27	Enabled	Inactive
3:	31	VLAN 31	Enabled	Inactive
3:	32	VLAN 32	Enabled	Inactive

VLAN Group. The Group to which this AutoTracker VLAN is assigned. The Group is specified when first creating an AutoTracker VLAN.

VLAN ID. An identification number that you assigned when you created this VLAN.

VLAN Description. A textual description that you entered to describe a VLAN when you created or modified it through **cratvl** or **modatvl**. This description is limited to 30 characters.

Admin Status. The Administrative Status for the VLAN may be enabled or disabled. You enable or disable the Administrative Status for a VLAN when you create or modify it. If the VLAN is enabled, the switch will use the policies you configured to filter traffic to the devices in this VLAN. If you disable the rule, then policies will not be used, but the parameters you set up for the VLAN will be saved.

Oper Status. The VLAN is shown as **Active** or **Inactive**. In order for an enabled VLAN to become “active” it must be able to assign a switch port to the VLAN. If the port rule is used for a VLAN, then the VLAN automatically becomes active. If any other rule is used (MAC address, protocol, etc.), then a frame matching the VLAN rule must first be received by a switch port before the VLAN is active. So, an Active VLAN requires the following:

- Admin Status must be enabled.
- A port must be assigned to the VLAN through either a port-based rule or by a device transmitting data that matches the VLAN policy.

Viewing Policy Configurations

Typing **viatrl** brings up the Policy Configuration Table, which shows the policies defined for the VLAN specified.

VLAN Group :	VLAN Id	Rule Num	Rule Type	Rule Status	Rule Definition
3:	5	1	PORT RULE	Disabled	2/7/Brg/1
3:	11	1	NET ADDR RULE	Enabled	IPX Addr = 11223344 IPX Encapsulation = Ethernet
3:	12	1	NET ADDR RULE	Enabled	DECNET Area = 13579
3:	22	1	PORT RULE	Enabled	2/7/Brg/1
3:	23	1	PORT RULE	Enabled	2/7/Brg/1
3:	24	1	MAC RULE	Enabled	082008:003002 082009:803728
3:	25	1	PROTOCOL RULE	Enabled	Protocol = IP
3:	26	1	NET ADDR RULE	Enabled	IP Addr = 131.1.2.3 IP Mask = 255.255.0.0
3:	27	1	USER RULE	Enabled	Offset = 64 Length = 2 Value = FFFF Mask = FFFF
3:	31	1	PROTOCOL RULE	Enabled	Protocol = IP
3:	32	1	NET ADDR RULE	Enabled	IPX Addr = 00000001 IPX Encapsulation = Ethernet

VLAN Group. The Group to which this AutoTracker VLAN is assigned. The Group number is specified when first creating the VLAN.

VLAN ID. An identification number that you assigned when you created this virtual LAN.

Rule Num. The number of the policy within the VLAN definition. Each rule defined for a VLAN is numbered sequentially in the order of creation. The rule number is needed when you want to modify or delete a rule definition.

Rule Type. The type of VLAN policy. The Rule Type can be a port policy (PORT RULE), MAC Address policy (MAC RULE), network address policy (NET ADDR RULE), Protocol policy (PROTOCOL RULE), or a user-defined policy (USER RULE). You set up VLAN policies when you create or modify the VLAN.

Rule Status. Indicates whether the rule for this row is Enabled or Disabled. If the rule is enabled, then the VLAN is using the rule definition to determine VLAN membership. If Disabled, then the VLAN is not using this rule to determine membership. Note that this Rule Status is different from the Admin Status for the VLAN since it controls only this specific rule within this specific VLAN. You can enable or disable the rule using the **modatvl** command.

Rule Definition. Details of this rule. For a Port Rule, this column lists the virtual interface for the Port included in the VLAN as

<slot>/<port>/<service>/<instance>

For example, the port defined for the first row in the table applies to the first bridge instance on port 7 on the module in slot 2 of the switch. For a MAC address rule, this column lists the MAC address for the device in the VLAN. For a Network Address Rule, the column will list the address (IP or IPX) and the IP Mask (IP) or the Encapsulation type (IPX). For a Protocol policy, the column list the protocol used to determine membership. And in a User-Defined rule, the offset, length, value, and mask are listed.

Viewing Virtual Ports' VLAN Membership

You can view the VLAN membership of each virtual interface in the switch. For physical LAN ports, the virtual interface is the same as a virtual port. However, when multiple services are set up for a physical port, then each service has a virtual port.

Type **viol** and a Virtual Interface Table displays similar to the one that follows. You can also specify just the slot and port number to narrow the range of ports displayed.

Virtual Interface VLAN Membership

Slot/Intf/Service/Instance	Group	Member of VLAN#
1 /1 /Rtr /1	1	1
1 /1 /Rtr /2	3	1
1 /1 /Rtr /3	3	23
1 /1 /Rtr /4	3	24
1 /1 /Rtr /5	3	25
1 /1 /Rtr /6	3	5
2 /1 /Brg /1	1	1
2 /2 /Brg /1	1	1
2 /3 /Brg /1	1	1
2 /4 /Brg /1	1	1
2 /5 /Brg /1	1	1
2 /6 /Brg /1	1	1
2 /7 /Brg /1	1	1 22
2 /8 /Brg /1	1	1
3 /1 /Brg /1	1	1
4 /1 /Brg /1	1	1
4 /2 /Brg /1	1	1
4 /3 /Brg /1	1	1
4 /4 /Brg /1	1	1
4 /5 /Brg /1	1	1
4 /6 /Brg /1	1	1
5 /1 /Brg /1	1	1

Slot/Intf/Service/Instance. Specifies the virtual interface for which AutoTracker VLAN information will be displayed. The **Slot** is the physical slot location to which the virtual interface maps. The **Intf** is the physical port to which the virtual interface maps. The **Service** is the service type for this interface. The service type may be a Router (**Rtr**), Bridge (**Brg**), Classical IP (**CIP**), FDDI Trunk (**Trk**), or an 802.10 Trunk (**T10**). **Instance** is the specific instance of this service type. These different instances are identified numerically. The first instance of a service type belonging to a physical port is identified as 1, the second instance is identified as 2, etc.

Group. The Group to which this virtual interface is assigned. The Group is specified when first creating an AutoTracker VLAN.

Member of VLAN #. The AutoTracker VLANs to which this virtual interface belongs. An interface may belong to more than one VLAN. For example, a port may contain devices using the IP Protocol and could match the Port policy of one AutoTracker VLAN and the Protocol policy of another AutoTracker VLAN. Also, physical ports always remain members of the default VLAN #1.

View VLAN Membership of MAC Devices

The **fwtl** command displays a table of learned MAC addresses and the VLAN membership of those MAC addresses. Follow these steps to view this table.

1. Enter **fwtl**.
2. The following prompt displays:

Enter Slot/Interface (return for all ports) :

Enter the slot and port for which you want to view MAC Address/VLAN information. You can also press **<Enter>** to view information on all ports in the switch.

3. The following message and prompt displays:

Total number of MAC addresses learned for Group 1: 4
Maximum number of entries to display [20] :

The top line displays the number of MAC addresses learned on this switch. This number indicates the potential number of entries you can display in the Learned MAC Address Table. The second line allows you to indicate how many of these MAC addresses you want to display. Enter the number of MAC entries you want to display or press **<Enter>** to select the default in brackets [20].

4. The Learned MAC Address/VLAN Membership Table displays as follows:

MAC Address	Slot/Intf/Service/Instance	AT VLAN Membership
0020DA:05F623	4/ /1 /Brg 1	1
0020DA:021533	4/ /1 /Brg 1	1
0020DA:0205B3	4/ /1 /Brg 1	1
0020DA:06BAD3	4/ /1 /Brg 1	1
0020DA:05F610	4/ /1 /Brg 1	1

MAC Address. The MAC address for which virtual interface and VLAN membership information will be displayed.

Slot/Intf/Service/Instance. Specifies the virtual port for which AutoTracker VLAN information will be displayed. The **Slot** is the physical slot location to which the MAC address maps. The **Intf** is the physical port to which the MAC address maps. The **Service** is the service type for this MAC address. The service type may be a Router (**Rtr**), Bridge (**Brg**), Classical IP (**CIP**), FDDI Trunk (**Trk**), or an 802.10 Trunk (**T10**). **Instance** is the specific instance of this service type. These different instances are identified numerically. The first instance of a service type belonging to a physical port is identified as 1, the second instance is identified as 2, etc.

AT VLAN Membership. The AutoTracker VLANs to which this MAC Address belongs. An MAC address may belong to more than one VLAN. For example, let's say a MAC device runs on an IPX network. It could be included in a MAC Address policy for one AutoTracker VLAN and the IPX Protocol Policy of another VLAN.

Creating a VLAN for Banyan Vines Traffic

Banyan Vines uses a fixed encapsulation for each network interface. For this reason, it is straightforward to create a VLAN for Banyan Vines traffic. For Ethernet traffic, Banyan Vines uses Ethernet II encapsulation; Token Ring uses LLC; FDDI uses SNAP. This procedure describes how to create a VLAN for Ethernet, Token, *and* FDDI traffic. Follow these steps to create a Banyan Vines VLAN:

1. Type **cratvl** at any prompt.

2. The following prompt displays:

Enter the VLAN Group id for this VLAN (1):

Enter the number for the Group to which this Banyan Vines VLAN will belong.

3. The following prompt displays:

Enter the VLAN Id for this VLAN (2):

Enter the number that will identify this VLAN within the Group specified above. By default the system displays the next available VLAN ID number. Press **<Enter>** to accept this default.

4. The following prompt displays:

Enter the new VLAN's description:

Enter a textual description that will help you identify the VLAN. For example, you might call the VLAN, "Banyan Vines VLAN." You may use up to 30 characters for this description.

5. The following prompt displays:

Enter the Admin Status for this vlan (Enable (e) / Disable (d)):

Enter whether or not you want the Administrative Status for this VLAN to be enabled or disabled. Once enabled, the switch begins using the policies you defined. A disabled VLAN is still defined (name, number, policies intact), but the switch keeps the VLAN disabled. The enable/disable status may be changed at a later time using the **modatvl** command.

6. The following menu displays:

Select rule type:

1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP PORT Rule

Enter rule type (1):

Press **3** and press **<Enter>**.

7. The following prompt displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Type **e** to enable this rule. When enabled, the VLAN will begin using the rule to determine membership of devices.

8. The following prompt displays:

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type
6. Protocol specified by DSAP and SSAP
7. Protocol specified by SNAP
```

Enter protocol type (1):

Enter a **5** to define a protocol by ether-type and press **<Enter>**.

9. The following prompt displays:

```
Enter the Ether-type value in hex:
```

10. Enter **0bad** as the Ether-type value for Ethernet II encapsulation.

11. The following prompt displays:

```
Configure more rules for this vlan (y/n):
```

Enter a **Y**. You still need to set up rules for LLC and SNAP traffic.

12. The following prompt displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP PORT Rule
```

Enter rule type (1):

Press **3** and press **<Enter>**.

13. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Type **e** to enable this rule.

14. The following prompt displays:

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type
6. Protocol specified by DSAP and SSAP
7. Protocol specified by SNAP
```

Enter protocol type (1):

Enter a **6** to define a protocol by DSAP and SSAP and press **<Enter>**.

15. The following prompt displays

Enter the DSAP value in hex:

Enter **bc** as the destination service access protocol (DSAP) value and press **<Enter>**.

16. The following prompt displays:

Enter the SSAP value in hex:

Again, enter **bc** as the source service access protocol (SSAP) value and press **<Enter>**.

17. The following prompt displays:

Configure more rules for this vlan (y/n):

Enter a **Y**. You still need to set up a rule for SNAP traffic.

18. The following prompt displays:

Select rule type:

1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP PORT Rule

Enter rule type (1):

Press **3** and press **<Enter>**.

19. The following prompt displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Type **e** to enable this rule.

20. The following prompt displays:

Select Protocol:

1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type
6. Protocol specified by DSAP and SSAP
7. Protocol specified by SNAP

Enter protocol type (1):

Enter a **7** to define a protocol by SNAP and press **<Enter>**.

21. The following prompt displays:

Enter the SNAP value in hex

Enter 00000080c4 as the desired SNAP value and press **<Enter>**.

22. The following prompt displays:

Configure more rules for this vlan (y/n):

Enter an **N**. You are done setting up rules for this VLAN. A prompt similar to the following displays:

VLAN 1:2 created successfully

23. The following prompt displays:

Enable IP (y):

Enter an **N**.

24. The following prompt displays:

Enable IPX (y):

Enter an **N**. The Banyan Vines traffic VLAN is complete.

23 Multicast VLANs

Multicast VLANs enable you to control the flooding of multicast traffic in your network. For example, you can define a multicast VLAN for all users that want to receive CNN Newscasts or any other video feed or combination of feeds.

You define the multicast traffic to be transmitted by specifying a multicast address. You define the recipients of the multicast traffic by specifying ports and/or specific MAC addresses. The members of a multicast VLAN consist of the ports specified to **receive** the multicast traffic and the ports to which MAC address recipients are connected. Instructions for creating multicast VLANs begin on page 23-4.

Note the difference between multicast VLANs and AutoTracker VLANs. In AutoTracker VLANs, devices are assigned to VLANs by examination of the frames that **originate** from those devices. The members of an AutoTracker VLAN consist of source devices that fit the VLAN's policies and the ports to which those source devices are connected.

There are several differences between the configuration of multicast VLANs and the configuration of AutoTracker VLANs. The following is a summary of points to note when configuring multicast VLANs:

- You can not configure routing for multicast VLANs. Multicast VLANs are independent broadcast domains for multicast traffic originating from a multicast address and transmitted to one or more recipients.
- Multicast VLANs allow three rules: Port, MAC Address, and multicast policy.
- There is not a default multicast VLAN. Therefore, you can define rules for all 32 available multicast VLANs. All ports (even those that eventually become part of a multicast VLAN) start off in the standard AutoTracker default VLAN #1, but they only get assigned to a multicast VLAN if you explicitly assign them to one.
- All multicast VLANs include the multicast policy. This policy specifies the multicast address. You use the other two rules—Port and MAC Address—to define the destination of the multicast traffic.

How Devices are Assigned to Multicast VLANs

If the recipients of the multicast traffic were defined using the port rule, each specified port is then marked as a member of the multicast VLAN.

If the recipients of the multicast traffic were defined using the MAC address rule to specify the MAC addresses of the receiving devices, no action is taken until a frame is received from one of those devices. When such a frame is received, the switch learns the device, adds its MAC address to the filtering database, and marks the port on which the frame was received as a member of the multicast VLAN. Note that the MAC address does not itself become a member of the multicast VLAN, even though it is a recipient of the multicast traffic. Only ports are members of multicast VLANs.

When the switch receive multicast traffic that has an address specified as a multicast address for the multicast VLAN, the traffic is switched to the ports defined as VLAN members.

◆ Please Take Note ◆

The source port of the multicast traffic (i.e., the port through which multicast traffic enters the switch) can be a member of any Group. The source port does *not* need to be a member of the same Group as recipient ports. Note that the source port does not become a member of the multicast VLAN.

Although some leakage may occur before devices are assigned to AutoTracker VLANs, no leakage occurs in conjunction with device assignment to multicast VLANs.

◆ Please Take Note ◆

There is no default multicast VLAN. Unless you explicitly create multicast VLANs, none will exist.

Multicast VLANs and Multicast Claiming

The goal of multicast claiming and multicast VLANs is the same—to free the MPX module from processing multicast traffic. Both methods off-load multicast traffic processing to the switching modules. However, multicast VLANs can be seen as a refinement to multicast claiming.

Multicast claiming claims the MAC addresses of all source devices sending multicast traffic and places those MAC addresses in the CAMs of all switching modules in a switch. Instead of claiming all multicast traffic, multicast VLANs claim only the traffic from the multicast address you specify. In addition, this multicast address is only placed in the CAMs of switching modules with destination ports that are part of the multicast VLAN.

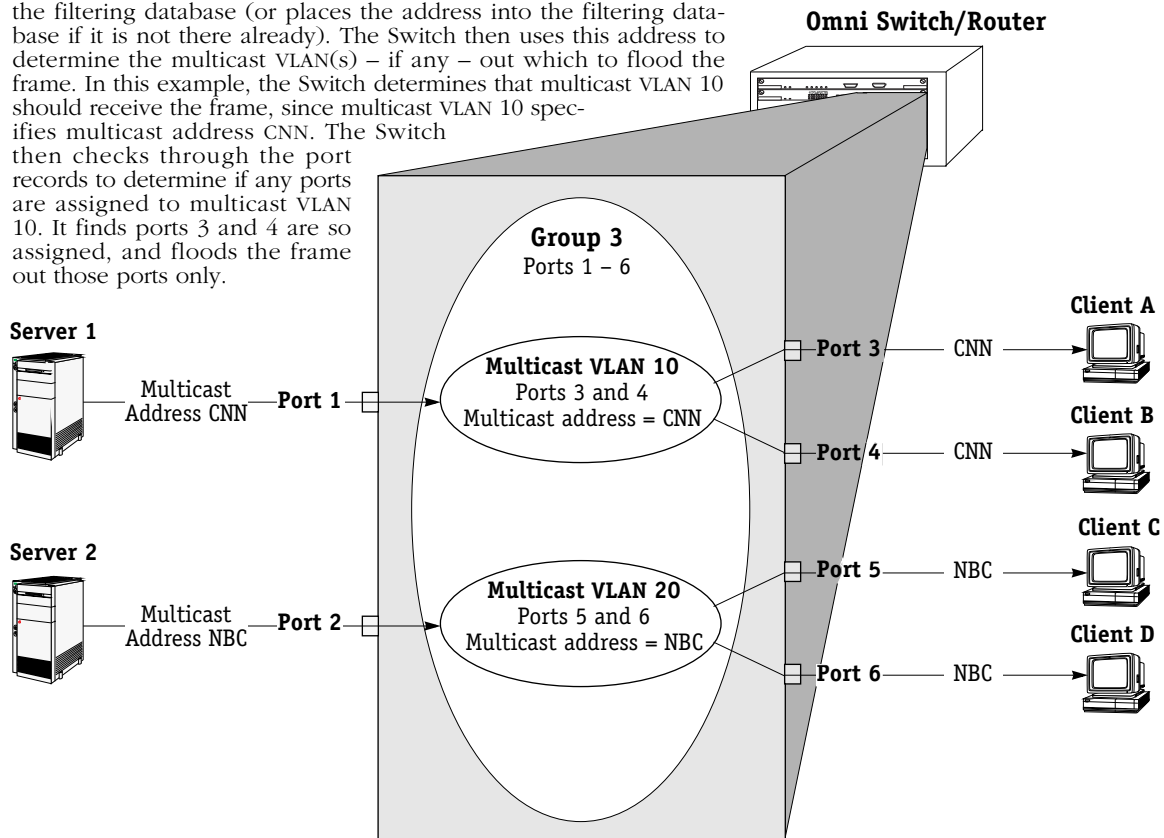
Frame Flooding in Multicast VLANs

Multicast traffic is flooded as follows in an environment that includes multicast VLANs:

- If the destination address is a multicast address, **and**
- if the destination multicast address is in the filtering database, **and**
- if the destination multicast address is a specified multicast address for a multicast VLAN, **then**

flood the traffic on all ports that have at least one multicast VLAN in common with the destination multicast address. This is illustrated below. If any of the conditions described above are untrue, the traffic is flooded as it is for normal AutoTracker VLANs.

When the Switch receives a frame with multicast destination address CNN from Server 1, the Switch locates the CNN multicast address in the filtering database (or places the address into the filtering database if it is not there already). The Switch then uses this address to determine the multicast VLAN(s) – if any – out which to flood the frame. In this example, the Switch determines that multicast VLAN 10 should receive the frame, since multicast VLAN 10 specifies multicast address CNN. The Switch then checks through the port records to determine if any ports are assigned to multicast VLAN 10. It finds ports 3 and 4 are so assigned, and floods the frame out those ports only.



For this Example, the Port Records are:

Port	VLAN Membership	MVLAN Membership
1	1	none
2	1	none
3	1	10
4	1	10
5	1	20
6	1	20

The port records show the VLAN and multicast VLAN (MVLAN) membership of each port. This table is for informational purposes only—it is not available as a UI command.

For this Example, the Filtering Database is:

MAC Address	Port	VLAN Membership	Type
CNN	n/a	10	MVLAN
NBC	n/a	20	MVLAN
Server 1	1	1	BRIDGE
Server 2	2	1	BRIDGE
Client A	3	1	BRIDGE
Client B	4	1	BRIDGE
Client C	5	1	BRIDGE
Client D	6	1	BRIDGE

The filtering database is a record of source MAC addresses, their ports of entry into the switch, and their VLAN membership. Note that the ports of entry for multicast addresses CNN and NBC are irrelevant in the filtering database. This table is for informational purposes—it is not available in the UI.

Creating Multicast VLANs

You create multicast VLANs through the AutoTracker menu options. Creating a multicast VLAN includes the following steps:

- A.** Entering basic information such as the name and number for the multicast VLAN. See *Step A. Entering Basic Information* on page 23-5 for instructions on this step.
- B.** Defining the multicast address. You define one or more multicast addresses that define the multicast stream(s) for the multicast VLAN. See *Step B. Defining the Multicast Address* on page 23-6 for instructions on this step.
- C.** Defining the recipients of multicast traffic. You may define these recipients as virtual ports or as specific MAC addresses. See *Step C. Defining the Recipients of Multicast Traffic* on page 23-7 for instructions on this step.

These steps are explained in detail below.

Step A. Entering Basic Information

1. To begin setting up a multicast VLAN type **crmcvl** at any prompt.
2. The following prompt displays:

Enter the VLAN Group id for this VLAN (1):

Enter the number for the Group to which this multicast VLAN will belong. You can create up to 32 multicast VLANs and up to 31 AutoTracker VLANs in a single Group.

3. The following prompt displays:

Enter the VLAN Id for this VLAN (5):

Enter the number that will identify this multicast VLAN within the Group specified above. Up to 32 multicast VLANs may belong to the same Group. By default the system displays the next available VLAN ID number.

◆ **Note** ◆

Unlike AutoTracker VLANs, you can configure rules for the multicast VLAN #1. There is not a default multicast VLAN, so multicast VLAN #1 is treated the same as the other 31 possible multicast VLANs.

Press **<Enter>** to accept this default.

4. The following prompt displays:

Enter the new VLAN's description:

Enter a textual description that will help you identify the multicast VLAN. For example, if you know this multicast VLAN will be composed of only workstations receiving CNN news feeds, you might call the multicast VLAN "CNN MVLAN." You may use up to 30 characters for this description.

5. The following prompt displays:

Enter the Admin Status for this vlan (Enable (e) / Disable (d)):

Enter whether or not you want the Administrative Status for this multicast VLAN to be enabled or disabled. Once enabled, the switch begins using the policies you defined. A disabled multicast VLAN is still defined (name, number, policies intact), but the switch keeps the multicast VLAN disabled. The enable/disable status may be changed at a later time using the **modmcvl** command.

◆ **Note** ◆

A multicast VLAN may not always be operational even when its Admin Status is enabled. A multicast VLAN's operation may be disabled by its switches because devices in the multicast VLAN cease transmitting data, among other reasons.

After you enter the administrative status, additional prompts display that allow you define the multicast address. See the next section, *Step B. Defining the Multicast Address* on page 23-6 for further instructions.

Step B. Defining the Multicast Address

The multicast address is an address that identifies a multicast traffic stream, such as CNN News.

◆ Please Take Note ◆

The source port of the multicast traffic (i.e., the port through which multicast traffic enters the switch) can be a member of any Group. The source port does *not* need to be a member of the same Group as recipient ports. Note that the source port does not become a member of the multicast VLAN.

1. After you enter the administrative status for this multicast VLAN, the following prompt displays:

Configure the Multicast Address Rule
Set Rule Admin Status to [(e)nable/(d)isable] (d):

Indicate whether you want to enable or disable this multicast Address Rule. If you enable this rule, AutoTracker will use the address to flood multicast traffic. Enter an **E** (enable) or a **D** (disable) and press **<Enter>**.

If you disable the rule, then this address will not be used to flood multicast traffic, but the parameters you set up will be saved. This Admin Status is different from the Admin Status for the multicast VLAN as it controls only this specific rule within this specific multicast VLAN. You can enable or disable the rule at a later time using the **modmctl** command.

2. The following prompt displays:

Enter the Multicast addresses (AABBCC:DDEEFF) in Canonical format
(Enter save to end):

Enter one or more multicast addresses, separated by spaces. The address must be a multicast address. If you enter too many characters, the system truncates the address. The switch will flood all traffic from the address(es) you specify here to the ports and/or MAC addresses you define as recipients in Step C.

All multicast MAC addresses must consist of 12 hex digits. In all valid multicast addresses, the least significant bit of the most significant byte is set to 1. Addresses with this bit unset will be rejected.

Most Significant Byte

x x x x x x x 1



least significant bit
must be set to 1

Structure of Multicast Address

When you have entered the final MAC address press **<Enter>**, and type **save** at the prompt.

Next, a menu displays prompting you to select the rules governing membership in this multicast VLAN. Go on to the next section, *Step C. Defining the Recipients of Multicast Traffic* on page 23-7 to continue setting up this multicast VLAN.

Step C. Defining the Recipients of Multicast Traffic

You can define the recipients of multicast traffic by virtual port or MAC address. You define these recipients as policies for this multicast VLAN. The available policies for recipients are Port and MAC Address. You can use both rules within a single multicast VLAN. For example, you might want to flood multicast traffic to all devices attached to one switch port, but only a few devices attached to other switch ports. In this case, you could use a Port rule for the devices on the port where all devices receive the multicast traffic, and then the MAC address rule to flood multicast traffic only to specific devices attached to the other ports on the switch.

Follow the directions in one of the following sections for the rule type you want to define.

Defining Recipients By Port

After you define the multicast address, the following menu displays:

```

Select rule type:
1. Port Rule
2. MAC Address Rule
3. Multicast Address Rule

```

```
Enter rule type (1):
```

1. Press **<Return>**.
2. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Indicate whether or not you want to enable this rule. Type **e** to enable or **d** to disable. If you enable the rule, the multicast VLAN will use it to determine membership of devices. If you disable the rule, then this rule will not be used in assigning devices to this multicast VLAN, but the parameters you set up for the multicast VLAN will be saved. This Admin Status is different from the Admin Status for the multicast VLAN as it controls only this specific rule within this specific multicast VLAN. You can enable or disable the rule at a later time using the **modmctl** command.

3. The following prompt displays:

```
Enter the list of port in Slot/Int/Service/Instance format:
```

Enter the ports that you want to receive multicast traffic for this multicast VLAN. You may enter multiple ports at a time. You can include a total of 255 ports per switch in a port-based multicast VLAN. Use the **<slot>/<port>** format. For example, to include port 7 from the module in slot 2, you would enter **2/7**. (The service and instance numbers are not necessary for specifying physical ports. They are only necessary when specifying logical or virtual ports, which normally only differ from physical ports in more complex configurations, such as ATM LAN Emulation.)

4. The following prompt displays:

```
Configure more rules for this vlan (y/n):
```

You can set up multiple rules for the same multicast VLAN. Enter a **Y** here if you want to set up more rules in addition to the port rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this multicast VLAN. If you enter **N**, you will receive a message, similar to the one below, indicating that the multicast VLAN was set up.

```
VLAN 3:23 created successfully
```

5. If you are done setting up rules for this multicast VLAN, then your multicast VLAN is set up. You can monitor activity on these multicast VLANs through other AutoTracker commands. See later sections in this chapter for information on these commands.

Defining Recipients By MAC Address

After you define the multicast address, the following menu displays:

Select rule type:

- 1. Port Rule**
- 2. MAC Address Rule**
- 3. Multicast Address Rule**

Enter rule type (1):

1. Press **2** and **<Return>**.
2. The following prompt displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Indicate whether or not you want to enable this rule. Type **e** to enable or **d** to disable. If you enable the rule, the multicast VLAN will use it to determine membership of devices. If you disable the rule, then this rule will not be used in assigning devices to this multicast VLAN, but the parameters you set up for the multicast VLAN will be saved. This Admin Status is different from the Admin Status for the multicast VLAN as it controls only this specific rule within this specific multicast VLAN. You can enable or disable the rule at a later time using the **modmctl** command.

3. The following prompt displays:

Enter the list of MAC addresses (Enter save to end):

Enter the MAC addresses that you want to receive multicast traffic for this multicast VLAN. Separate addresses by a space. When you have entered the final MAC address, leave a space and type **save**.

4. The following prompt will display:

Configure more rules for this vlan (y/n):

You can set up multiple rules for the same multicast VLAN. Enter a **Y** here if you want to set up more rules in addition to the port rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this multicast VLAN. If you enter **N**, you will receive a message, similar to the one below, indicating that the multicast VLAN was set up.

VLAN 3:24 created successfully

5. If you are done setting up rules for this multicast VLAN, then your multicast VLAN is set up. You can monitor activity on these multicast VLANs through other AutoTracker commands. See later sections in this chapter for information on these commands.

Modifying Multicast VLANs

After you set up a multicast VLAN you can modify its Admin Status, description, rules, and the Admin Status of each of the rules. You use the **modmctl** command to modify a multicast VLAN as follows:

modmctl <Group Number>:<VLAN Number>

You must specify the Group and multicast VLAN number and they must be separated by a colon. For example, to modify multicast VLAN 2 in Group 2, you would specify:

modmctl 2:2

After entering a valid **modmctl** command, a screen similar to the following sample displays:

```

VLAN  2: 2 is defined as:
  1.   Description    = MVLAN 2
  2.   Admin Status  = Enabled
  3.   Rule Definition
      Rule Num  Rule Type    Rule Status
        1      Port Rule    Enabled
        2      Multicast Rule Enabled
Available options:
  1.   Set VLAN Admin Status
  2.   Set VLAN Description
  3.   Add more rules
  4.   Delete a rule
  5.   Set rule Admin Status
  6.   Quit
Option =

```

The first half of the display shows the current configuration of this multicast VLAN. For example, this sample shows multicast VLAN 2 in Group 2 with a description, “MVLAN 2.” The multicast VLAN is Enabled and a Port Rule has been set up and it is enabled.

The second half of the display shows a list of the multicast VLAN attributes you can modify. You can modify basic information such as the Admin Status and Description. You can also add rules, delete rules, and enable or disable a rule. To modify an attribute, enter the number next to the option you want to modify and press **<Enter>**.

The following sections describe each of the six Available Options for the **modmctl** command.

Changing a VLAN’s Admin Status

1. At the **Option=** prompt enter a **1** and press **<Enter>**.
2. The following prompt displays:

Set Admin Status to ((e)nable/(d)isable):

Type an **e** to enable the multicast VLAN or a **d** to disable it. An enabled VLAN starts using policies to direct data flow. A disabled multicast VLAN is saved, but can not become active.

The system returns to the **Available Options** menu. You can modify more attributes for this multicast VLAN, or quit modifying the multicast VLAN by typing a **6**.

Changing a VLAN's Description

1. At the **Option=** prompt enter a **2** and press **<Enter>**.
2. The following prompt displays:

Enter a new description:

Type in the revised description for this multicast VLAN. The description can be up to 30 characters long. Press **<Enter>** when you have completed the new description.

The system returns to the **Available Options** menu. You can modify more attributes for this multicast VLAN, or quit modifying the multicast VLAN by typing an **6**.

Adding More Policies for This VLAN

1. At the **Option=** prompt enter a **3** and press **<Enter>**.
2. The following menu displays:

Select rule type:

1. **Port Rule**
2. **MAC Address Rule**
3. **Multicast Address Rule**

Enter rule type (1):

This is the same menu used by the **crmcvl** command. This menu has three options, some of which contain multiple branching options. This menu is documented fully in the section, *Step C. Defining the Recipients of Multicast Traffic* on page 23-7. Please consult this section for information on this menu.

When have entered all new rule types, the system returns to the **Available Options** menu. You can modify more attributes for this multicast VLAN, or quit modifying the multicast VLAN by typing an **6**.

Deleting A Policy for This VLAN

1. At the **Option=** prompt enter a **4** and press **<Enter>**.
2. The following menu displays:

Enter rule number to delete:

The rule number is listed with other information on the multicast VLAN just after you entered the **modmctl** command. Find the number corresponding to the rule you want to delete and enter it at this prompt and press **<Enter>**. The rule is deleted and the system returns to the **Available Options** menu. You can modify more attributes for this multicast VLAN, or quit modifying the multicast VLAN by typing a **6**.

Changing the Admin Status for a VLAN Policy

1. At the **Option=** prompt enter a **5** and press **<Enter>**.
2. The following menu displays:

Enter rule number:

The rule number is listed with other information on the multicast VLAN just after you entered the **modmcvl** command. Find the number corresponding to the rule you want to change and enter it at this prompt and press **<Enter>**.

3. The following menu displays:

Set Rule Admin Status to ((e)nable/(d)isable):

Type an **e** to enable this rule or a **d** to disable it. If the rule is enabled, the multicast VLAN will start using the rule criteria to segment data traffic.

The system returns to the **Available Options** menu. You can modify more attributes for this multicast VLAN, or quit modifying the multicast VLAN by typing a **6**.

Deleting a Multicast VLAN

You can delete a multicast VLAN. When you delete a multicast VLAN, multicast traffic is no longer flooded to the recipients you defined. Follow these steps to delete a multicast VLAN.

1. Type **rmmcvl** followed by the Group number, a colon (:), and the multicast VLAN number that you want to delete. For example to delete multicast VLAN 2 in Group 3, you would type:

rmmcvl 3:2

2. The following prompt displays:

Delete VLAN 3:2 ? (n):

Enter a **y** and press **<Enter>** to complete the deletion of the multicast VLAN. A message display confirming the deletion.

VLAN 3:2 deleted

Modifying a Multicast Address Policy

After you create a multicast VLAN, you can modify the multicast address policy by adding more addresses through the **modmctl** command. However, you can not add an existing multicast address. Follow the steps outlined in *Modifying Multicast VLANs* on page 23-9 and the steps for *Adding More Policies for This VLAN* on page 23-10. Continue with the procedure below.

The following menu displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Multicast Address Rule
```

```
Enter rule type (1):
```

1. Press **3** and **<Return>**.
2. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Indicate whether or not you want to enable this rule. Type **e** to enable or **d** to disable. If you disable the rule, then the multicast addresses you enter will not be used to flood traffic, but the parameters you set up for the multicast VLAN will be saved. This Admin Status is different from the Admin Status for the multicast VLAN as it controls only this specific rule. You can enable or disable the rule at a later time using the **modmctl** command.

3. The following prompt displays:

```
Enter the list of MAC addresses (Enter save to end):
```

Enter one or more multicast addresses. Separate addresses by a space. When you have entered the final multicast address, leave a space and type **save**.

4. The following prompt will display:

```
Configure more rules for this vlan (y/n):
```

You can set up multiple rules for the same multicast VLAN. Enter a **Y** here if you want to set up more rules in addition to the multicast address rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this multicast VLAN. If you enter **N**, you will receive a message, similar to the one below, indicating that the multicast VLAN was set up.

```
VLAN 3:24 created successfully
```

Viewing Multicast VLANs

You can view the current status of all multicast VLANs in the switch using the **mcvl** command. Type **mcvl** and a table similar to the following displays:

VLAN Group :	VLAN Id	VLAN Description	Admin Status	Operational Status
3:	5	MVLAN 5	Enabled	Active
3:	11	MVLAN 11	Enabled	Inactive
3:	12	MVLAN 12	Enabled	Inactive
3:	22	MVLAN 22	Enabled	Active
3:	23	MVLAN 23	Enabled	Active
3:	24	MVLAN 24	Enabled	Inactive
3:	25	MVLAN 25	Enabled	Inactive
3:	26	MVLAN 26	Enabled	Inactive
3:	27	MVLAN 27	Enabled	Inactive
3:	31	MVLAN 31	Enabled	Inactive
3:	32	MVLAN 32	Enabled	Inactive

VLAN Group. The Group to which this multicast VLAN is assigned. The Group is specified when first creating a multicast VLAN.

VLAN ID. An identification number that you assigned when you created this multicast VLAN.

VLAN Description. A textual description that you entered to describe a multicast VLAN when you created or modified it through **crmcvl** or **modmcvl**. This description is limited to 30 characters.

Admin Status. A multicast VLAN can be enabled or disabled. You enable or disable a multicast VLAN when you create or modify it. If the multicast VLAN is enabled, AutoTracker floods multicast traffic to the recipients you specified when setting up the multicast VLAN. If the multicast VLAN is disabled, the multicast traffic is not flooded as you specified; however, the parameters you set up for the multicast VLAN are saved.

Oper Status. The multicast VLAN is shown as active or inactive. In order for an enabled multicast VLAN to become “active” it must be able to assign a switch port to the multicast VLAN. If the port rule is used for a multicast VLAN, then the multicast VLAN automatically becomes active. If you defined multicast traffic recipients by MAC address only, then a frame destined for a defined MAC address must first be received by a switch port before the multicast VLAN is active. An active multicast VLAN requires the following:

- Admin Status must be enabled.
- A port must be assigned to the multicast VLAN through either a port-based rule or by a device transmitting data that matches the multicast VLAN policy.

Viewing Multicast VLAN Policies

You can view the current multicast VLAN policies and their status using the **vimcrl** command. Type **vimcrl** and a Policy Configuration Table displays similar to the following:

VLAN Group :	VLAN Id	Rule Num	Rule Type	Rule Status	Rule Definition
3:	5	1	PORT RULE	Disabled	2/7/Brg/1
3:	5	2	MCAST	Disabled	072467:0034ab
3:	22	1	PORT RULE	Enabled	2/7/Brg/1
3:	22	2	MCAST	Enabled	080027:0135de1
3:	23	1	PORT RULE	Enabled	2/7/Brg/1
3:	23	2	MCAST	Enabled	050034:000017
3:	24	1	MAC RULE	Enabled	082008:003002 082009:803728
3:	24	2	MCAST	Enabled	053967:0126af5

VLAN Group. The Group to which this multicast VLAN is assigned. The Group is specified when first creating a multicast VLAN.

VLAN ID. An identification number that you assigned when you created this multicast VLAN.

Rule Num. The number for this rule within the multicast VLAN definition. Each rule defined for a multicast VLAN is numbered sequentially in the order of creation. The rule number is needed when you want to modify or delete a rule definition.

Rule Type. The type of multicast VLAN rule. For multicast VLANs, the rule type can be PORT RULE, MAC RULE, or MULICAST RULE. Each multicast VLAN by definition will contain a multicast rule. The multicast rule defines the multicast address. In addition, the multicast VLAN contains either a Port-based rule, MAC address rule, or both a Port and MAC address rule. The Port and MAC address rules define the recipients of multicast traffic.

Rule Status. Indicates whether the rule for this row is Enabled or Disabled. If the rule is enabled, then the switch is using the rule definition to determine multicast traffic flooding. If Disabled, then the switch is not using this rule to regulate multicast traffic flow. Note that this Rule Status is different from the Admin Status for the multicast VLAN since it controls only this specific rule within this specific multicast VLAN. You can enable or disable the rule using the **modmctl** command.

Rule Definition. Details of this rule. For a Port Rule, this column lists the virtual interface for the Port that is a recipient of the multicast traffic as

```
<slot>/<port>/<service>/<instance>
```

For example, the port defined for the first row in the table applies to the first bridge instance on port 7 on the module in slot 2 of the switch. For a MAC address rule, this column lists the MAC address for the recipient of the multicast traffic. For a multicast Rule, this column lists the multicast address.

Viewing the Virtual Interface of Multicast VLANs

You can view the multicast VLAN membership of each virtual interface in the switch. In most cases the virtual interface is the same as a virtual port. However, when multiple services are set up for a virtual port, then each service may be split into one or more instances.

Type **vimcvi** and a Virtual Interface Table displays similar to the one that follows. You can also specify just the slot and port number to narrow the range of ports displayed.

Virtual Interface VLAN Membership

Slot/Intf/Service/Instance				Group	Member of VLAN#
1	/1	/Rtr	/1	1	1
1	/1	/Rtr	/2	3	23
1	/1	/Rtr	/3	3	24
2	/1	/Brg	/1	1	23
2	/7	/Brg	/1	1	22
4	/1	/Brg	/1	1	24
5	/1	/Brg	/1	1	22

Slot/Intf/Service/Instance. Specifies the virtual interface for which multicast VLAN information will be displayed. The **Slot** is the physical slot location to which the virtual interface maps. The **Intf** is the physical port to which the virtual interface maps. The **Service** is the service type for this interface. The service type may be a Router (**Rtr**), Bridge (**Brg**), Classical IP (**CIP**), FDDI Trunk (**Trk**), or an 802.10 Trunk (**T10**). **Instance** is the specific instance of this service type. These different instances are identified numerically. The first instance of a service type belonging to a physical port is identified as 1, the second instance is identified as 2, etc.

Group. The Group to which this virtual interface is assigned. The Group is specified when first creating a multicast VLAN.

Member of VLAN #. The multicast VLANs to which this virtual interface belongs. An interface may belong to more than one multicast VLAN. For example, if you set up a multicast VLAN for CNN News and another for NBC News, you may want certain ports to receive both multicast traffic streams.

24 AutoTracker VLAN Application Examples

This chapter provides specific examples of AutoTracker VLANs in various network configurations. These examples illustrate basic concepts about AutoTracker and highlight issues that can arise when AutoTracker is used in different network situations.

- *Application Example 1* illustrates a network organized according to logical policies and explains the benefits of a logical network organization.
- *Application Example 2* explains unique characteristics of IPX networks that must be considered when using AutoTracker IPX network address VLANs.
- *Application Example 3* highlights an issue concerning translated frames and AutoTracker IPX network address VLANs.
- *Application Example 4* explains how routing works generally in IPX networks and explains how to avoid an exception condition in which AutoTracker can affect the behavior of an IPX-routed network.
- *Application Example 5* explains why a port-based policy may be required for a VLAN – in addition to any other policies defined for that VLAN – to establish communications in some network situations, such as traversing a backbone.

Application Example 1

VLANs Based on Logical Policies

Example 1 shows a network organized logically. The network is organized according to IP networks, but this organization is achieved through the application of logical policies rather than physical segmentation. The use of logical policies provides the flexibility of moving IP users from segment to segment and preserving their original VLAN membership – without reconfiguring AutoTracker or the workstations.

Group and VLAN Membership


The network shown in Example 1 contains one Group – Group 2 – that consists of ports 1, 2, 3, and 4. Note that a Group defines a physical area – a set of ports – within the network. When VLANs with logical policies are created within a Group, the logical policies are applied to traffic received from all ports within the Group – but not to traffic from ports outside the Group – to determine if any source device should be a VLAN member.

As shown on the facing page, two VLANs were created within Group 2, each with a logically-based Network Address policy. The Network Address policy for VLAN A defines IP network 125.0.0.0 and the Network Address policy for VLAN B defines IP network 130.0.0.0. All traffic received on ports 1, 2, 3, and 4 will be checked for possible membership in these two VLANs.

Routing was enabled on both VLAN A and VLAN B so that traffic can move between the two VLANs, as is shown in this example by the presence of the internal IP router.

Benefits

This network configuration shown in this example provides flexibility. As explained on the following page, this logical network organization enables the Network Manager to move IP users between segments while preserving their original VLAN membership – without reconfiguring AutoTracker or the workstations.



Workstation 130.0.0.11 has been moved from the segment connected to port 4 to the segment connected to port 2. When workstation 130.0.0.11 transmits its first frame from its new location, the switch automatically places it into its original VLAN, VLAN B, because VLAN B has a network address rule that places all devices with network address 130.0.0.0 into VLAN B.

Both VLAN A and VLAN B are now active on port 2. In addition, VLAN B is now active on multiple ports – ports 2, 3, and 4. However, this does not cause confusion.

As an example, if workstation 125.0.0.1 (in VLAN A) wants to talk to workstation 130.0.0.11 (in VLAN B), workstation 125.0.0.1 ARPs for workstation 130.0.0.11's MAC address. The address returned is that of workstation 125.0.0.1's default gateway, which is VLAN A's internal IP router, 125.0.0.35. Workstation 125.0.0.1 transmits its frame to this address and the internal IP router routes the frame to VLAN B.

When VLAN B's internal IP router receives the frame addressed to workstation 130.0.0.11, it ARPs for workstation 130.0.0.11's MAC address if it does not already know it. The switch's filtering database identifies the port through which this MAC address can be reached. The frame sent by workstation 125.0.0.1 to workstation 130.0.0.11 is correctly transmitted to port 2.

Application Example 2

VLANs in IPX Networks

Example 2 illustrates the use of AutoTracker VLANs in IPX networks – specifically, VLANs based on IPX network address rules. IPX networks have unique characteristics that must be considered when configuring VLANs based on network address rules.

Encapsulation Type in IPX Networks

The encapsulation type a MAC station uses is very important in IPX networks, because a close relationship exists between encapsulation type and IPX network number. In IPX networks, a network number and an encapsulation type are configured for each segment. When two IPX servers share the same LAN segment, they must have the same network number and the same encapsulation type in order to communicate. In addition, only clients and servers that use the same encapsulation type can communicate. (The Omni Switch/Router removes this restriction somewhat through MAC-layer translations, which will not be discussed at this time.)

In summary, network number and encapsulation type define a broadcast domain in an IPX network that is analogous to a LAN – or a VLAN. (Remember that VLANs have the same characteristics as LANs, with the exception that VLANs can span multiple segments as LANs cannot.)

An encapsulation type is configured within each IPX client prior to bootup on the network. An IPX client acquires its network number dynamically from an IPX server (or from an intervening router) using a “Get_Nearest_Server” mechanism. Upon bootup, each client sends a query seeking the nearest server that uses the same encapsulation type as the client. Only those servers using the same encapsulation type respond to the query. (An intervening router can also respond to the query: routers traditionally interconnect LAN segments and can use different encapsulation types for different networks.) This means that IPX clients do not know their network numbers at bootup, but rather acquire their network numbers after they have communicated with IPX servers or with an intervening router.

VLAN Assignment in IPX Networks

The close relationship between encapsulation type and network number in IPX networks is the main reason AutoTracker’s IPX network address policy requires you to specify both a network number and an encapsulation type. The Omni Switch/Router assigns devices to IPX network address VLANs as follows:


- **IPX servers.** Frames from an IPX server always contain information on the server’s network number, so the Omni Switch/Router can always assign IPX servers to the correct VLAN based on the server’s network number.
- **IPX clients.** As explained previously, IPX clients do not know their network number at bootup and so cannot, initially, be assigned to VLANs based on their network number. For this reason the Omni Switch/Router initially assigns clients to IPX network address VLANs based on their encapsulation type. An example of this is shown on the facing page. Once an IPX client communicates with a server or an intervening router, learns its network number and begins transmitting frames with that number, it is removed from all previously-assigned IPX network address VLANs (but not from VLANs of other policy types) and placed into the correct IPX network address VLAN according to network number.

So How Do I Avoid Conflicts?

As an example, IPX defines four different types of Ethernet encapsulation: Ethernet-II, 802.2, SNAP, and IPX 802.3 (also referred to as “raw”). So, what do you do to avoid conflicts when you have more than four servers and they use different encapsulation types? The solution is to put each server into a different VLAN, as shown in the example on the facing page.

continued ...

IPX VLAN Assignment at Bootup



IPX Servers
 IPX servers always know their network number, so IPX servers are assigned to VLANs according to network number.

IPX Client Workstations
 When IPX clients boot up, their encapsulation types are known but their network numbers are not. Therefore, IPX clients are initially assigned to VLANs according to encapsulation type. This is the reason Clients 1 and 2 (which use Ethernet-II encapsulation) are assigned to VLANs B and C (which both specify Ethernet-II encapsulation).

Once an IPX client communicates with a server or an intervening router, learns its network number and begins transmitting frames with that number, it is removed from all previously-assigned IPX VLANs and placed into a single IPX VLAN according to network number. Client 1 and Client 2 will be reassigned to either VLAN B or VLAN C when their respective network numbers are known.

IPX Client	VLAN Membership
Client 1	both B & C initially, then either B or C when network number is known
Client 2	both B & C initially, then either B or C when network number is known
Client 3	D
Client 4	E
Client 5	F
Please note that all ports in Group 3 are also members of ports 3's default VLAN #1.	

Application Example 2

In this example one Group was created – Group 3 – that includes all ports to which IPX servers and clients are connected. Within this Group five VLANs were created, one for each server:



When the Omni Switch/Router receives frames from the five servers, each server is assigned to the appropriate VLAN and no conflict occurs. IPX routing is enabled for each VLAN – with appropriate framing specified – so that traffic can route between the VLANs.

When a client workstation boots up and queries for a server, the Omni Switch/Router assigns the client to the appropriate VLAN(s) based on encapsulation type. If the client uses 802.2 encapsulation, SNAP encapsulation, or IPX 802.3 encapsulation, VLAN assignment is simple: the client is assigned to VLAN D (802.2 encapsulation), VLAN E (SNAP encapsulation), or VLAN F (IPX 802.3 encapsulation), respectively.

However, when a client workstation using Ethernet-II encapsulation boots up and queries for a server, the Omni Switch/Router initially assigns the client to both VLAN B and VLAN C, since both of these VLANs specify Ethernet-II encapsulation. However, the Omni Switch/Router recognizes that the client's frame is a "Get_Nearest_Server" query and remembers that the client is in search of its network number. While the client remains in this transitional state, it remains assigned to all VLANs that specify Ethernet-II encapsulation. Once the client has received response from a server or servers or from an intervening router, the client selects its network number and begins transmitting frames with the network number embedded. The Omni Switch/Router detects these frames, removes the client from all previously-assigned IPX network address VLANs (but not from VLANs of other policy types) and assigns it to the proper IPX network address VLAN according to network number.

Please Take Note

IPX clients often are not particular about the server to which they attach. However, clients can select a preferred server if the **/PS** (preferred server name) option is included in their start-up script.

Why is this Solution Recommended?

As as been explained, isolating each IPX server in its own IPX network address VLAN is the recommended way to avoid conflicts. No problems occur if a client receives broadcast and multicast traffic from multiple servers, especially for the brief period that the client remains in a transitional state in search of a server.

Problems do occur if two servers with different network numbers and the same encapsulation type are members of the same VLAN, because each server will detect the other's frames, notice conflicting network numbers for the same VLAN, and respond with a router configuration error. For this reason it is not advisable to create four VLANs based on IPX network address policies within the same Group, each configured for one of the four encapsulation types. It is important to isolate the servers, but it is not important to isolate the clients – at least immediately.

While it is not important to isolate IPX clients immediately at bootup, it is desirable to isolate them as soon as possible. Isolating clients – rather than letting them remain in multiple VLANs that specify the same encapsulation type – increases efficiency and reduces broadcast and multicast traffic in the network. If a client remains in multiple VLANs that specify the same encapsulation type, the client receives all broadcast and multicast traffic from each server using that encapsulation type, even though the client only communicates with the server that shares its network number. In addition, when a VLAN is extended across a WAN backbone, it is wasteful and inefficient to transmit unnecessary frames across the WAN. For these reasons, as soon as a client learns its network number and begins transmitting frames with that number, the Omni Switch/Router removes the client from all previously-assigned IPX network address VLANs and assigns it to a single IPX VLAN according to network number.

Application Example 3

IPX Network Address VLANs and Translated Frames

Application Example 3 shows two IPX networks connected over a bridged FDDI ring spanning two Omni Switch/Routers. VLAN B exists in both switches and specifies an IPX network address policy of network number 100 and Ethernet-II encapsulation.

The Problem

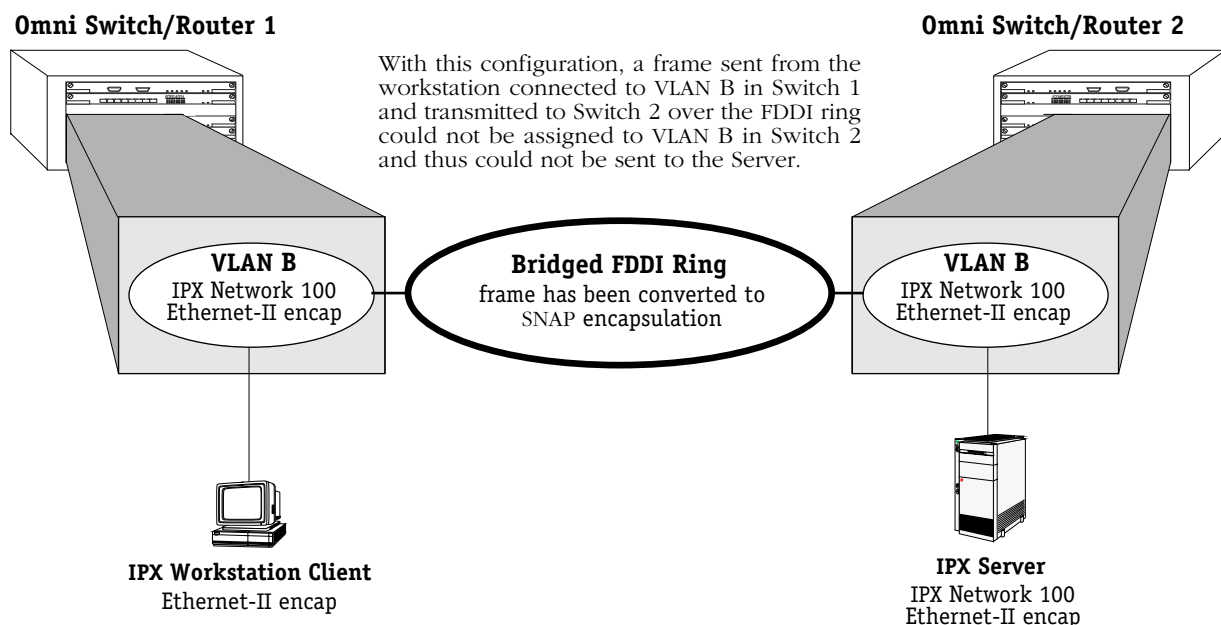
In the figure below, when the IPX client connected to Switch 1 boots up and sends a frame seeking a server, Switch 1 assigns the frame to VLAN B (since VLAN B specifies Ethernet-II encapsulation) and then converts the frame to SNAP encapsulation so that the frame can traverse the FDDI ring. When the frame arrives at Switch 2, the network number is not available (since, as previously explained, IPX clients do not know their network number at bootup) and the frame's encapsulation is no longer Ethernet-II – it is now SNAP. Because the IPX network address policy selects VLAN members according to network number and encapsulation, Switch 2 cannot assign the frame to VLAN B and send it to the IPX server.

The Solution

The solution for this problem is to specify a second encapsulation type for VLAN B in addition to Ethernet-II – for example, SNAP encapsulation. If VLAN B specifies Ethernet-II or SNAP encapsulation, the frame will match the network address policy for VLAN B when it arrives at Switch 2 and can thus be assigned to VLAN B and sent to the server. Note that the second encapsulation type must be specified for VLAN B in both Switches, to accommodate data transmission in either direction.

Please Take Note

This situation can occur whenever translations occur, such as with bridged FDDI rings or Token Rings. If you are using FDDI trunking you do not need to specify a second encapsulation policy for IPX network address VLANs, because trunked frames are not translated.



Application Example 4

Routing in IPX Networks

How Routing Works Generally

AutoTracker “activates” a VLAN – and its internal router interface – when the first port is assigned to the VLAN. If a VLAN has a port policy, AutoTracker assigns the specified port(s) and activates the VLAN immediately. If a VLAN has a logical policy, AutoTracker assigns the first port to the VLAN when a frame is received from a source device that matches the VLAN’s policy. When such a frame is received, the source device – and the port to which that device is connected – are assigned to the VLAN and the VLAN is activated.

Until a port is assigned to a VLAN, that VLAN is maintained in an inactive state and its internal router port is inactive – even if routing was enabled by the user. Use of a VLAN’s routing service is “on-demand” and AutoTracker does not enable routing until a port is present that might require it. When AutoTracker assigns the first port to a particular VLAN, it activates that VLAN and its routing service (as long as routing was enabled by the user).

Once AutoTracker has established devices’ VLAN assignments and activated the appropriate VLAN routing services, it does not participate in the routing process. Routing works correctly as long as the policies of the IPX protocol were followed – with the exception below.

The Exception

There is one scenario in which AutoTracker affects the behavior of an IPX-routed network. This situation occurs when an IPX server is a member of any VLAN with IPX network address policies **and** IPX routing is enabled on the Group’s default VLAN #1. An exception condition arises in this situation because all ports in a Group are always members of that Group’s default VLAN #1 in addition to any other VLANs of which they are members. As a result, default VLAN #1 is always active.

The figure on the facing page illustrates this problem situation. In this figure, three VLANs within Group 2 – one of which is default VLAN #1 – have IPX routing enabled, as indicated by the presence of the internal IPX router. VLANs 2 and 3 both have IPX network address policies. When IPX Server A is connected to the Omni Switch/Router on port 1, the Server is assigned to VLAN 2 (per the network address policy) and port 1 becomes a member of VLAN 2. When IPX Server B is connected to the Omni Switch/Router on port 2, the Server is assigned to VLAN 3 (per the network address policy) and port 2 becomes a member of VLAN 3. However, ports 1 and 2 are also members of the Group’s default VLAN #1, so port 1 is now a member of VLAN 1 and VLAN 2 and port 2 is now a member of VLAN 1 and VLAN 3.

When IPX Server A sends broadcasts, they are restricted to VLAN 2 because of the network address policies. When IPX Server B sends broadcasts, they are restricted to VLAN 3, also because of the network address policies. However, when the internal IPX router sends out broadcasts on VLAN 1 the broadcasts are flooded out all ports in the Group, because all ports in the Group are, by default, members of VLAN 1. IPX Server A responds to this with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 2. IPX Server B also responds with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 3.

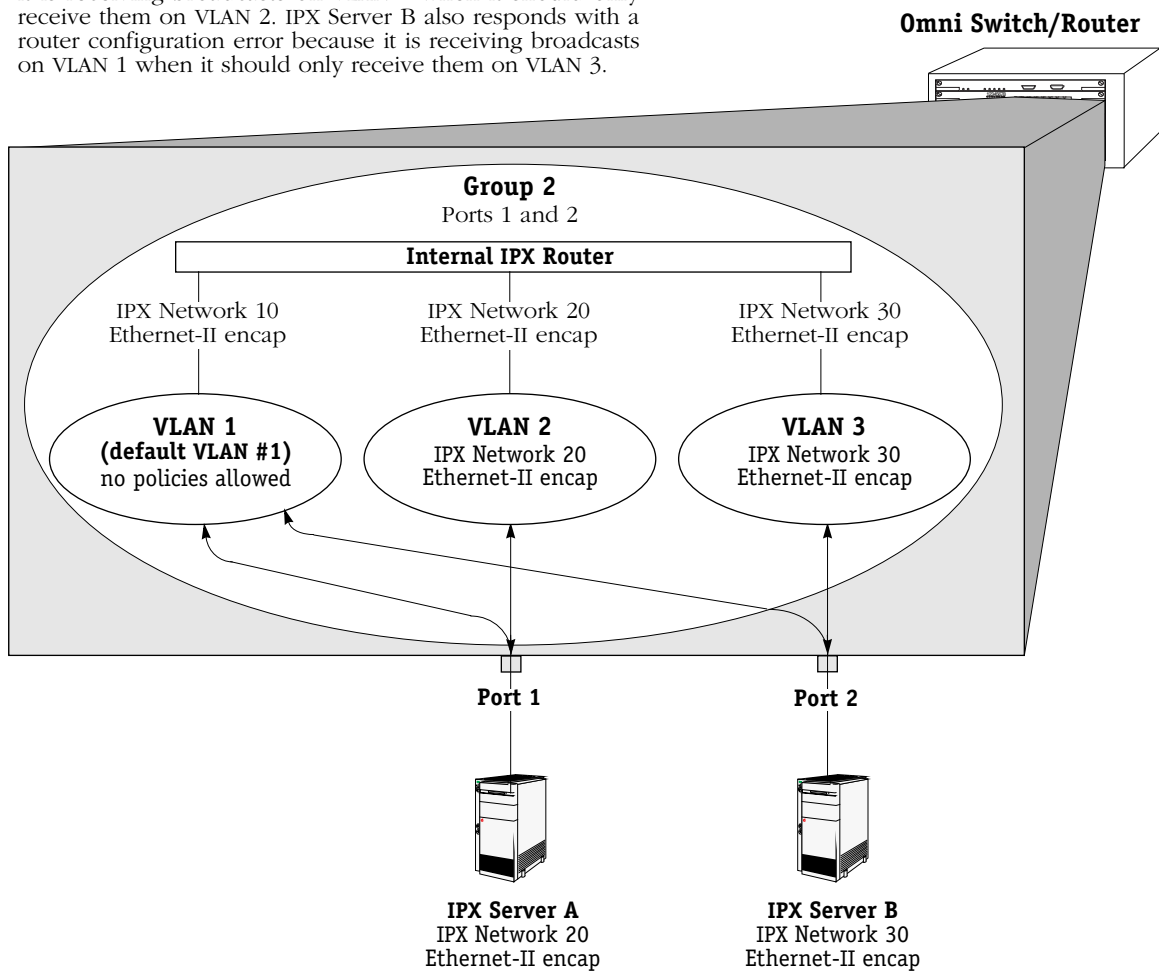
The Solution

The solution for this problem is to disable IPX routing on default VLAN #1. Because of this, when your network includes IPX servers that are members of IPX network address VLANs and IPX routing is enabled, you should configure your network such that disabling IPX routing on default VLAN #1 is not a problem.

Important Note

If you enable routing for a Group, you are actually enabling routing for that Group's default VLAN #1. For this reason, do not enable routing for any Group in which an IPX server is a member of an IPX network address VLAN.

When the internal IPX router sends out broadcasts on VLAN 1, they are flooded out all ports in the Group because, by default, all ports in the Group are members of VLAN 1. IPX Server A responds with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 2. IPX Server B also responds with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 3.



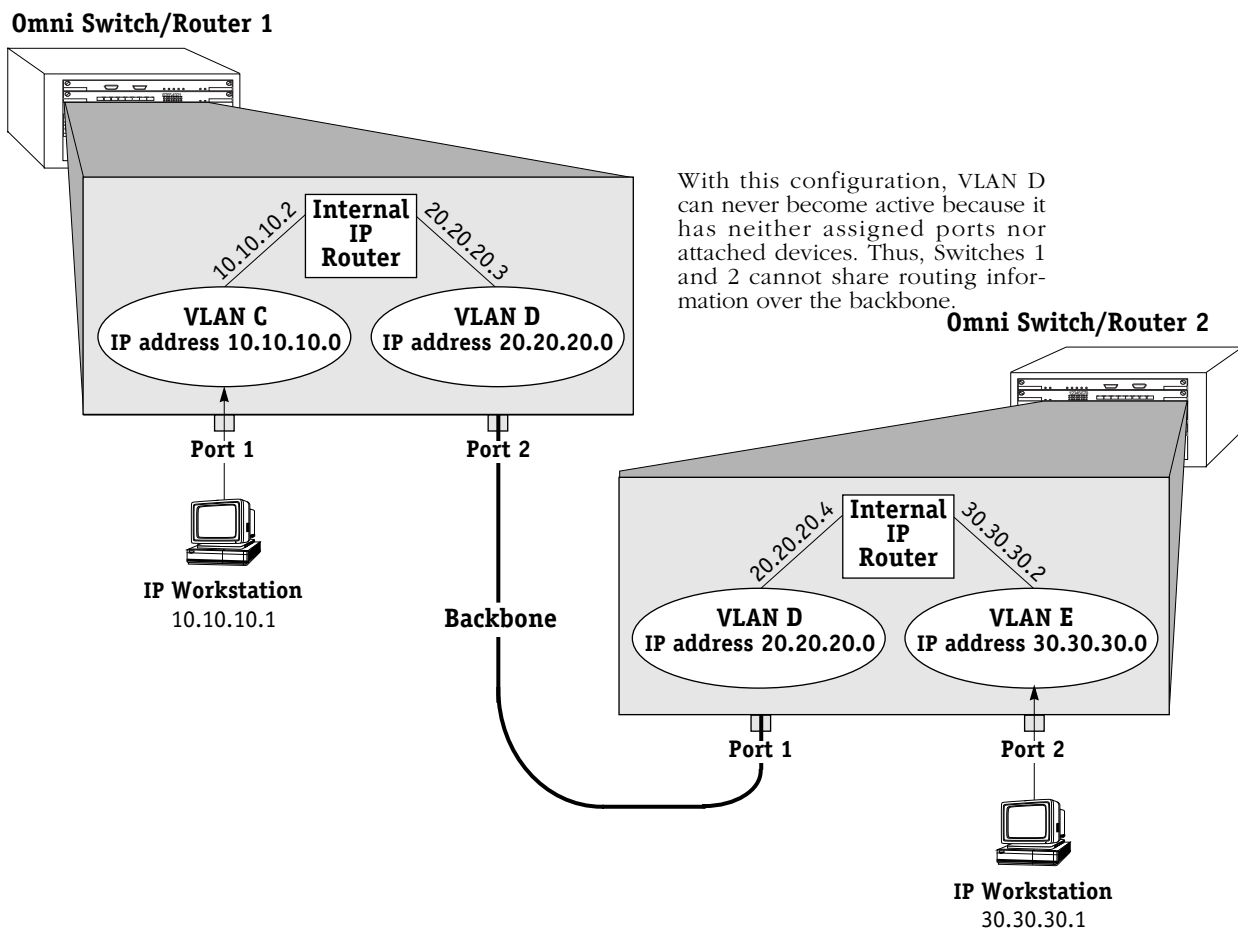
Application Example 5

Traversing a Backbone

Application Example 5 illustrates why port-based policies may be required to establish communications in some network situations, such as traversing a backbone. This necessity arises because, as explained in *How Routing Works Generally* on page 24-8, AutoTracker does not activate a VLAN – or its internal router interface – until a port is assigned to that VLAN. AutoTracker assigns ports to VLANs with port policies immediately. However, AutoTracker only assigns ports to VLANs with logical policies when a frame is received from a source device that matches the VLAN’s policies. This means that, in some network situations, you may need to assign a port policy to a VLAN to force it active.


The figure below illustrates the problem that can occur. The network below contains two Omni Switch/Routers in which three IP network address VLANs exist: VLAN C (IP address 10.10.10.0), VLAN D (IP address 20.20.20.0), and VLAN E (IP address 30.30.30.0). VLAN D spans both Omni Switch/Routers, but has no assigned devices. Routing is enabled for all three VLANs. A backbone connects port 2 on Omni Switch/Router 1 to port 1 on Omni Switch/Router 2.

When IP workstation 10.10.10.1 transmits a frame VLAN C and its internal router activate. When IP workstation 30.30.30.1 transmits a frame VLAN E and its internal router activate. All subsequent traffic on VLAN C is transmitted to IP workstation 10.10.10.1 and all subsequent traffic on VLAN E is transmitted to IP workstation 30.30.30.1. VLAN D cannot activate because there are no devices that match its network address policy and it has no ports assigned. Because VLAN D is not active, Switches 1 and 2 cannot exchange routing information. Switch 1 will not be aware of network 30 and Switch 2 will not be aware of network 10.



The Solution

The recommended solution is to add a port policy to VLAN D, as is shown in the figure below. A port policy can be defined in addition to any other policies defined for a VLAN. If VLAN D has a port policy that includes port 2 on Switch 1 and port 1 on Switch 2 – the ports to which the backbone is connected – VLAN D and its internal router will activate immediately in both Switch 1 and Switch 2. Traffic (i.e., routing information) can then flow between Switch 1 and Switch 2 over the backbone. Switch 1 will be aware of network 30 and Switch 2 will be aware of network 10.



Please Take Note

Refer to Chapter 20, “Configuring Group and VLAN Policies,” for information on original and current port policy functionality.

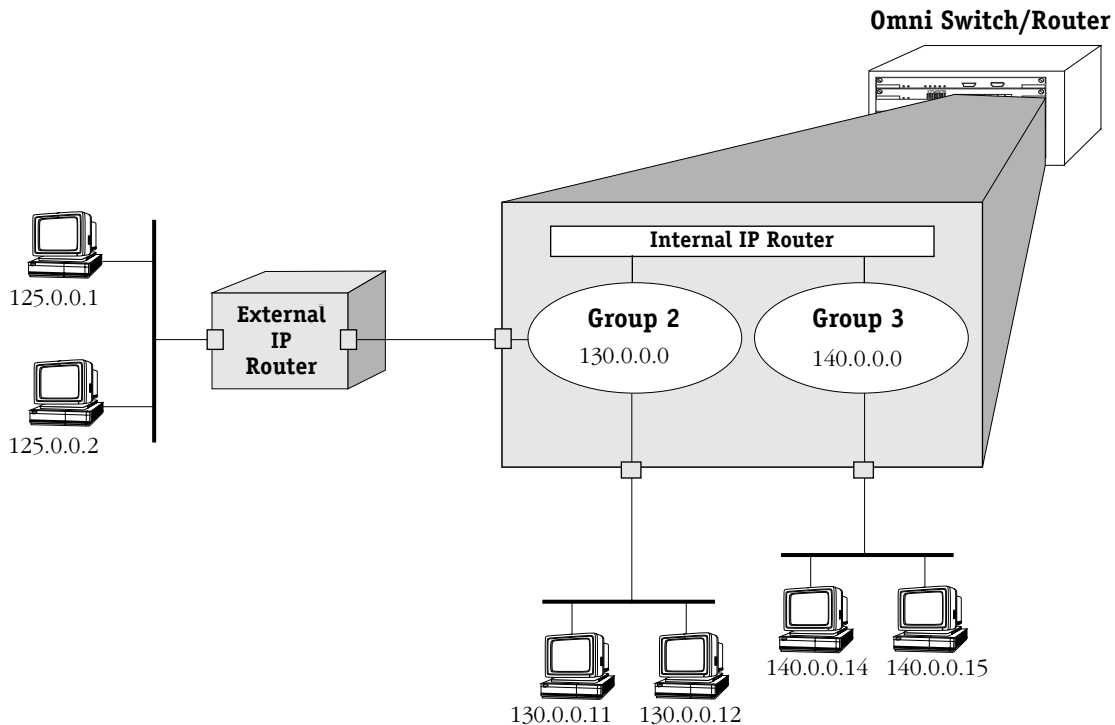
25 IP Routing

Introduction

This chapter gives an overview of IP routing and includes information about configuring static routes and viewing/configuring TCP/IP protocols such as Telnet and the Routing Information Protocol (RIP). IP routing requires at least one virtual router port to be configured on the switch. For information about configuring virtual router ports, see Chapter 19, “Managing Groups and Ports.”

When IP routing is enabled on the switch, the switch exchanges routing information with external IP routers in the network, and stations connected to groups and VLANs with virtual router ports can communicate. Groups or VLANs that do not have router ports with routing enabled are essentially firewalled from each other.

In the example shown here, stations connected to each group can communicate if a virtual router port is created for each group and each router port on the switch has IP routing enabled. Stations in group 2 and group 3 communicate with stations attached to the external IP router if a default route to that router is configured on the switch or the switch learns about the external router through RIP or some other routing protocol.



IP Routing Overview

In switching, traffic may be transmitted from one media type to another within the same broadcast domain (or group/VLAN). Switching happens at layer 2, the physical layer; routing happens at layer 3, the network layer. In routing, traffic may be transmitted across groups/VLANs, and broadcast or multicast traffic is prevented from being transmitted across those domains (unless some other mechanism is set up on the switch, such as UDP forwarding or IP multicast routing).

In IP routing, the switch builds routing tables to keep track of optimal destinations for traffic it receives that is destined for remote networks. The switch also sends and receives routing messages, or advertisements, to/from other routers in the network. When the switch receives a packet to be routed, it strips off the MAC header and examines the IP header of the packet. It looks up the source/destination address in the routing table, and then adds the appropriate MAC address to the packet.

Calculating routing tables and stripping/adding MAC headers to packets is performed by switch software unless a Hardware Routing Engine (HRE) or HRE-X is installed. The HRE or HRE-X significantly improves routing performance. See Chapter 1, “Omni Switch/Router Chassis and Power Supplies,” for information about the HRE-X. On the Omni S/R, IP routing has a fastpath mechanism with some additional statistics display on the IP Statistics and Errors screen available through the **ips** command (see *Viewing IP Statistics and Errors* on page 25-12).

IP is associated with several layer 3 and layer 4 protocols. Some of these protocols are built into the base code loaded into the switch. Others are included as part of Advanced Routing software. Some protocols are specifically used for routing; others are used by any host or end station that has an IP address. A brief overview of supported IP protocols is included here.

Routing Protocols

When IP routing is enabled, the switch uses routing protocols to build routing tables that keep track of stations in the network and to decide the best path for forwarding data. These routing protocols include:

- Routing Information Protocol (RIP)—An interior gateway protocol that defines how routers exchange information in an autonomous system. RIP makes routing decisions using a “least-cost path” method. RIP services are performed by a program operating in the switch called RouteD. RIP and RIP II services are also available from a program called GateD, which is part of Alcatel’s optional Advanced Routing software. RIP, whether performed by RouteD or GateD, allows the switch to learn routing information from other, neighboring RIP routers.
- Open Shortest Path First (OSPF)—An interior gateway protocol that provides a routing function similar to RIP but which uses different techniques to determine the best route for a datagram. OSPF services are provided by GateD, part of Alcatel’s optional Advanced Routing software.
- Border Gateway Protocol (BGP)—An exterior gateway protocol that provides for routing between autonomous systems. BGP is not part of the base code but is included in the Advanced Routing software.

Transport Protocols

IP is both connectionless (it routes each datagram separately) and unreliable (it does not guarantee delivery of datagrams). This means that a datagram may be damaged in transit, or thrown away by a busy router, or simply never make it to its destination. The resolution of these transit problems is to use a layer 4 transport protocol:

- Transmission Control Protocol (TCP)—A major data transport mechanism that provides reliable, connection-oriented, full-duplex data streams. While the role of TCP is to add reliability to IP, TCP relies upon IP to do the actual delivering of datagrams.
- User Datagram Protocol (UDP)—A secondary transport-layer protocol that uses IP for delivery. However, UDP is not connection-oriented so it does not provide reliable end-to-end delivery of datagrams. But some applications can safely use UDP to send datagrams that don't require the extra overhead added by TCP.

Application-Layer Protocols

- Bootstrap Protocol (BOOTP)/Dynamic Host Configuration Protocol (DHCP)—May be used by an end station to obtain an IP address. The switch provides a UDP relay that allows BOOTP requests/replies to cross different networks. See Chapter 26, “UDP Forwarding.”
- Simple Network Management Protocol (SNMP)—Used to manage nodes on a network. SNMP is discussed in Chapter 13, “Configuring SNMP.”
- Telnet—Used for remote connection to a device. The **telnet** command is described in this chapter.
- File Transfer Protocol (FTP)—Enables transferring files between hosts.

Additional IP Protocols

- Internet Control Message Protocol (ICMP)—Specifies the generation of error messages, test packets, and informational messages related to IP. ICMP supports the **ping** command used to determine if hosts are online.
- Address Resolution Protocol (ARP)—Used to find the IP address that corresponds to a given physical (MAC) address.
- Internet Group Management Protocol (IGMP)—Tracks multicast group membership. See the Multicast Services section of the *Advanced Routing User Manual*.
- Resource ReSerVation Protocol (RSVP)—Signals Quality of Service (QoS) requests in an IP network. For more information, see the *Switched Network Services User Manual*.

Setting Up IP Routing on the Switch

IP routing is enabled on a per-port basis by creating a virtual IP router port for a group/VLAN. The switch does not do any routing unless the virtual router port has IP routing enabled (routing is enabled by default). The steps for setting up IP routing on the switch are given here:

Step 1. Configuring a Virtual Router Port

A virtual router port may be created when you set up or modify a group/VLAN through the **crpg** command or **modvl** command described in Chapter 19, “Managing Groups and Virtual Ports.” To create a virtual router port, enable IP routing and specify an IP address for the router port.

When routing is enabled on the port, the switch creates routing tables and address translation tables so it knows how to forward traffic. The switch keeps track of router ports and any other routers in the network. The switch uses the Address Resolution Protocol (ARP) to match IP addresses with MAC addresses. It uses routing protocols, such as the Routing Information Protocol (RIP), to determine the best path for forwarding traffic. (Other routing protocols are available in the Advanced Routing software package.) It also periodically sends/receives routing messages to/from other routers to keep its routing tables updated.

◆ Important Note ◆

When Spanning Tree and IP routing are both enabled, packets are not forwarded unless the Spanning Tree Status for the port to which packets are to be forwarded has progressed from Listening to Learning to Forwarding. For example, if IP is enabled on VLAN 42 that has ports 1/1-3 attached to it and you want to forward to a host from port 1/2. Use the **vi 1/2** command to determine if the Spanning Tree Protocol has entered the Forwarding state for that port.

Step 2. Configuring Optional IP Routing Parameters

Optional configuration for IP routing includes the following:

- Static routes. These are routes that are manually added to the routing table and may be used rather than dynamic routes (which are learned through routing protocols like RIP).
- RIP filters. Controls the operation of RIP by minimizing the number of entries that will be added to the routing table.

Static routes and RIP filters are described in this chapter. This chapter also describes how to view various IP statistics as well as the routing table. It includes information about how to ping another IP host in the network, how to telnet to a remote system, and how to trace an IP route.

Step 3. Configuring Other IP Routing Features

There are several optional features that may be used with IP routing. Some features are included as part of the base code and are described in this user manual. Other features are available as optional switch software and are described in separate user manuals. The features are listed here:

- UDP forwarding—Forwards UDP broadcasts/multicasts across groups/VLANs. See Chapter 26, “UDP Forwarding.”
- GateD—Provides gateway protocols, including RIP, OSPF, and BGP/CIDR. See the *Advanced Routing User Manual*.
- Virtual Router Redundancy Protocol (VRRP)—Used to back up static IP routes. See the *Advanced Routing User Manual*.
- IP Firewall—Enables the switch to act as a gateway to provide security for all data entering and exiting the switch to and from its attached physical ports, as well as internally between groups and VLANs that are defined in the switch. See the *Switched Network Services User Manual*.
- Multicast services—Includes IP multicast switching (IPMS) and IP multicast routing (MrouteD). See the *Advanced Routing User Manual*.
- IP Control—Manages IP addresses through Lightweight Directory Access Protocol (LDAP), DHCP, and Domain Name Service (DNS). See the *Switched Network Services User Manual*.

The Networking Menu

The Networking menu contains commands that control, and are related to, the routing protocols that are run on the switch.

To switch to, and to display, the **Networking** menu, enter the following commands:

```
networking
?
```

If you have enabled the verbose mode, you do not need to enter the question mark (?).

A screen similar to the following displays:

Command	Networking Menu
-----	-----
snmps	View SNMP statistics
snmpc	Configure SNMP
Names	Configure the DNS resolver
probes	Display all RMON probes
events	Display all logged RMON events
IP	Enter IP networking command sub-menu.
IPX	Enter IPX networking command sub-menu
Gated	Enter Gated menu/control Gated
IPMR	Enter the IPMR routing sub-menu
IPMS	Enter the IPMS networking command sub-menu
VRRP	Enter the VRRP menu
QoS	Enter the QoS menu
Policy	Administer the SNS policy sub-menu
LDAP	Configure the SNS LDAP server sub-menu
Monitor	Enter port monitor utility command sub-menu
chngmac	Configure router port's MAC address on selected Group
RD	Routing Domain Management Menu

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The commands in this menu are described throughout this manual as follows:

- The **snmps** and **snmpc** commands are described in Chapter 13, “Configuring SNMP.”
- The **Names**, **probes**, **events**, and **chngmac** commands are described in Chapter 14, “RMON and DNS Resolver.”
- The IP submenu is discussed in this chapter. The IPX submenu is described in Chapter 27, “IPX Routing.”
- The Gated, IPMR, IPMS, VRRP, and RD submenus are available if Advanced Routing software is loaded on the switch. See the *Advanced Routing User Manual* for more information.
- The QoS, Policy, and LDAP submenus are available if Switched Network Services software is loaded on the switch. See the *Switched Network Services User Manual* for more information.
- The Monitor submenu is described in Chapter 19, “Managing Groups and Ports.”

The IP Submenu

The **ip** command in the Networking menu is used to display the IP submenu. To display the IP submenu, enter the following commands:

```
ip
```

```
?
```

If you have enabled the verbose mode, you don't need to enter the question mark (?).

A screen similar to the following displays:

Command	IP Menu			
xlat	View the address translation table			
ips	View IP stats & errors			
ipr	View IP routes			
aisr	Add an IP static route			
risr	Remove an IP static route			
icmps	View ICMP stats & errors			
ping	Ping a system			
udps	View UDP stats and errors			
udpl	View the UDP listener table			
rips	View RIP stats and errors			
tcps	View TCP-related statistics			
tcpc	View the TCP Connection table			
telnet	Remote login to another system using TELNET			
traceroute	Trace an IP route			
relay	Use 'relayc' or 'relays'			
fwconfig	Configure the IP Firewall			
ripflush	Flush all routes obtained by RIP			
ipfilter	Add/delete an IP RIP filter			
ipf	Display IP RIP filters			
ipmac	View the IP to MAC Address Association table			
ipclass	Turn on/off IP Class Address Checking			
ipdirbrcast	Turn on/off IP directed broadcast			
Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

This chapter describes all of the above commands with the exception of **fwconfig**, **relayc**, **relays**, and **ipclass** commands. The **fwconfig** command is described in the *Switched Network Services User Manual*. The relay commands, **relayc** and **relays**, are described in Chapter 26, "UDP Forwarding." The **ipclass** command is described in the *Advanced Routing User Manual*.

Viewing the Address Translation (ARP) Table

The **xlat** command is used to access the ARP (Address Resolution Protocol) Table. This table contains a listing of IP addresses and their corresponding translations to MAC addresses (or slot/port for WAN interfaces). Submenu commands are used to add entries to the table, to delete them, show all the entries currently in the table, to flush “temporary” entries, to display specific entries by either MAC or IP address, and to quit out of the **xlat** submenu.

To begin working with the ARP Table, enter the following command:

xlat

A screen similar to the following displays:

ARP Table Functions

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

The default command is **show** which is used to display all entries in the table. The **quit** command is used to exit out of this submenu and return to the main system prompt.

Displaying All Entries in the ARP Table

At the above prompt, press **<Enter>** to select **Show**, the default command.

A screen similar to the following displays:

Address Translation Table

IP Address	at	Physical Address
90.0.0.1	at	3/1, dlci=32
198.206.184.34	at	00:05:02:c0:7f:11
198.206.184.254	at	00:20:da:6a:98:40

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

The fields on this screen have the following meanings:

IP Address

The IP address, in dotted-decimal format, of a specific host or other device.

Physical Address

The MAC address, in hexadecimal format, of the specific host or other device that corresponds to the IP address in the left-hand column.

Adding Entries to the ARP Table

The **add** subcommand is used to manually add an IP address entry to the ARP Table. To be able to manage your switch over an IP network connection, you will need at least one IP address configured for the switch.

Follow the steps below to add an address to the ARP Table.

1. Enter **add**.

The following prompt displays:

Host name or IP addr to add:

Enter the name of the host or its IP address.

2. The following prompt displays:

Physical address (format aa:bb:cc:dd:ee:ff):

Enter the host's physical address in hexadecimal format.

3. The following prompt displays:

Publish (i.e., proxy for) this entry? (y/n) (n):

Enter **y** to publish (i.e., proxy for) this ARP entry. This feature allows the switch to answer all ARP requests directed at the hosts on a subnetwork. As the "proxy" for these hosts, the switch responds with its own MAC address whenever ARP requests come in for any of the hosts on the subnetwork. Enter **n** if you do not want this ARP entry to act as a proxy.

4. The following prompt displays:

Is this entry permanent (ie. flush will not remove it) (y/n)? (n) :

Enter **y** if this entry is to be permanent (that is, you do not want it to be removed by the **Flush** subcommand). Enter **n** if the entry is to be temporary (that is, you want to allow it to be removed by the **Flush** subcommand). All of the entries in the table, whether they are permanent or temporary, survive across switch reboots. Therefore, you must use the **Delete** subcommand when you want to remove permanent entries from the table.

5. The following prompt displays:

Use trailer encapsulation on this host (y/n)? (n) :

Enter **y** if you want to use trailer encapsulation on this host. Enter **n** if you do not want to use trailer encapsulation on this host.

6. The system then confirms the addition to the table (an example is shown below).

ARP table entry for host 198.206.184.35 successfully added

7. The **xlat** submenu redisplay:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Deleting Entries from the ARP Table

The **Delete** subcommand is used to delete a “permanent” IP address from the ARP Table. Follow the steps below to delete an address from the ARP Table.

1. Enter **delete**.

The following prompt displays:

Host name or IP addr to delete:

Enter the host name or address that you wish to delete.

2. The system will then confirm the deletion from the table (an example is shown below).

ARP table entry for host 198.206.184.35 successfully deleted

3. The **xlat** submenu redisplay:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Flushing Temporary Entries from the ARP Table

The **Flush** subcommand is used to delete “temporary” IP addresses from the ARP Table. Follow the steps below to flush all temporary addresses from the ARP Table.

1. Enter **flush**.

The following prompt displays:

Flushing all non-permanent ARP table entries...done

2. The **xlat** submenu redisplay:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Finding a Specific IP Address in the ARP Table

The **Macfind** subcommand is used to locate a specific IP address in the ARP Table *based on a known MAC address*. (The **Ipfind** subcommand, discussed next, is used to find a specific MAC address based on a known IP address).

Follow the steps below to display a specific IP address in the ARP Table.

1. Enter **macfind**.

The following prompt displays:

MAC address to find (format aa:bb:cc:dd:ee:ff):

2. Enter the known MAC address (for example, 00:05:02:c0:7f:11).

A prompt similar to the following displays which shows the IP address that is related to the MAC address you entered:

Corresponding IP address: 198.206.184.34

3. The **xlat** submenu will then be redisplayed:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Finding a Specific MAC Address in the ARP Table

The **ipfind** subcommand is used to locate a specific MAC address in the ARP Table *based on a known IP address or host name*. (The **Macfind** subcommand, discussed above, is used to find a specific IP address based on a known MAC address).

Follow the steps below to display a specific MAC address in the ARP Table.

1. Enter **ipfind**.

The following prompt displays:

Hostname or IP address to find:

2. Enter the known IP address or host name (for example, 198.206.184.34).

A prompt similar to the following displays which shows the MAC address that is related to the IP address entered:

Corresponding MAC address: 00:05:02:c0:7f:11

3. The **xlat** submenu redisplay:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Viewing IP Statistics and Errors

The **ips** command is used to monitor IP datagram traffic and errors. The **ips** command displays *cumulative* IP statistics and errors. The statistics show the cumulative totals since the last time the switch was powered on or since the last reset of the switch was executed.

To display information about IP statistics and errors, enter the following command:

```
ips
```

The Omni Switch/Router (OmniS/R) includes fastpath code that enhances the speed of IP routing. Fastpath statistics are included on the IP Statistics and Errors screen:

IP Statistics and Errors	
Default Time to Live	32
Reassembly Timeout (seconds)	1
Total Datagrams Recvd/Forwarded	513342 / 513283
Fastpath Datagrams Received	513281
Fastpath Datagrams Forwarded	513280
Fastpath Inbound Discards	1
Fastpath Utilization	100%
PDU's Requested for Transmit	4294931545
PDU's Needing Reassembly	0
PDU's Successfully Reassembled	0
PDU's Needing Fragmentation	0
Fragments created	0
IP Errors (Discards due to the following problems)	
Header errors	0
Address errors	45994
Unknown/Unsupported Protocol	0
Local discards inbound/outbound	0 / 0
Unknown Route	45994
Reassembly Failures	0
Fragmentation Failures	0

The fields on this screen have the following meanings:

Default Time to Live

The default time, in seconds, assigned to each outgoing IP datagram before it is discarded as expired.

Reassembly Timeout (seconds)

The time, in seconds, to wait for all fragments to arrive before discarding datagrams.

Total Datagrams Recvd/Forwarded

The total number of input IP datagrams received, including those received in error.

HRE Datagrams Forwarded

The total number of IP datagrams forwarded by the HRE (Hardware Routing Engine).

Fastpath Datagrams Received

(Displays for Omni S/R.) The number of IP datagrams received by the fastpath code.

Fastpath Datagrams Forwarded

(Displays for Omni S/R) The number of IP datagrams forwarded to their destination without using the MPX.

Fastpath Inbound Discards

(Displays for Omni S/R) The number of bad packets received and discarded. Typically this value should be zero.

Fastpath Utilization

(Displays for Omni S/R) The percentage of total datagrams received that are forwarded by the fastpath code.

PDU Requested for Transmit

The total number of IP datagrams which transmit local IP user-protocols (including ICMP) supplied to IP in requests for transmission, not including forwarded datagrams.

PDU's Needing Reassembly

The number of IP datagram fragments that needed to be reassembled by this switch.

PDU's Successfully Reassembled

The number of IP datagrams successfully reassembled by this switch.

PDU's Needing Fragmentation

The number of IP datagrams requiring fragmentation by this switch.

Fragments created

The number of IP datagram fragments that have been generated as a result of fragmentation by this switch.

Header errors

The number of input IP datagrams discarded due to errors in their IP header, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discarded in processing their IP options, etc.

Address errors

The number of input IP datagrams discarded because the IP header destination field contained an invalid address.

Unknown/Unsupported Protocol

The number of local addresses, unsupported protocols, datagrams received successfully but discarded because of an unknown or unsupported protocol.

Local discards inbound/outbound

The number of packets discarded, both inbound and outbound, though they had no errors to prevent their being transmitted (lack of buffer space, etc.).

Unknown Route

The number of packets received and discarded by IP because IP was unable to route them.

Reassembly Failures

The number of failures detected by the IP reassembly algorithm for all reasons (timed out, error, etc.) This value is not necessarily a count of the discarded fragments.

Fragmentation Failures

The number of IP datagrams discarded because they needed to be fragmented but could not be. This situation could happen when a large packet has the "Don't Fragment" flag set.

Viewing the IP Forwarding Table

The `ipr` command is used to display the IP Forwarding Table. The entries in the table show the routes entered by a routing protocol, if the switch is running any of the supported protocols, and the static routes that you may have entered manually. You can also add to, or remove static routes from, the IP Forwarding Table (see *Adding an IP Static Route* on page 25-17 and *Removing an IP Static Route* on page 25-19).

To display the IP Forwarding Table, enter the following command:

```
ipr
```

A screen similar to the following displays:

10 routes in forwarding table

IP FORWARDING TABLE

Network	Mask	Gateway	Metric	Group VLAN Id:	Protocol
0.0.0.0	255.0.0.0	198.206.184.254	1	1:1	STATIC
10.0.0.0	255.0.0.0	10.0.0.1	1	6:1	STATIC
11.0.0.0	255.0.0.0	11.0.0.1	1	5:1	STATIC
90.0.0.0	255.0.0.0	90.0.0.3	1	4:1	DIRECT
127.0.0.0	255.0.0.0	127.0.0.1	0	1:2	LOOPBACK
127.0.0.1	255.255.255.255	127.0.0.1	0	2:3	LOOPBACK
196.196.7.0	255.255.255.0	196.196.7.42	1	3:1	RIP
198.206.187.0	255.255.255.0	198.206.183.0	1	1:4	STATIC
198.206.184.0	255.255.255.0	198.206.184.42	1	1:1	RIP
203.229.229.0	255.255.255.0	203.229.229.250	1	2:1	RIP

If routing domains are configured on the switch, the `ipr` command will display the forwarding table for the default routing domain only. Routing domains are part of Advanced Routing software and are not part of the base code. For more information about routing domains, see Chapter 14, "Routing Domains," in the *Advanced Routing User Manual*.

To display the forwarding table for a routing domain other than the default domain, enter the `ipr` command with the relevant routing domain ID. For example:

```
ipr 2
```

The screen display is similar to the following:

4 routes in forwarding table

IP FORWARDING TABLE for Routing Domain 2

Network	Mask	Gateway	Metric	Group VLAN Id:	Protocol
0.0.0.0	255.0.0.0	198.206.184.254	1	1:1	STATIC
10.0.0.0	255.0.0.0	10.0.0.1	1	6:1	STATIC
11.0.0.0	255.0.0.0	11.0.0.1	1	5:1	STATIC
90.0.0.0	255.0.0.0	90.0.0.3	1	4:1	DIRECT

Viewing the IP Forwarding Table

The fields on the IP Forwarding Table have the following meanings:

Network

The destination network IP address.

Mask

The IP subnet mask.

Gateway

The network address of the gateway (the router from which this address was learned).

Metric

The metric associated with this network. Generally, this is a RIP “hop” count, or the number of hops the network is away from this router.

Group VLAN Id

The group and VLAN number from which this IP address was learned.

Protocol

The way in which this route was learned, for example, through RIP.

Adding an IP Static Route

The **aisr** command is used to add IP static routes to the switch's IP Forwarding Table. You might want to add a static route to send traffic to a router other than the one determined by the routing protocols.

In order to add a static route, you will need to know the host/net IP address and the gateway IP address which will be used to route traffic to the external IP address. If routing domains are configured on the switch and you want to add the route to a particular domain other than the default, you will need to know the relevant routing domain ID (RDID). For more information about routing domains, see Chapter 14, "Routing Domains," in the *Advanced Routing User Manual*.

Follow the steps below to add an IP static route.

1. Enter **aisr**. The prompt that displays depends on whether routing domains are configured on the switch.

If routing domains *are* configured on this switch, the following prompt displays:

**Routing Domains (RD) are configured on this switch.
List the RD(s) you want this route applied to? (default: none) :**

If you do not want to apply the new route to a particular routing domain, press **Enter**. To apply the route you are adding to an existing routing domain, enter the desired routing domain ID (RDID) and go to step 3.

If routing domains *are not* configured on this switch or if you are applying this route to the default domain, the following prompt displays:

Do you want to see the current route table? (y or n) (y) :

2. Enter a **y** at this prompt (or press **Enter**) to display the current forwarding table.

A screen similar to the following displays:

IP FORWARDING TABLE

Network	Mask	Gateway	Metric	Group VLAN Id:	Protocol
0.0.0.0	255.0.0.0	198.206.184.254	1	1:1	STATIC
10.0.0.0	255.0.0.0	10.0.0.1	1	6:1	STATIC
11.0.0.0	255.0.0.0	11.0.0.1	1	5:1	STATIC
90.0.0.0	255.0.0.0	90.0.0.3	1	4:1	DIRECT
127.0.0.0	255.0.0.0	127.0.0.1	0	1:2	LOOPBACK
127.0.0.1	255.255.255.255	127.0.0.1	0	2:3	LOOPBACK
196.196.7.0	255.255.255.0	196.196.7.42	1	3:1	RIP
198.206.187.0	255.255.255.0	198.206.183.0	1	1:4	STATIC
198.206.184.0	255.255.255.0	198.206.184.42	1	1:1	RIP
203.229.229.0	255.255.255.0	203.229.229.250	1	2:1	RIP

Destination IP address of host or network :

3. At the prompt for the destination IP address, enter the address of the host or network to which you are setting up a route. For a "default" route, use an entry of 0.0.0.0 as the IP address (or just enter the word **default**).
4. If you entered an IP address, a prompt similar to the following displays:

Host or network mask (255.255.255.000) :

Enter the mask (or just press **<Enter>** to accept the default mask).

5. The following prompt displays:

IP address of next hop :

Enter the IP address of the next hop (the gateway) router to the destination IP address. The gateway address must be on the same network as one of the VLANs (that is, it must be a directly connected network).

A message will confirm the creation of the static route:

Route successfully added

Removing an IP Static Route

The **risr** command is used to remove IP static routes from the switch's IP Forwarding Table.

Follow the steps below to remove an IP static route.

1. Enter **risr**. The prompt that displays depends on whether routing domains are configured on the switch. For more information about routing domains, see Chapter 14, "Routing Domains," in the *Advanced Routing User Manual*.

If routing domains *are* configured on this switch, the following prompt displays:

**Routing Domains (RD) are configured on this switch.
List the RD(s) you want this route applied to? (default: none) :**

If you are removing a route from an existing domain, press **Enter**. To remove a route from an existing routing domain, enter the desired routing domain ID (RDID) and go to step 3.

If routing domains are not configured on this switch, or if you are applying this route to the default domain, the following prompt displays:

Do you want to see the current route table? (y or n) (y) :

2. Enter a **y** at this prompt (or just press **<Enter>**) to display the current forwarding table.

A screen similar to the following displays:

IP FORWARDING TABLE

Network	Mask	Gateway	Metric	Group VLAN Id:	Protocol
0.0.0.0	255.0.0.0	198.206.184.254	1	1:1	STATIC
10.0.0.0	255.0.0.0	10.0.0.1	1	6:1	STATIC
11.0.0.0	255.0.0.0	11.0.0.1	1	5:1	STATIC
90.0.0.0	255.0.0.0	90.0.0.3	1	4:1	STATIC
127.0.0.0	255.0.0.0	127.0.0.1	0	1:2	LOOPBACK
127.0.0.1	255.255.255.255	127.0.0.1	0	2:3	LOOPBACK
196.196.7.0	255.255.255.0	196.196.7.42	1	3:1	RIP
198.206.187.0	255.255.255.0	198.206.183.0	1	1:4	STATIC
198.206.184.0	255.255.255.0	198.206.184.42	1	1:1	RIP
203.229.229.0	255.255.255.0	203.229.229.250	1	2:1	RIP

Destination IP address of host or network :

3. At the prompt for the destination IP address, enter the IP address of the host or network that you want to remove.

4. A prompt similar to the following displays:

Host or network mask (255.255.255.000) :

Enter the mask (or just press **<Enter>** to accept the default mask).

5. The following prompt displays:

IP address of next hop :

Enter the IP address of the next hop (the gateway) router to the destination IP address.

A message will confirm the deletion of the static route:

Route successfully deleted

Viewing ICMP Statistics and Errors

The `icmps` command is used to monitor ICMP activity.

To display information about ICMP statistics and errors, enter the following command:

```
icmps
```

A screen similar to the following displays:

ICMP Statistics		
	In	Out
Total ICMP Messages	1	1
Redirect Messages	0	0
Echo Messages	1	0
Echo Reply Messages	0	1
Time Stamp Messages	0	0
Time Stamp Reply Messages	0	0
Address Mask Messages	0	0
Address Mask Reply Messages	0	0

ICMP Errors		
	In	Out
Errors	0	0
Destination Unreachable Msgs	0	0
Time Exceeded Msgs	0	0
Parameter Problems	0	0
Source Quenches	0	0

The following field descriptions pertain to both the “in” and “out” statistics:

Total ICMP Messages

The total number of ICMP messages which this switch received or attempted to send out.

Redirect Messages

The number of ICMP Redirect messages sent/received by this switch.

Echo Messages

The number of ICMP Echo messages sent/received by this switch to see if a destination is active and reachable.

Echo Reply Messages

The number of ICMP Echo Reply messages received by this switch.

Time Stamp Messages

The number of Time Stamp Request messages sent/received by this switch requesting/receiving a reply with timestamp.

Time Stamp Reply Messages

The number of Time Stamp Reply messages sent/received by this switch.

Address Mask Messages

The number of Address Mask Reply messages that were sent/received by this switch in an attempt to determine the subnet mask for a network.

Address Mask Reply Messages

The number of Address Mask Reply messages that were sent/received by this switch.

Errors

The number of ICMP messages this switch sent/received but was unable to process because something was wrong (for example, a checksum failure).

Destination Unreachable Msgs

The number of ICMP “destination unreachable” messages that were sent/received. These occur when the gateway is unable to route a datagram to its destination.

Time Exceeded Msgs

The number of “time exceeded” messages that were sent/received. These occur when a packet is dropped because the Time-to-Live counter reaches zero. When a large number of these messages are encountered this is a symptom that packets are looping, that congestion is severe, or that the Time-to-Live counter is set too low. These messages also occur when all the fragments trying to be reassembled don't arrive before the reassembly timer expires.

Parameter Problems

The number of messages sent/received which indicate that an illegal value has been detected in a header field. These messages can indicate a problem in the sending host's IP software or possibly in the gateway's software.

Source Quenches

The number of messages sent/received which tell a host that is sending too many packets. A host should attempt to reduce its transmissions upon receiving these messages.

Using the PING Command

The **ping** command is used to test the reachability of IP network destinations. A fast ping command (**fping**) is also available for repeating the last ping request sent from the switch. The commands send an ICMP echo request to a destination and then wait for a reply.

Follow the steps below to issue an IP ping request.

1. Enter **ping**.

A screen similar to the following displays:

Host () :

Enter the IP address of the host that you want to “ping.”

2. The following prompt displays:

Count (0 for infinite) (0) :

Enter the number of frames to be transmitted (0 equals “infinite”). To abort an “infinite” transmission once it is in progress, just press **Enter** again.

3. The following prompt displays:

Size (64) :

Enter the desired size of the data portion of the packet. You can specify a packet size or a range of packet sizes up to 8148. If you give a range, the switch will increment the packet size by 1 each time up to the top of the range. It will then wrap and continue from the bottom size of the range again until the total number of frames specified in the count has been sent. You can also set the increment by which the packet size is increased each time by entering a comma and an increment number after the size. For example, an entry of

1-100,5

will send out the number of frames specified in the “Count” prompt, starting with a frame size of 1 and incrementing up to a frame size of 100 in steps of 5. Note that if the “Count” is too small, the 100-byte frame size may never be reached. If the count is large enough, the packet size will wrap and go back to 1.

4. The following prompt displays:

Timeout (1) :

Enter the number of seconds the program is to wait for a response before timing out.

Viewing UDP Statistics and Errors

The **udps** command is used to display a listing of UDP statistics and errors. The **udps** command displays cumulative statistics since the last time the switch was powered on or since the last reset of the switch was executed.

To display information about UDP statistics and errors, enter the following command:

```
udps
```

A screen similar to the following displays:

```
Total UDP datagrams received           :   831  
Total UDP datagrams transmitted      :    22  
Total Datagrams received w/unknown applications :    0  
Total UDP datagrams w/other Errors    :    0
```

The fields on this screen have the following meanings:

Total UDP datagrams received

The total number of UDP datagrams delivered to UDP applications.

Total UDP datagrams transmitted

The total number of UDP datagrams sent from this switch.

Total UDP datagrams received w/unknown applications

The total number of datagrams for which there was no application at the destination.

Total UDP datagrams w/other Errors

The total number of UDP datagrams that could not be delivered for reasons other than lack of application at the destination.

Viewing the UDP Listener Table

The **udpl** command is used to display the UDP Listener Table. This table contains information about the switch's UDP end-points on which a local application is currently accepting datagrams. The UDP Listener Table shows the local IP addresses for each UDP listener and the local port number for this listener. An IP address of zero (0.0.0.0) indicates that it is listening on all interfaces.

To view the UDP Listener Table, enter the following command:

```
udpl
```

A screen similar to the following appears:

UDP Listener Table			Recv-Q	Send-Q
Local Address/Port				
0.0.0.0	/	162	0	0
0.0.0.0	/	161	0	0
0.0.0.0	/	520	0	0
0.0.0.0	/	1024	0	0

Local Address/Port

The local IP address, and the local port number, for this UDP connection. In the case of a connection in the listen state, which is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used.

Recv-Q and Send-Q

For the SNMP Traps (port 162) this is the number transmitted (there is no receive).

For the SNMP Requests (port 161) this is the number of Request PDUs sent and the number of Response PDUs received.

For RIP (port 520) this is the number of packets received and transmitted.

Viewing RIP Statistics and Errors

The **rips** command is used to display RIP statistics and errors. This command displays cumulative statistics since the last time the switch was powered on, or since the last reset of the switch was executed.

To display information about RIP statistics and errors, enter the following command:

```
rips
```

A screen similar to the following displays:

```

                                RIP Statistics
Rtr (Group ID:VLAN ID 1:1) IP Address 198.206.182.115 RIP Mode silent
In          4769          Out          0
Transmit Error    0      Non-zero field    0
Bad Version      0      Bad Metric      0
Bad Family       0      Bad Size       0
Bad Address      0      Bad Command    0
```

The fields on this screen have the following meanings:

In/Out

The total number of RIP packets received and transmitted on a per-virtual-LAN basis.

Transmit Error

The total number of RIP packets that were unable to be sent.

Bad Version

The total number of RIP messages delivered to the switch that were not version 1.

Bad Family

The number of packets received on this VLAN whose family ID was not of the Internet family.

Bad Address

The number of received packets whose IP address was not a Class A, B, or C.

Non-zero Field

The number of received packets whose mandated “must-be-zero” fields were not zero.

Bad Metric

The number of received packets with a routing entry’s metric that was out of range.

Bad Size

The number of received packets that were not compatible with the expected size.

Bad Command

The number of received packets whose command field was not a “request” or “response.”

Viewing TCP Statistics

The **tcps** command is used to monitor TCP traffic activity and check TCP configuration parameters. To reconfigure TCP parameters, see *Viewing the TCP Connection Table* on page 25-29.

To display information about TCP activity, enter the following command:

```
tcps
```

A screen similar to the following displays:

```

TCP Statistics

Round Trip Algorithm Used      : RSRE (MIL-STD-1778)
Retransmission Min/Max Timeout : 300/3000
Max Connections Allowed       : Unlimited
Active Opens                   : 76
Passive Opens                  : 43
Attempt Fails                  : 0
Established Resets            : 5
Currently Established         : 3
Total Segments Received       : 1117
Total Segments Sent           : 832
Total Segments Retransmitted   : 0
Total Segments Received w/err : 0
Total Segments Sent w/RST flag : 0

```

The fields on this screen have the following meanings:

Round Trip Algorithm Used

The algorithm used to determine the Timeout value used for retransmitting unacknowledged octets. The value is: RSRE (MIL-STD-1778).

Retransmission Min/Max Timeout

The minimum/maximum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.

Max Connections Allowed

The maximum number of connections allowed. Currently, the number is unlimited.

Active Opens

The number of times TCP connections have made a direct transition to the “synSent” state from the “closed” state (refer to RFC 973).

Passive Opens

The number of times TCP connections have made a direct transition to the “synReceived” state from the “listen” state (refer to RFC 973).

Attempt Fails

The number of times TCP connections have made a direct transition to the “closed” state from either the “synSent” state or the “synReceived” state, plus the number of times TCP connections have made a direct transition to the “listen” state from the “synReceived” state.

Established Resets

The number of times TCP connections have made a direct transition to the “closed” state from either the “established” state or the “closeWait” state.

Currently Established

The number of TCP connections for which the current state is either “established” or “closeWait”.

Total Segments Received

The total number of segments received, including those received in error. This count includes segments received on currently established connections.

Total Segments Sent

The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets.

Total Segments Retransmitted

The number of TCP segments transmitted containing one or more previously transmitted octets.

Total Segments Received w/err

The total number of TCP segments that are in error; for example, bad TCP checksums.

Total Segments Sent w/RST flag

The number of TCP segments containing the RST flag.

Viewing the TCP Connection Table

The **tcpc** command is used to check the current TCP connections available in the TCP Connection Table.

To display the TCP Connection Table, enter the following command:

```
tcpc
```

A screen similar to the following displays:

TCP Connection/Listener Table						
Local Address/Port	Remote Address/Port	Recv-Q	Send-Q	Conn State		
127.0.0.1 / 1090	27.0.0.1 / 1091	0	0	ESTABLISHED		
127.0.0.1 / 1091	127.0.0.1 / 1090	0	322	ESTABLISHED		
198.206.184.42 / 23	198.206.184.34 / 2057	0	0	ESTABLISHED		
0.0.0.0 / 23	0.0.0.0 / 0	0	0	LISTEN		
0.0.0.0 / 21	0.0.0.0 / 0	0	0	LISTEN		

The fields on this screen have the following meanings:

Local Address/Port

The local IP address for this TCP connection and the local port for this TCP connection. In the case of a connection in the listen state which is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used.

Remote Address/Port

The remote IP address/the remote port number for this TCP connection.

Recv-Q

The number of segments received on this port.

Send-Q

The number of segments sent on this port.

Conn State

Describes the state of the TCP connection, as defined in RFC 973. Possible values are: closed, listen, synSent, synReceived, established, finWait1, finWait2, closeWait, lastAck, closing, time-Wait, and deleteTCB.

Using the TELNET Command

The **telnet** command is used to connect to another system. All of the standard TELNET commands are supported by the software in the switch.

To initiate a TELNET session, enter the following command:

```
telnet
```

A screen similar to the following displays:

```
telnet>
```

To display a listing of the TELNET commands, enter the following command:

```
?
```

A screen similar to the following displays:

```
Commands may be abbreviated. Commands are:
```

close	close current connection
display	display operating parameters
mode	try to enter line or character mode ('mode ?' for more)
open	connect to a site
quit	exit telnet
send	transmit special characters ('send ?' for more)
set	set operating parameters ('set ?' for more)
unset	unset operating parameters ('unset ?' for more)
status	print status information
toggle	toggle operating parameters ('toggle ?' for more)
environ	change environment variables ('environ ?' for more)
?	print help information

Enter the desired commands to establish and conduct your TELNET session.

Cancelling a Telnet request

If you initiate a Telnet session to an IP address that is not responding, after several seconds the switch will respond with the following message:

```
telnet: Unable to connect to remote host: S_error_ETIMEDOUT
```

If you don't want to wait for the switch to timeout on its own, you can cancel your request for a Telnet session by typing either **Ctrl-J** or **Ctrl-C**.

Tracing an IP Route

The **tracert** command is used to find the IP route from the local switch to a specified IP address destination. This command displays the individual hops to the destinations as well as some timing information. When using the **tracert** command, you must enter the name of the destination as part of the command line.

As an example, we might want to trace the route to “corporate.com.” To do so, we would enter this command:

```
tracert corporate.com
```

A screen similar to the following displays:

```
tracert to corporate.com (198.206.185.7),30 hops max,40 byte packets  
1 branch-wan-gw.CORPORATE.COM (198.206.181.252) 16 ms 0 ms 16 ms  
2 10.254.1.253 (10.254.1.253) 98 ms 81 ms 98 ms  
3 198.206.185.7 (198.206.185.7) 121 ms 81 ms 98 ms
```

Each number displayed corresponds to an individual hop. The time needed to reach that hop is shown (in milliseconds) after the hop’s IP address. The time may be followed by one of the following codes:

- !** The TTL of the received ICMP message is less than or equal to 1.
- !H** The host was unreachable.
- !N** The network was unreachable.
- !P** The protocol was unreachable.

If the time is replaced by an asterisk (*), no response was received from the host during the default 3-second timeout period.

Flushing the RIP Routing Tables

The **ripflush** command is used to flush all entries in the RIP Routing Table. All existing routes, with the *exception* of static and direct routes, are removed from the table by entry of the **ripflush** command.

To flush the RIP Routing Table, enter the following command:

ripflush

No message is displayed; the system prompt simply reappears.

Configuring IP RIP Filters

The **ipfilter** command is used to add or delete an IP RIP Output or Input filter. The IP RIP Filtering feature gives you a means of controlling the operation of the IP RIP protocol. By using IP RIP filters, you can minimize the number of entries that are put into the IP Forwarding Table as well as improve overall network performance by eliminating unnecessary traffic.

Two types of IP RIP filters are available:

- **RIP Input** filters control which IP networks are allowed into the switch's IP Forwarding Table whenever IP RIP updates are received.
- **RIP Output** filters control the list of IP networks that are included in the RIP Updates sent out by the switch on any interface. Thus, RIP Output filters effectively control which networks the router advertises in the RIP updates it generates.

Here are some example uses of IP RIP filters:

- RIP Input and Output filters can be used to isolate entire network segments (and/or routers) in order to make the network "appear" differently to the network's various segments.
- RIP Input and Output filters can be used to reduce the overall amount of WAN traffic that is needed to advertise routes that should not be used by a particular network segment.

◆ Important Note ◆

The IP RIP Filtering feature works *only* with the switch's standard RIP routing protocol. If you elect to use Alcatel's Advanced Routing feature (GateD) to provide RIP routing functionality in your switch, you will not be able to activate IP RIP Filtering.

Adding a "Global" IP RIP Filter

Follow the steps below to add a "global" IP RIP Output or Input filter.

1. Enter **ipfilter**.

A screen similar to the following displays:

Selecting global IP filter:

Add or delete entry {add(a), delete(d)} (a) :

Enter **a** (or just press **Enter**) to select to add a filter.

2. The following prompt displays:

**Filter type{RIP Output(ro),
RIP Input(ri)} (ro) :**

Enter **ro** (or just press **Enter**) to add a RIP Output filter. Enter **ri** to add a RIP Input filter.

3. The following prompt displays:

Filter action {block(b), allow(a)} (a) :

Enter **a** (or just press **Enter**) to set the filter action to "allow."

Enter **b** to set the filter action to "block."

- The following prompt displays:

IP address (default: all networks) :

Enter the IP address of the network that is to be allowed or blocked by the filter (or just press **Enter** to use the default of all networks). If you choose the default you will *not* be prompted for the network mask (as is shown in the next step).

- The following prompt displays:

IP network mask (default: 255.255.255.0) :

Enter the IP network mask of the network that is to be allowed or blocked by the filter (or just press **Enter** to use the default mask of 255.255.255.0). Note that the default mask will vary depending on the class of the IP address you entered above.

- A message displays indicating that the filter was successfully added:

ipfilter successfully added

Adding an IP RIP Filter For a Specific Group or VLAN

Follow the steps below to add an IP RIP Output or Input filter for a specific Group or VLAN.

- Enter the Group and VLAN numbers after the command like this: **ipfilter 1:1**.

A screen similar to the following displays:

Selecting IP filter for interface 1:1 :

Add or delete entry {add(a), delete(d)} (a) :

Filter action {block(b), allow(a)} (a) :

IP address (default: all networks) :

IP network mask (default: 255.255.255.0) :

ipfilter successfully added

Enter **a** (or just press **Enter**) to select to add a filter.

- The following prompt displays:

**Filter type{RIP Output(ro),
RIP Input(ri)} (ro) :**

Enter **ro** (or just press **Enter**) to add a RIP Output filter. Enter **ri** to add a RIP Input filter.

- The following prompt displays:

Filter action {block(b), allow(a)} (a) :

Enter **a** (or press **Enter**) to set the filter action to “allow.” Enter **b** to set the filter action to “block.”

- The following prompt displays:

IP address (default: all networks) :

Enter the IP address of the network that is to be allowed or blocked by the filter (or press **Enter** to use the default of all networks). If you choose the default you will *not* be prompted for the network mask (as is shown in the next step).

5. The following prompt displays:

IP network mask (default: 255.255.255.0) :

Enter the IP network mask of the network that is to be allowed or blocked by the filter (or just press **Enter** to use the default mask of 255.255.255.0). Note that the default mask will vary depending on the class of the IP address you entered above.

6. If the Group:VLAN is a WAN routing service, the following prompt displays:

Do you wish to apply this filter to a specific WAN endpoint? (n): y
Frame Relay VC or PPP Peer {vc(v), peer(p)} (v):

Enter **y** to apply this filter to a specific WAN endpoint.

7. The following prompt displays:

Frame Relay VC or PPP Peer {vc(v), peer(p)} (v):

Enter **v** (or just press **Enter**) to apply this filter to a Frame Relay VC.

Enter **p** if you want to apply this filter to a PPP Peer.

8. If you choose to apply the filter to a Frame Relay VC, this prompt will appear:

Slot/port:

Enter the slot and port numbers to which you want to apply this filter.

9. You will then be prompted for the virtual circuit (VC) to which to apply this filter:

VC:

Enter the VC to which you want to apply this filter.

10. If you choose to apply a filter to a PPP Peer, this one prompt will appear:

Peer ID:

Enter the Peer ID to which you want to apply this filter.

A message will appear indicating that the filter was successfully added.

IP RIP Filter Precedence

Whenever you use multiple “allow” filters you must first define a filter to block all RIPs. Then, any other “allow” filters of the same type must be *at least* as specific in all areas in order for the filters to work. Note that filtering precedence is related only to “allow” filters. Multiple “block” filters can be defined with varying specificity in each of the areas of the filter. The filtering done by the configurable parameters (Address/Mask) in the “allow” filter must be at least as specific as the filtering defined in the “block” filter.

Deleting IP RIP Filters

Follow the steps below to delete an existing IP RIP Output or Input filter.

1. Enter **ipfilter**.

A screen similar to the following displays:

Selecting global IP filter:

Add or delete entry {add(a), delete(d)} (a) :

Enter **d** to select to delete a filter.

2. A screen similar to the following displays:

Displaying all filters:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.0.0.0	255.0.0.0	A	global
2	RIP IN	99.0.0.0	255.0.0.0	B	2:1

Entry number to delete? (default: none) :

This screen contains a list of the existing IP RIP filters. The fields on this screen are described in the next section (see *Displaying IP RIP Filters* on page 25-37).

3. Enter the index number of the filter that you want to delete. If you decide at this point that you want to abort out of the deletion process, simply press **Enter** to accept the default of “none”.
4. A message will confirm the deletion of the filter:

ipfilter successfully deleted

Displaying IP RIP Filters

The **ipf** command is used to display a list of all existing IP RIP Output and Input filters. See *Configuring IP RIP Filters* on page 25-33 for complete information on creating these filters.

Displaying a List of All IP RIP Filters

To display the listing of all existing IP RIP filters, enter the following command:

```
ipf
```

A screen similar to the following displays:

Displaying all filters:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.0.0.0	255.0.0.0	A	global
2	RIP IN	99.0.0.0	255.0.0.0	B	2:1
3	RIP OUT	All Networks		B	5:1 (3/1/32)
4	RIP IN	All Networks		B	6:1 (P1)

This screen contains a list of the existing IP RIP filters. The fields on this screen have the following meanings:

#

Indicates the index number assigned to identify this filter.

Type

Indicates the type of filter, either RIP Input (**RIP IN**) or RIP Output (**RIP OUT**).

Network

Indicates the IP address that is to be filtered (entered in dotted-decimal format). An entry of “All Networks” means that all addresses are to be filtered.

Mask

The IP network mask of the network to be filtered (entered in dotted-decimal format). This field is blank if the network entered is “All Networks.”

Md

Indicates the filter’s mode of operation, either to “allow” traffic (**A**) or to “block” traffic (**B**).

GP:VL (s/p/vc) or (Peer ID)

The first number (**GP**) is the Group associated with this entry. The second number (**VL**) is the VLAN associated with this entry. When a filter applies to all interfaces, this field will say “global.” If an entry refers to a Frame Relay interface, column headings for slot, port, and virtual circuit (**s/p/vc**) may be displayed when the filter is applied to a particular virtual circuit rather than to the entire VLAN. If an entry refers to a PPP interface, the Peer ID (**Peer ID**) may be displayed when the filter is applied to a particular PPP Peer.

Displaying a List of "Global" IP RIP Filters

To display a listing of just the global IP RIP filters, enter the following command:

```
ipf global
```

A screen similar to the following displays:

Displaying global filters:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.99.99.99	255.0.0.0	A	global

Displaying a List of Specific IP RIP Filters

To display a listing of IP RIP filters for a specific interface, you can specify other parameters along with the **ipf** command. The format for the command in this case is:

```
ipf <type> <GP:VL>
```

The type is one of these codes:

ri for RIP INput

ro for RIP OUTput

For example, to display a list of the filters defined for Group 2, VLAN 1, you would enter:

```
ipf 2:1
```

A screen similar to the following would be displayed:

Displaying filters for interface 2:1:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP IN	99.0.0.0	255.0.0.0	B	2:1

As another example, to display a list of all global RIP Output filters, you would enter:

```
ipf ro global
```

A screen similar to the following would be displayed:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.0.0.0	255.0.0.0	A	global

Viewing the IP-to-MAC Address Table

The **ipmac** command is used to display the IP-to-Mac Address Association Table. This table contains a listing of IP addresses and their associated MAC (Media Access Control) addresses together with the slot/port from which the information was learned. The information in this table is learned from ARP (Address Resolution Protocol) messages received on “leaf” ports. A “leaf” port is one on which Spanning Tree has been disabled or on which no Spanning Tree BPDUs have yet been received.

The **ipmac** command can be very helpful in resolving certain problems. For example, in large networks where hosts are frequently moved around, users can experience connectivity problems. In this situation, the **ipmac** command can be used to help locate a particular IP workstation. Another use is to help resolve duplicate IP addresses on a network. The program checks all ARP messages, whether they are received on a “leaf” port or not, against those in its table to see if a duplicate IP address exists. If a duplicate is detected, an SNMP trap message is generated and the duplicate can easily be seen in the table produced by the **ipmac** command.

The **ipmac** command can be entered alone in which case it will display all entries currently in the table, or you may enter a specific IP address along with the command to show only the information related to that IP address. An optional parameter (-f) can be entered to flush the table. Each of these uses of the **ipmac** command is illustrated below.

Displaying All Entries in the IP-to-MAC Table

To display the list of all the entries in the IP-to-MAC table, enter the following command:

```
ipmac
```

A screen similar to the following displays:

IP to MAC ADDRESS ASSOCIATION TABLE

IP Address	MAC Address	Slot / Intf
192. 168. 10. 1	0020DA:6DE610	4 / 5
172. 16. 0. 5	0020DA:76D3D0	3 / 2
172. 16. 0. 7	00E029:00D41E	3 / 2
172. 16. 0. 41	0000C0:24FFEC	3 / 2
172. 16. 0. 47	00A0C9:0AA907	3 / 2
172. 16. 0. 28	0020DA:7AE9D3	3 / 2
172. 16. 0. 45	080020:8AE301	3 / 2
172. 16. 0. 60	0020DA:73C3A0	3 / 2
172. 16. 30. 00	0020AF:04BA57	3 / 2
172. 16. 41. 03	0000C0:AD8EE9	3 / 2
172. 16. 50. 12	080020:7B79E1	3 / 2
172. 16. 255.254	0020DA:6F97E5	3 / 2
*****	0020DA:032273	5 / 1
192. 168. 10. 1	0020DA:7AEA60	3 / 2
198. 206. 182.222	0020DA:7F48A0	3 / 2

The fields on this screen have the following meanings:

IP Address

The IP address learned from ARP messages received on “leaf” ports. A series of asterisks (*****) in this field indicates that the preceding entry is a duplicate to this entry. In the example screen shown above, the address 172.16.255.254 is assigned to two MAC addresses.

MAC Address

The MAC address corresponding to the listed IP address.

Slot/Intf

The slot number and interface number from which the IP and MAC addresses were learned.

Displaying Information for a Specific IP Address

To display the entry in the IP-to-MAC table for a specific IP address, enter the desired IP address after the command. For example, to locate the entry for IP address 192.168.10.1, enter the following command:

```
ipmac 192.168.10.1
```

A screen similar to the following displays:

IP to MAC ADDRESS ASSOCIATION TABLE

IP Address	MAC Address	Slot / Intf
192.168. 10. 1	0020DA:6DE610	4 / 5

Flushing Entries from the Table

To flush all the entries in the IP-to-MAC table, enter the following command:

```
ipmac -f
```

The system prompt redisplay.

Enabling/Disabling Directed Broadcasts

An IP directed broadcast is an IP datagram that has all zeroes or all 1's in the host portion of the destination IP address. The packet is sent to the broadcast address of a subnet to which the sender is not directly attached. The datagram is routed through the network as a unicast packet. When it arrives at the subnet, it is converted into a broadcast packet.

Directed broadcasts are used in denial-of-service *smurf* attacks. In a smurf attack, a continuous stream of ping requests are sent from a falsified source address to a directed broadcast address, resulting in a large stream of replies, which can overload the host of the source address.

By default, the switch drops directed broadcasts. Typically, directed broadcasts should not be enabled.

To enable directed broadcasts to be routed through the switch:

1. At the system prompt, enter the **ipdirbcast** command.
2. Enter **y** to enable direct broadcasts.

Path MTU Discovery

All Gigabit Ethernet modules and all Mammoth-based Ethernet modules on the Omni Switch/Router in Release 4.0 and later support path Maximum Transmission Unit (MTU) discovery. In path MTU discovery, the Ethernet frame (datagram) size is set to the largest size that does not require fragmentation anywhere along the path from a source host to its destination. This frame size, known as a Path MTU (PMTU), is thus equal to the minimum of the MTUs of each hop in the path.

◆ Note ◆

MTU discovery is *not* supported on token ring, FDDI, WAN, or non-Mammoth Ethernet modules. However, token ring and FDDI can be used as intermediate links (e.g., trunking or bridging) between remote switches.

Path MTU discovery is active all of the time and is part of the switch's operating system; you do not need configure it.

The source host initially assumes that the PMTU of a path is the MTU of the first hop. It sends all datagrams with the "Don't Fragment" (DF) bit set. If a switch/router along the path receives a datagram that is too large to forward without fragmentation, the following steps will be executed:

1. The switch/router that cannot forward these datagrams (i.e., the constricting hop) will discard them.
2. The constricting hop will send ICMP destination unreachable messages to the source host with a code that indicates fragmentation is needed and the "Don't Fragment" (DF) bit in the Internet Protocol (IP) header has been set. This message (known as a "Datagram Too Big" message) contains the PMTU of the constricting hop.
3. After receiving a "Datagram Too Big" message, the source host reduces the size of the MTU so it matches the PMTU of the constricting hop.
4. The MTU discovery process ends when datagrams can be sent without fragmentation. However, the source host will *not* reduce the size of a datagram below 68 octets.

26 UDP Forwarding

UDP is a connectionless transport protocol that is used for applications that do not require the establishment of a session and end-to-end error checking, such as email and file transfer. This chapter describes the UDP relay function in the switch, which allows UDP broadcast packets to be forwarded across groups and VLANs that have IP routing enabled. The UDP relay allows you to use nonroutable protocols in a routing environment. (For information about IP routing, see Chapter 25, “IP Routing.”)

◆ Note ◆

BOOTP/DHCP relay has previously been available on the switch. It is now part of an expanded feature that includes relays for NetBIOS and generic services.

The relay may be configured for the following services:

- Bootstrap Protocol (BOOTP)/Dynamic Host Configuration Protocol (DHCP)
- NetBIOS Name Server (NBNS)
- NetBIOS Datagram Distribution Server (NBDD)
- Generic applications, such as Trivial File Transfer Protocol (TFTP)

The UDP services, their corresponding well-known port numbers, and configurable options on the switch are listed here.

Service	UDP Port No.	Configurable Options
BOOTP/DHCP	67/68	Next-hop address (up to 8) Forward delay Maximum hops
NBNS	137	Next-hop address (up to 8) Forwarding VLANs (up to 32)
NBDD	138	Next-hop address (up to 8) Forwarding VLANs (up to 32)
Generic	user-configured	Next-hop address (up to 8) Forwarding VLANs (up to 32)

UDP Relay and RIF Stripping

Routing Information Field (RIF) stripping is required for transparent bridge ports in source route environments and may also be useful in non-source route environments.

In a source route environment, where RIF stripping is enabled for transparent bridging to Ethernet, UDP relay clients should not be more than one switch away from the DHCP server. (In RIF stripping, 2 bytes are stripped from the RIF and each bridge adds 2 bytes to the RIF. Packets with a RIF greater than 2 bytes are discarded.)

In non-source route environments, RIF stripping may be required if DHCP clients are token ring stations. Token ring stations may have packets with RIFs even though source routing is not enabled on the station. RIF stripping is required if there is bridging to Ethernet, FDDI, or 802.3 LANE anywhere along the path between the client and the DHCP server. RIF stripping should be enabled on the first non-token ring port in the path. The number of bridges on the path does not matter.

UDP Relay Hardware/Software Support

The UDP forwarding feature has the following hardware/software support:

- UDP relay is supported on any Omni Switch/Router (OmniS/R).
- To relay DHCP requests from authentication clients in a default group to a DHCP server in an authenticated group, the **avlbootpmode** command must be used in addition to the **relayc** command described in this chapter. See the Authentication Services chapter of the *Switched Network Solutions User Manual* for information about the **avlbootpmode** command.

UDP Relay Configuration Screen

To configure any of the UDP relays, use the **relayc** command. The **relayc** command is listed in the IP submenu. (For more information about IP commands, see Chapter 25, “IP Routing.”) The screen display is similar to the following:

UDP Relay Configuration

```
1) BOOTP/DHCP Enabled      : No
2) NBNS Enabled            : No
3) NBDD Enabled            : No
4) +Generic Services Menu
```

Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :

Use the UDP Relay Configuration screen to enable any of the relays and display more configuration options for enabled relays. The following sections describe each UDP service and how to configure each of the relays using the User Interface (UI). A UDP statistics screen may also be displayed.

◆ Note ◆

For general information about the UI, see Chapter 4, “The User Interface.”

BOOTP/DHCP Relay

The switch supports a UDP relay function that allows Bootstrap Protocol (BOOTP) and Dynamic Host Configuration Protocol (DHCP) packets to pass between AutoTracker Groups.

◆ **Note** ◆

A BOOTP/DHCP relay may be configured for authenticated groups as well. See *BOOTP/DHCP Relay and Authentication* on page 26-5 and the Authentication Services chapter of the *Switched Network Solutions User Manual*.

Through UI software, you can turn the relay function on or off and specify the IP addresses of DHCP servers, the delay before the relay forwards a request, and the maximum number of hops a packet may be forwarded through the network.

Alternately the relay function may be provided by an external router connected to the switch; in this case, the relay would be configured on the external router.

Overview of DHCP

DHCP provides a framework for passing configuration information to Internet hosts on a TCP/IP network. It is based on the Bootstrap Protocol (BOOTP), adding the ability to automatically allocate reusable network addresses and additional configuration options. DHCP consists of the following two components:

- A protocol for delivering host-specific configuration parameters from a DHCP server to a host.
- A mechanism for allocating network addresses to hosts.

DHCP is built on a client-server model in which a designated DHCP server allocates network addresses and delivers configuration parameters to dynamically configured hosts. It supports the following three mechanisms for IP address allocation:

Automatic	DHCP assigns a permanent IP address to a host.
Dynamic	DHCP assigns an IP address to a host for a limited period of time (or until the host explicitly relinquishes the address).
Manual	The network administrator assigns a host's IP address and DHCP simply conveys the assigned address to the host.

A particular network will use one or more of these mechanisms, depending on the policies of the network administrator.

For information about configuring DHCP servers, see the IP Control chapter of the *Switched Network Solutions User Manual*.

DHCP and the OmniS/R

The unique characteristics of the DHCP protocol require a good plan before setting up the switch in a DHCP environment. Since DHCP clients initially have no IP address, placement of these clients in an AutoTracker VLAN is hard to determine. In simple networks (i.e., one group, one VLAN) AutoTracker rules do not need to be deployed to support the BOOTP/DHCP relay functionality.

In multiple group configurations, AutoTracker rules can be deployed to strategically support the relay function. Two types of AutoTracker IP policies are appropriate for DHCP environments. The first is the IP protocol policy that puts all IP type frames into a single VLAN regardless of network address. The second is the IP network policy that groups IP users based on their specific IP address.

Besides AutoTracker rules, the network administrator must be aware that some network environments may contain DHCP-ready and non-DHCP clients. Such configurations are supported by the switch's BOOTP relay function.

BOOTP/DHCP Relay and Source Routing

In source route environments (where VLAN framing type is set for source routing) and DHCP clients are not directly attached to the switch but have one or more bridges between them, the **mpx.cmd** file must be modified so that replies from the DHCP server can get through the bridge.

Typically a router caches the client's RIF information for source routing when the client responds to an ARP, but if the client does not yet know its IP address it cannot reply to an ARP and no RIF information is cached on the router. Unicast replies to the client before the RIF is cached are discarded by the router. Forcing the BOOTP reply to be broadcast eliminates this problem.

Use the **edit** command to make this change to the **mpx.cmd** file (see Chapter 7, "Managing Files," for instructions on using the **edit** command).

Add the following command:

```
bootpBcastReply=1
```

Reboot the switch to force the broadcast. Replies from the DHCP server to the client will be broadcast from the router as STE or ARE packets so they can be sent through the bridge.

BOOTP/DHCP Relay and Authentication

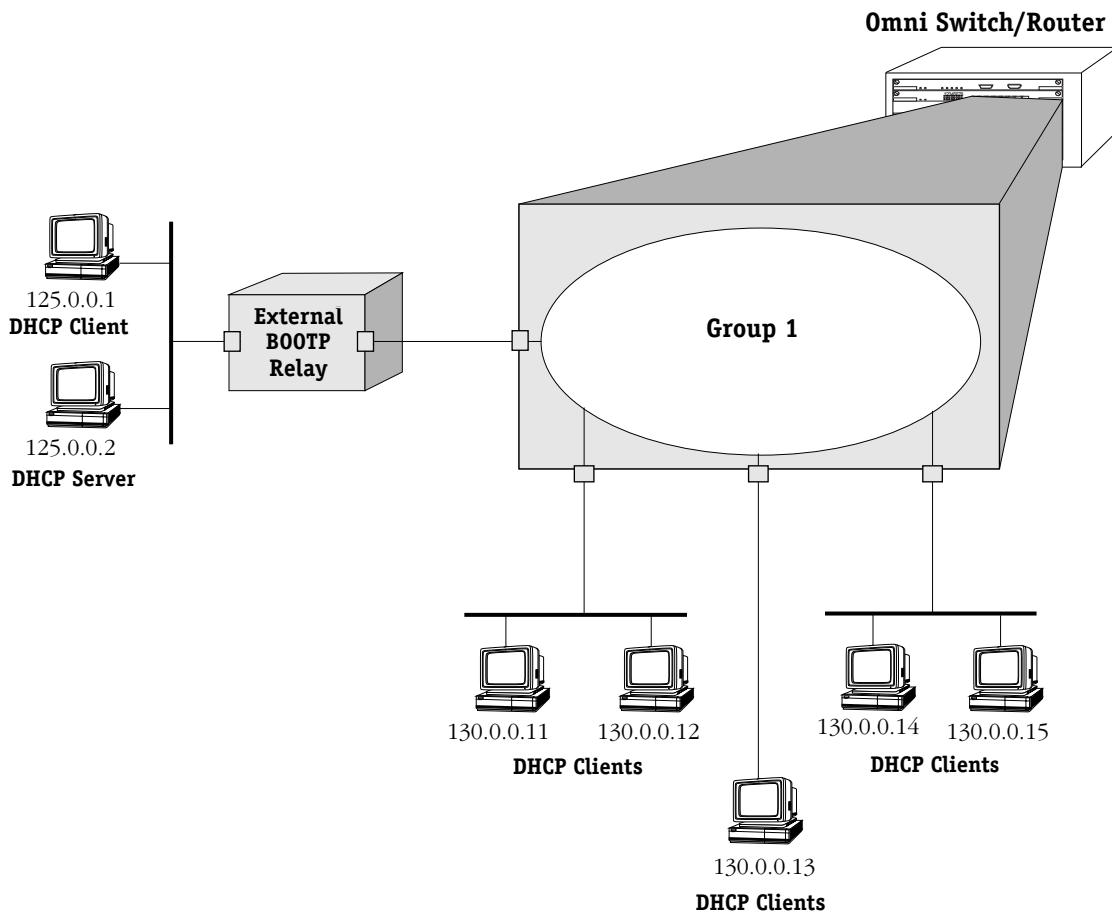
Authentication clients may use DHCP to get an IP address. For Telnet authentication clients, an IP address is required for authentication. The DHCP server may be located in the default group, an authenticated group, or both. If authentication clients will be getting an IP address from a DHCP server located in an authenticated group, a BOOTP/DHCP relay should be set up using the **relayc** command described in this chapter.

In addition, the router port address of the authenticated group must also be configured for the relay through the **avlbootpmode** command. See the Authentication Services chapter of the *Switched Network Solutions User Manual* for more information about this command.

External BOOTP Relay

The BOOTP relay may be configured on a router that is external to the switch. In this application example the switched network has a single AutoTracker Group configured with multiple segments. All of the network hosts are DHCP-ready, meaning they obtain their network address from the DHCP server. The DHCP server resides behind an external network router, which supports the BOOTP relay functionality.

One requirement for routing DHCP frames is that the router must support BOOTP relay functionality to be able to forward DHCP frames. In this example, BOOTP relay is supported within an external router, which forwards request frames from the incoming router port to the outgoing router port attached to the Omni Switch/Router.



DHCP Clients are Members of the Same VLAN


The external router inserts the subnet address of the first hop segment into the DHCP request frames from the DHCP clients. This subnet address allows the DHCP server to locate the segment that the requesting client resides on. In this example, all clients attached to the Omni Switch/Router are DHCP-ready and will have the same subnet address (130.0.0.0) inserted into each of the requests by the router's BOOTP relay function. The DHCP server will assign a different IP address to each of the clients. The switch does not need an IP address assigned and all DHCP clients will be members of either a default VLAN or an IP protocol VLAN.

Internal BOOTP/DHCP Relay

The internal BOOTP/DHCP relay is configured using the UDP forwarding feature in the switch, available through the **relayc** command. See *UDP Relay Configuration Screen* on page 26-3.

Example 1

This application example shows a network with two AutoTracker Groups, each with multiple segments. All network clients are DHCP-ready and the DHCP server resides on just one of the groups. This example is much like the first application example, except that the BOOTP relay function is configured inside the switch.



DHCP Clients in Two Groups

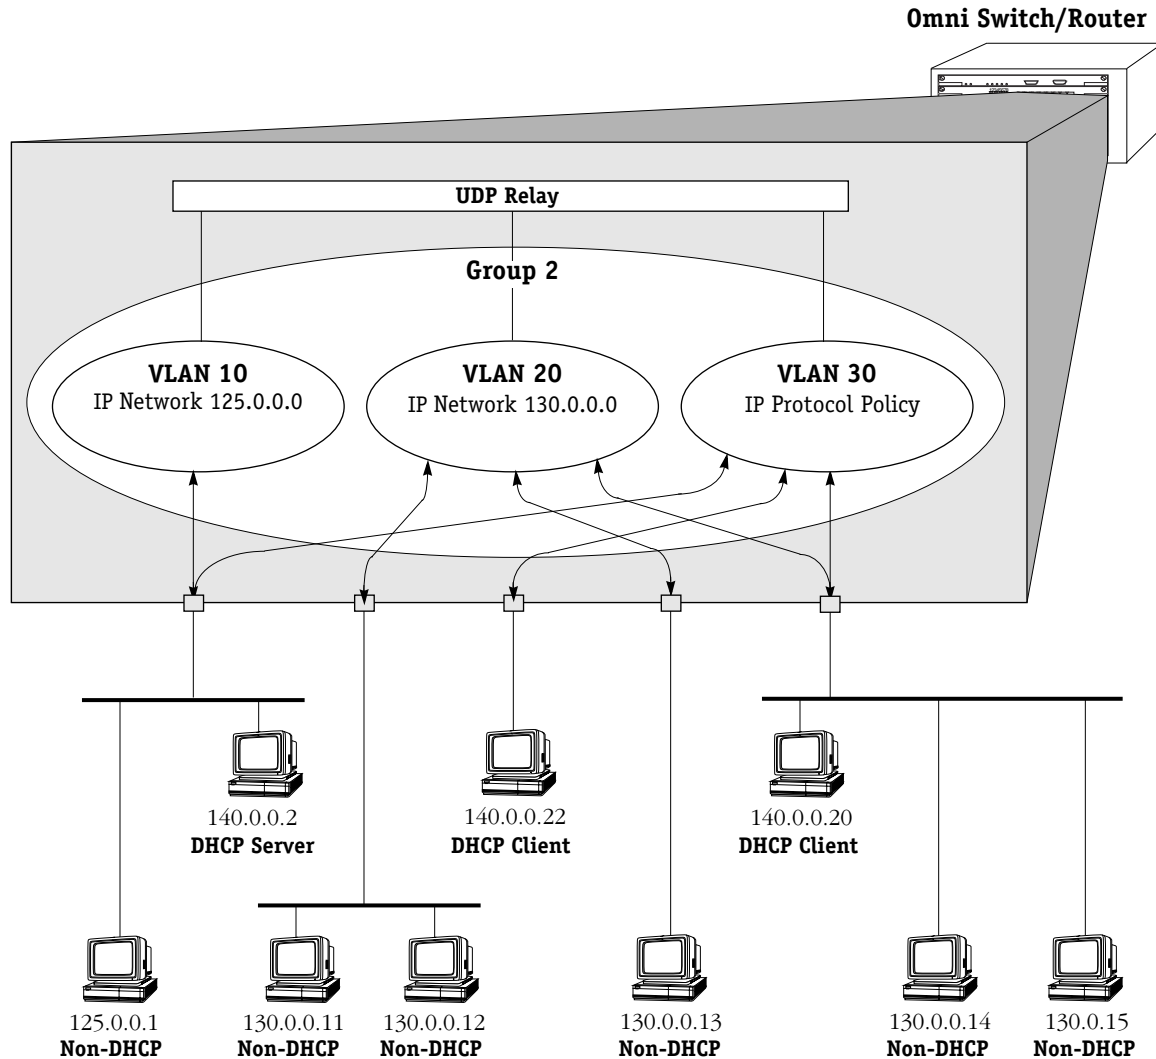
During initialization, each network client forwards a DHCP request frame to the DHCP server using the local broadcast address. For those stations locally attached, the frame will simply be switched.

In the example above, the DHCP server and clients in the same group must be members of the same VLAN so that the policies match (they could also all be members of the default VLAN). One way to accomplish this is to use an IP protocol policy that places all IP frames in the same VLAN. An IP network policy would not work in this case because the DHCP clients will not have an IP network address until *after* they communicate with the DHCP server.

Because the clients in group 3 are not on the same segment as the DHCP server, they must request an IP address via the BOOTP relay routing entity in the switch. When a DHCP request frame is received by the BOOTP relay entity, it will be forwarded from group 3 to group 2. All the DHCP-ready clients in group 3 must be members of the same VLAN, and the switch must have the BOOTP relay function configured.

Example 2

This application example has a single group in a network with a mix of DHCP-ready and non-DHCP clients. DHCP-ready and non-DHCP clients can coexist in the same network, group, or segment. There are two types of AutoTracker policies defined within the group—IP address and IP protocol.



AutoTracker IP Policy Places DHCP Clients in Same VLAN

Whenever AutoTracker receives an IP frame, it examines the frame for an IP network layer policy match. In the case of DHCP, the client generates an IP frame without an IP address. Without an IP address, AutoTracker will not be able to place the client into a VLAN based on IP address. Therefore, the client will become a member of the VLAN defined by a general IP Protocol policy (i.e., VLAN 30).

In this example, the VLAN defined by an IP protocol policy is used *as a mechanism to group the DHCP server and associated clients*. The DHCP server is local, so all clients requesting an IP address will be allocated an IP address on the same subnet.

◆ Note ◆

This configuration works if you require only one DHCP subnet. All clients received on the same router port will be assigned to the same VLAN.

Note that the client's request frames will also be received and forwarded by the BOOTP relay if it is configured.

The non-DHCP workstations will be assigned to VLANs defined by Network Address policies. These workstations already have manually configured IP addresses. They don't require a server to dynamically assign them an address. AutoTracker will move these workstations into the VLANs with IP network address policies (VLAN 10 and VLAN 20).

It is true that these non-DHCP workstations also match the IP protocol policy. However, Network Address policies have precedence over IP protocol policies. If AutoTracker finds a match on a Network Address policy, it does not look for a protocol policy match.

Enabling BOOTP/DHCP Relay

To enable UDP relay for BOOTP/DHCP:

At the prompt for the UDP Relay Configuration screen (the UDP Relay Configuration screen is displayed using the **relayc** command described in *UDP Relay Configuration Screen* on page 26-3), enter the following:

1=y

The screen redisplay with more configuration options for BOOTP/DHCP.

UDP Relay Configuration

```

1) BOOTP/DHCP Enabled           : Yes
  11) Server Address(list/add/delete) : UNSET
  12) Forward Delay              : 3
  13) Maximum Hops               : 4
2) NBNS Enabled                 : No
3) NBDD Enabled                 : No
4) +Generic Services Menu

```

Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :

The parameters are defined here.

Server Address

This parameter allows you to list, add, or delete the server address(es) to which the BOOTP/DHCP relay will forward. The default is **UNSET**. When you have configured at least one valid address, the value redisplay as **SET**. Up to 8 addresses may be configured. *The server address cannot be an internal DHCP server configured through the IP Control feature. For more information about IP Control, see the **Switched Network Solutions User Manual**.*

Forward Delay

The amount of time (typically in seconds, but determined by the client) the BOOTP/DHCP relay will wait before forwarding a request to the server address. This delay gives a local server a chance to respond to a client before the relay forwards it further out in the network. This value may range from 1 to 65535.

Maximum Hops

The maximum number of relays that a packet can go through while traversing the network. This limit keeps packets from “looping” through the network. Set this value to the maximum number of BOOTP/DHCP relays you expect packets to traverse. This value may range from 1 to 16.

Configuring BOOTP/DHCP Relay Parameters

At least one server address must be configured for the BOOTP/DHCP relay. To configure the server address:

1. On the UDP Relay Configuration screen prompt, enter

11=a

A screen displays similar to the following:

FORWARD TO Server List

Item	Server address	Server Name (if known)
------	----------------	------------------------

Enter IP address or host name of server to be added to the list ['h' for help/<ret> to exit]

2. Enter the IP address, which may be a specific host on the network or a subnet broadcast address. The address should be in dotted decimal format (i.e., 198.206.181.12) or hexadecimal address (i.e., 0xc6ceb501). Alternately you may enter a host name (i.e., system.com) if the DNS resolver is enabled on the switch through the **res** command. The screen redisplay with the entry.
3. Repeat the previous step to add all the addresses to which you want to forward to. Press **Enter** when you are finished adding addresses. The screen redisplay with the Server Address field set to **SET**.
4. Make any changes to Forward Delay or Maximum Hops.
5. Enter **s** to save your changes. If the relay has just been enabled, the system initializes the relay. If the relay is already running, it is stopped and reinitialized with the changes.
6. Enter **q** to quit the UDP Relay Configuration screen.

By default, Alcatel's implementation of BOOTP rejects packets less than 300 bytes. To prevent BOOTP from discarding packets smaller than 300 bytes add the following line to the **mpx.cmd** file:

bootpSizeCheck=0

This line must appear before the **cminit** line.

NetBIOS Relays

The switch supports a UDP relay function that allows Network Basic Input/Output System (NetBIOS) messages to be sent across groups or VLANs.

Overview of NetBIOS

NetBIOS is an applications interface that allows computers on Ethernet or token ring LANs to communicate with one another. An enhanced version of the protocol is used by networking operating systems such as LAN Manager and Windows NT.

With NetBIOS, each client and host in the LAN has a unique NetBIOS name. Stations in a NetBIOS network broadcast queries to verify that their names are unique on the LAN. Names may be verified by using the NetBIOS Name Server (NBNS) protocol, which sends messages to a well-known UDP port (137). Name requests are sent to an IP subnet broadcast address or the unicast address of the server.

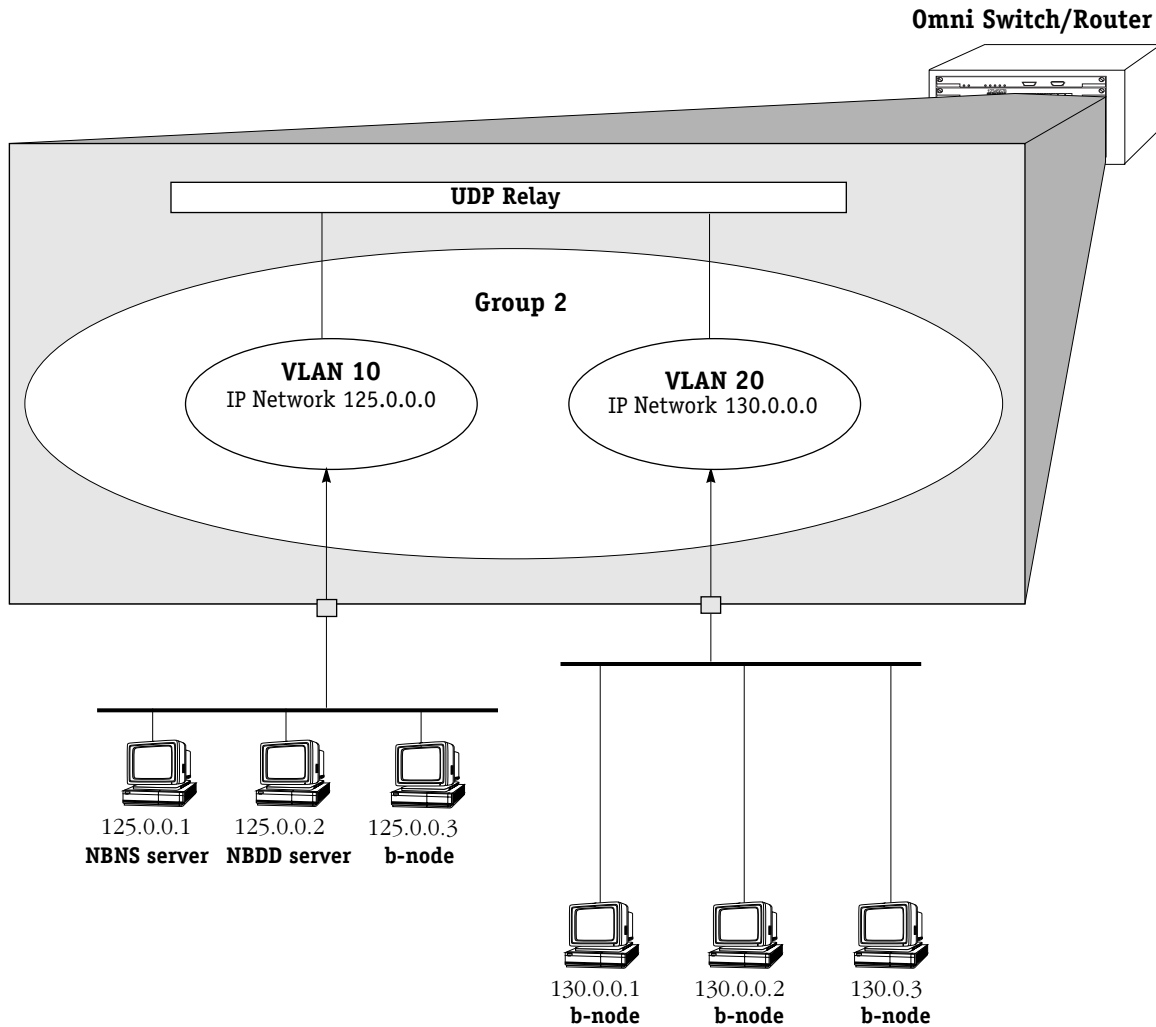
The NetBIOS protocol also has a datagram service that allows an application to exchange datagrams with a specific application or to broadcast and receive datagrams on a VLAN. A NetBIOS Datagram Distribution (NBDD) server may be installed in the network to provide this service, which uses a well-known UDP port number (138).

NetBIOS may be configured to run over TCP/IP using the various modes defined in RFC_1001 and RFC_1002. These modes are defined in terms of *nodes* and determine how NetBIOS stations (or nodes) in the network register their NetBIOS names and resolve (or map) these names to IP addresses. Each NetBIOS PC has a node type. The nodes are defined as follows:

- Broadcast node (b-node)—These nodes use broadcast for name registration and resolution. Since routers do not forward broadcast traffic, b-node clients in different networks will not be able to communicate
- Point-to-point node (p-node)—In this environment, each station knows the address of the server. Unicast queries are sent directly to the name and datagram servers. This method does not use broadcast.
- Mixed-mode node (m-node)—This mode uses a combination of b-node and p-node.

NetBIOS Relay Application

The UDP relay function in the switch extends b-node functionality across the internet. An example application is shown here.



NetBIOS Relay Application

In this example, NBNS and NBDD relays should be configured to forward to VLAN 10 and VLAN 20. The VLANs must be selected for forwarding, or you can configure the subnet address of the VLANs as next-hop addresses. The server addresses could be configured as next-hop addresses, but since the assignment of the NBNS and NBDD servers are by definition dynamic, configuring the VLAN number or the IP address of the VLAN ensures that the relay will function even if a server address changes.

Configuring NBNS Relay

Each NetBIOS PC has a name unique for its local network. If you are using NetBIOS broadcast queries to resolve names and NetBIOS clients are located in different groups or VLANs, you should configure UDP relay for NBNS.

The relays are enabled from the UDP Relay Configuration screen and are configured in similar ways. The UDP Relay Configuration screen is displayed using the **relayc** command described in *UDP Relay Configuration Screen* on page 26-3. To enable the NBNS relay, enter **2=y**. (To enable the NBDD relay, enter see *Configuring NBDD Relay* on page 26-16.)

The UDP Relay Configuration screen redispays similar to the following:

```

UDP Relay Configuration
1) BOOTP/DHCP Enabled           : Yes
  11) Server Address{list/add/delete} : UNSET
  12) Forward Delay              : 3
  13) Maximum Hops               : 4
2) NBNS Enabled                 : Yes
  21) Next-hop Address {list/add/delete} : UNSET
  22) Forward to VLANs {list/add/delete} : UNSET
3) NBDD Enabled                 :No
4) +Generic Services Menu

```

Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :

Either a Next-hop Address *or* a Forward to VLANs value must be configured for the relay.

Next-hop Address

Use this parameter to list, add, or delete the server address(es) to which the NBNS UDP relay will forward. The default is **UNSET**. The value redispays as **SET** when at least one address is configured. Up to 8 addresses may be configured. You can list, add, or delete addresses by entering **21=l**, **21=a**, or **21=d** on the command line.

Forward to VLANs

Use this parameter to list, add, or delete VLANs to which the NBNS UDP relay will forward. This default is **UNSET**. This value redispays as **SET** when at least one VLAN is configured. Up to 32 VLANs may be configured. You can list forwarding VLANs, or add or delete VLANs from the forwarding list by entering **22=l**, **22=a**, or **22=d** on the command line. Entries marked with an asterisk indicate the VLANs to which the relay will forward.

Next-Hop Addresses for NBNS

At least one next-hop address (or a forwarding VLAN as described in *Forwarding VLANs for NBNS Relay* on page 26-15) must be configured.

To *add* a next-hop address for NBNS relay:

1. On the UDP Relay Configuration screen command, enter the following:

21=l

A screen similar to the following displays:

FORWARD TO Server List		
Item	Server address	Server Name (if known)
1)	172. 28. 5.212	

Enter IP address or host name of server to be added to list ['h' for help/<ret> to exit]:

2. Enter the IP address of the next hop. Enter the address in dotted decimal format (i.e., 198.206.181.12), a hexadecimal address (i.e., 0xc6ceb501). A host name (i.e., system.com) may be entered if the DNS resolver is enabled using the **res** command.

◆ Note ◆

This address may be the unicast address of the server or a subnet broadcast address of the subnet where the server is located. Using a unicast address is not recommended because an NBNS by definition may shift part or all of its responsibility to another node in the network segment.

3. Enter any additional addresses up to a maximum of 8. Press **<Enter>** to return to the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

To *delete* next-hop addresses for the NBNS relay:

1. Enter **22=d** at the command prompt of the UDP Relay Configuration screen. The FORWARD TO Server List displays.
2. Enter the item number that corresponds to the entry that you want to delete. Repeat this step to delete any additional entries.
3. Press **<Enter>** to return to the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

Forwarding VLANs for NBNS Relay

At least one forwarding VLAN (or a next-hop address as described in *Next-Hop Addresses for NBNS* on page 26-14) must be configured for NBNS relay.

To *select* forwarding VLANs for NBNS relay:

1. On the command line of the UDP Relay Configuration screen, enter the following:

22=i

A screen similar to the following displays:

```

Available/Selected VLANS
Item  Group ID:VLAN ID      MASK      IP ADDR
  1)   1:1              255.255. 0. 0    172. 23. 9.105  *
* = selected for forwarding

```

Enter item number of VLAN to be selected ['h'f or help/<ret> to exit] :

2. Enter the item number of the group/VLAN that you want to select. Repeat this step for all the groups/VLANs you want to select.
3. Press **<Enter>** to return to the command line for the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

To *deselect* forwarding VLANs:

1. On the UDP Relay Configuration screen, enter

22=d

The Available/Selected VLANs screen displays.

2. Enter the item number of the group/VLAN that you want to select. Repeat this step for all the groups/VLANs you want to select.
3. Press **<Enter>** to return to the command line for the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

Configuring NBDD Relay

If you want to send NetBIOS datagrams across networks, you should enable the NBDD relay. To enable the NBDD relay, enter **3=y** at the command prompt of the UDP Relay Configuration screen. The screen redisplay is similar to the following:

```

                                UDP Relay Configuration
1) BOOTP/DHCP Enabled           : Yes
   11) Server Address{list/add/delete} : UNSET
   12) Forward Delay              : 3
   13) Maximum Hops              : 4
2) NBNS Enabled                 : Yes
   21) Next-hop Address {list/add/delete} : UNSET
   22) Forward to VLANs {list/add/delete} : UNSET
3) NBDD Enabled                 : Yes
   31) Next-hop Address {list/add/delete} : UNSET
   32) Forward to VLANs {list/add/delete} : UNSET
4) +Generic Services Menu

```

Command {Item=Value/?/Help/Quit/Redraw/Save} (Redraw) :

Either a Next-hop Address *or* a Forward to VLANs value must be configured for the relay.

Next-hop Address

Use this parameter to list, add, or delete the server address(es) to which the NBNS UDP relay will forward. The default is **UNSET**. This value redisplay as **SET** when at least one address is configured. Up to 8 addresses may be configured. You can list, add, or delete addresses by entering **31=l**, **31=a**, or **31=d** on the command line.

Forward to VLANs

Use this parameter to list, add, or delete VLANs to which the NBNS UDP relay will forward. This default is **UNSET**. This value changes to **SET** when at least one VLAN is configured. Up to 32 VLANs may be configured. You can list forwarding VLANs, or add or delete VLANs from the forwarding list by entering **32=l**, **32=a**, or **32=d** on the command line. Entries marked with an asterisk indicate the VLANs to which the relay will forward.

Next-Hop Addresses for NBDD

At least one next-hop address (or a forwarding VLAN as described in *Forwarding VLANs for NBDD Relay* on page 26-18) must be configured for the relay.

To *add* a next-hop address for NBDD relay:

1. At the command prompt for the UDP Relay Configuration screen, enter the following:

32=a

A screen similar to the following displays:

```

FORWARD TO Server List
Item      Server address      Server Name (if known)
1)        172. 28.  5.212

```

Enter IP address or host name of server to be added to list ['h' for help/<ret> to exit]:

2. Enter the IP address of the next hop. Enter the address in dotted decimal format (i.e., 198.206.181.12), a hexadecimal address (i.e., 0xc6ceb501). A host name (i.e., system.com) may be entered if the DNS resolver is enabled using the **res** command.

◆ Note ◆

This address may be the unicast address of the server or a subnet broadcast address of the subnet where the server is located. Using a unicast address is not recommended because an NBNS by definition may shift part or all of its responsibility to another node in the network segment.

3. Enter any additional addresses up to a maximum of 8. Press **<Enter>** to return to the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

To *delete* next-hop addresses for the NBDD relay:

1. Enter **32=d** at the command prompt of the UDP Relay Configuration screen. The FORWARD TO Server List displays.
2. Enter the item number that corresponds to the entry that you want to delete. Repeat this step to delete any additional entries.
3. Press **<Enter>** to return to the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

Forwarding VLANs for NBDD Relay

You may select or deselect VLANs to which the NBDD relay will forward. At least one forwarding VLAN (or a next-hop address as described in *Next-Hop Addresses for NBDD* on page 26-17) must be configured for the relay.

To *select* forwarding VLANs for NBDD relay:

1. On the command line of the UDP Relay Configuration screen, enter the following:

32=a

A screen similar to the following displays:

Available/Selected VLANs

Item	Group ID:VLAN ID	MASK	IP ADDR	
1)	1:1	255.255. 0. 0	172. 23. 9.105	*

* = selected for forwarding

Enter item number of VLAN to be selected [**h** or help/<ret> to exit] :

2. Enter the item number of the group/VLAN that you want to select. Repeat this step for all the groups/VLANs you want to select.
3. Press <Enter> to return to the command line for the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

To *deselect* forwarding VLANs:

1. On the UDP Relay Configuration screen, enter

32=d

The Available/Selected VLANs screen displays. Asterisks indicate VLANs selected for forwarding.

2. Enter the item number of the group/VLAN that you want to deselect. Repeat this step for all the groups/VLANs you want to deselect.
3. Press <Enter> to return to the command line for the UDP Relay Configuration screen.
4. Enter **s** to save the changes.

Generic Service UDP Relay

UDP relay may be configured for generic services. Generic services may include applications such as Trivial File Transfer Protocol (TFTP), Domain Name System (DNS), IEN-116 Name Server. You will need to know the well-known UDP port number if you want to configure these services.

Generic Services Menu

To configure a relay for a generic service, on the command line for the UDP Relay Configuration screen, enter **4**. A menu similar to the following displays:

```
4) +Generic Services Menu
  41) +Modify existing Generic Services Menu
  42) +Delete existing Generic Service Menu
  43) +Add new Generic Service Menu
```

Submenu Command {Item/?/Help/Quit/Redraw} {Redraw} :

Adding a Generic Service

Use the Add new Generic Service Menu to create a new generic service. On the Generic Services Menu, enter **43**. A screen similar to the following displays:

```
43) +Add new Generic Service Menu
  431) Description of new Service      :
  432) Forwarded port                 : UNSET
  433) Next-hop Address {list/add/delete} : UNSET
  434) Forward to VLANs {list/add/delete} : UNSET
```

Command {Item/?/Help/Quit/Done/Redraw} {Redraw} :

The required parameters are Forwarded port, and *either* Next-hop Address *or* Forward to VLANs. A description of the generic service is optional.

The **Done** command on this screen saves the current changes but does not activate the relay. The relay will be reinitialized and activated with the changes when **Save** is entered on the UDP Relay Configuration screen.

Description of new Service

A description of the service you want to configure.

Forwarded port

The corresponding well-known UDP port number for the service. For example, TFTP uses port 69. The default is **UNSET**. When you set this parameter, the relevant port number displays.

Next-hop Address

Use this parameter to list, add, or delete the server address(es) to which the NBNS UDP relay will forward. The default is **UNSET**. Up to 8 addresses may be configured. The value redisplay as **SET** when at least one address is configured. You can list, add, or delete addresses by entering **433=l**, **433=a**, or **433=d** on the command line.

Forward to VLANs

Use this parameter to list, add, or delete VLANs to which the NBNS UDP relay will forward. This default is **UNSET**. This value redisplay as **SET** when at least one VLAN is configured. Up to 32 VLANs may be configured. You can list forwarding VLANs, or add or delete VLANs from the forwarding list by entering **434=l**, **434=a**, or **434=d** on the command line.

To configure a generic service:

1. On the Add new Generic Service menu, enter a description of the generic service. For example:

```
431=TFTP
```

2. Enter the relevant UDP port number. For example:

```
432=69
```

3. At least one next-hop address must be configured. To add an address, enter:

```
433=a
```

The screen displays similar to the following:

FORWARD TO Server List

Item	Server address	Server Name (if known)
------	----------------	------------------------

Enter IP address or host name of server to be added to list ['h' for help/<ret> to exit]:

4. Enter the next-hop address in dotted decimal format (i.e., 198.206.181.12), a hexadecimal address (i.e., 0xc6ceb501). A host name (i.e., system.com) may be entered if the DNS resolver is enabled using the **res** command.
5. When you are finished entering next-hop addresses, press **<Enter>** to return to the prompt for the Add new Generic Services menu.
6. Select any VLANs for the relay to forward to. At the prompt, enter

```
434=a
```

A screen similar to the following displays:

Available/Selected VLANs

Item	Group ID:VLAN ID	MASK	IP ADDR	
1)	1:1	255.255. 0. 0	172. 23. 9.105	*

* = selected for forwarding

Enter item number of VLAN to be selected ['h'f or help/<ret> to exit] :

7. Enter the item number of the group/VLAN that you want to select. Repeat this step for all the groups/VLANs you want to select. An asterisk displays next to all selected VLANs.
8. Press **<Enter>** to return to the Add new Generic Services menu. Add any other generic services in this way.
9. Enter **d** to keep the current changes and return to the Generic Services menu. Enter **d** to return to the UDP Relay Configuration screen.
10. Enter **s** to save the changes and reinitialize the relay.

Modifying a Generic Service

Use the Configured Generic Services screen to modify an existing generic service. On the Generic Services Menu, enter **41**. A screen similar to the following displays:

Configured Generic Services					
Item	State	Port Number	Description	Servers/Vlans	
(1)	enabled	80	TFTP	198.172. 5.	4

Enter item number of service to be modified ['h' for help/<ret> to exit] :

The parameters are defined here.

Item

A unique number assigned by the switch to the generic service in the order the services were configured using the Add new Generic Service screen.

State

The current state of the service, enabled or deleted. The service is enabled as soon as it is added using the Add new Generic Service screen.

Port Number

The well-known UDP number configured for the generic service on the Add new Generic Service screen.

Description

The description of the generic service configured on the Add new Generic Service screen.

Servers/Vlans

The servers or VLANs that the relay will forward to.

To modify an existing generic service:

1. On the Configured Generic Services screen, enter the item number of the relevant service. The Modify existing Generic Services Menu displays similar to the following:

```

41) +Modify existing Generic Service Menu
    411) Description of Service being modified : TFTP
    412) Forwarded port                       : 80
    413) Next-hop Address {list/add/delete}   : SET
    414) Forward to VLANs {list/add/delete}  : SET
  
```

Command {Item/?/Help/Quit/Done/Redraw} {Redraw} :

2. Modify any of the parameters in the same way you configured them (described in *Adding a Generic Service* on page 26-19).
3. Enter **d** to keep the current changes and return to the Generic Services Menu. (The relay will not be initialized with the changes until you save them on the UDP Relay Configuration screen.)

4. Enter **d** to return to the UDP Relay Configuration screen.
5. Enter **s** to save the changes and reinitialize the relay.

Deleting a Generic Service

To delete a generic service:

1. On the Generic Services Menu, enter **42**. The Configured Generic Services screen displays similar to the following:

Configured Generic Services				
Item	State	Port Number	Description	Servers/Vlans
(1)	enabled	80	TFTP	198.172. 5. 4

Enter item number of service to be deleted [**h** for help/<ret> to exit] :

The parameters are defined in *Modifying a Generic Service* on page 26-21.

2. Enter the item number of the service you want to delete. A message similar to the following displays:

Are you sure you want to delete item 1? [**y/n**] (n) :

3. Enter **y** to delete the service. The Configured Generic Services screen redisplay with the State parameter changed to **deleted**. At this point, the service is marked for deletion but has not actually been deleted from the configuration.

Configured Generic Services				
Item	State	Port Number	Description	Servers/Vlans
(1)	deleted	80	TFTP	198.172. 5. 4

Enter item number of service to be deleted [**h** for help/<ret> to exit] :

4. Select any other services to be marked for deletion. Press **<Enter>** to return to the Generic Services Menu.
5. Enter **q** to return to the UDP Relay Configuration screen.
6. Enter **s** to save the changes and delete the selected service(s).

Viewing UDP Relay Statistics

Use the **relays** command to display statistics about configured UDP relays. The **relays** command is listed in the IP submenu. For information about other IP commands, see Chapter 25, “IP Routing.”

The screen display for UDP statistics is similar to the following:

UDP RELAY PACKETS RECEIVED/TRANSMITTED						
SERVICE	PORT	PKTS RCVD	RCV RATE(pkts/s)	PKTS XMTD	XMT RATE(pkts/s)	
1	67/68	0	0.000	0	0.000	
2	137	6	0.010	0	0.000	

NOTE: Rates are average number of packets/s since last query.
Time since last query: 0 days, 0 hours, 10 minutes, 6 seconds.

UDP RELAY TRANSMIT PACKETS DISCARDED					
SERVICE	RVC PORT	DEST VLAN/SVR		PKTS	
1	67/68	172. 28.	5. 21	0	
1	67/68	198.172. 34.	2	0	
1	67/68	198.172. 34.	5	0	
2	137	172. 23.	9.105	6	
2	137	172. 28.	5.212	6	

The fields are defined here.

SERVICE. The number assigned by the switch to the UDP service, in order that the services were configured.

PORT. The well-known UDP port number associated with the type of service. For example, BOOTP/DHCP is 67/68. This number is manually configured for generic services.

PKTS RCVD. The total number of packets received by the relay for the indicated service.

RCV RATE(pkts/s). The average rate, in packets per second, that packets were received for the indicated service since the last time the **relays** command was entered.

PKTS XMTD. The total number of packets transmitted from the relay for the indicated service.

XMT RATE(pkts/s). The average rate, in packets per second, that packets were transmitted for the indicated service since the last time the **relays** command was entered.

RVC PORT. The UDP port number associated with the service

DEST VLAN/SVR. The IP address of the VLANs to which the indicated relay is forwarding. Forwarding VLANs are configurable for each type of relay.

PKTS. The number of packets forwarded to the indicated VLAN.


27 IPX Routing

Introduction

This chapter gives an overview of Internetwork Packet Exchange (IPX) routing and includes information about configuring static IPX routes as well as configuring Routing Information Protocol (RIP) and Service Advertising Protocol (SAP) filters and timers. IPX is a layer 3 protocol developed by Novell for interconnecting NetWare clients and servers. (NetWare is Novell's network server operating system.) IPX routing requires at least one IPX router port to be configured on the switch.

When IPX routing is enabled on the switch, the switch will be able to exchange routing information with IPX routers in the network, and stations connected to groups and VLANs with virtual IPX router ports will be able to communicate. Groups or VLANs that do not have IPX router ports with IPX routing enabled cannot communicate with each other.

In the example shown here, stations connected to each group will be able to communicate if a virtual IPX router port is created for each group and each router port on the switch has IP routing enabled. Stations in group 2 and group 3 will also be able to communicate with stations attached to the external IPX router if a static route to that router is configured on the switch or the switch learns about the external router through IPX RIP or SAP.



IPX Routing Overview

In IPX routing, the switch builds routing tables to keep track of optimal destinations for traffic it receives that is destined for remote IPX networks. The switch sends and receives routing messages, or advertisements, to/from other routers in the network. When the switch receives an IPX packet, it looks up the destination network number in its routing table. If the network is directly connected to the switch, the switch also checks the destination node address. The network number consists of eight hex digits, and the node address is typically the MAC address of the end station or server.

Creating routing tables is performed by switch software unless a Hardware Routing Engine (HRE) or HRE-X is installed. The HRE or HRE-X significantly improves routing performance. See Chapter 1, “Omni Switch/Router Chassis and Power Supplies,” for information about the HRE-X respectively.

IPX is associated with additional protocols built into the switch software. These are described in the next section.

IPX Protocols

The switch supports the following IPX protocols:

- **SPX** (Sequenced Packet Exchange) is a Transport-layer protocol that provides a reliable end-to-end communications link by managing packet sequencing and delivery. SPX does not play a direct role in IPX routing; it simply guarantees the delivery of routed packets.
- **IPX RIP** (Routing Information Protocol) is a layer 3 protocol used by NetWare routers to exchange IPX routing information. IPX RIP functions similarly to IP RIP. IPX RIP uses two metrics to calculate the best route: hop count and ticks. An IPX router periodically transmits packets containing the information currently in its own routing table to neighboring IPX RIP routers in order to advertise the best route to an IPX destination.
- **SAP** (Service Advertising Protocol) is a layer 3 protocol used by NetWare routers to exchange IPX routing information. SAP is similar in concept to IPX RIP. Just as RIP enables NetWare routers to exchange information about routes, SAP enables NetWare devices to exchange information about available network services. NetWare workstations use SAP to obtain the network addresses of NetWare servers. IPX routers use SAP to gather service information and then share it with other IPX routers.

Setting Up IPX Routing on the Switch

IPX routing is enabled on a per-port basis by creating a virtual IPX router port for a group/VLAN. The switch does not do any routing unless the virtual IPX router port has IPX routing enabled (routing is enabled by default). The steps for setting up IPX routing on the switch are given here:

Step 1. Configuring a Virtual Router Port

A virtual IPX router port may be created when you set up or modify a group/VLAN through the **crgrp** command or **modvl** command described in Chapter 19, “Managing Groups and Virtual Ports.” To create a virtual router port, you enable IPX routing and specify a network address for the router port.

◆ **Note** ◆

IP and IPX routing may be enabled on the same port.

IPX router ports on the switch must also be configured with a particular encapsulation type for Ethernet: 802.3, 802.2 or LLC, SNAP, or Ethernet II.

Step 2. Configuring Optional IPX Routing Parameters

Optional configuration for IPX routing includes the following:

- Static routes. These are routes that are manually added to the routing table and may be used rather than dynamic routes that are learned through RIP or SAP.
- IPX RIP and SAP filters. IPX RIP and SAP filters may be configured and displayed. The default timers for RIP and SAP may also be modified. Extended RIP and SAP packets may also be configured.

The IPX Submenu

The **ipx** command in the Networking menu is used to access a submenu containing all the IPX-related commands. For more information about the Networking menu, see Chapter 25, “IP Routing.”

To display the IPX submenu, enter the following commands:

```
IPX
?
```

If you have enabled the verbose mode, you don't need to enter the question mark (?).

A screen similar to the following displays:

Command	IPX Menu
ipxr	View IPX routes
ipxs	View IPX stats and errors
ipxsap	View IPX SAP bindery
aipxsr	Add an IPX static route
ripxsr	Remove an IPX static route
ipxoff	Turn off the IPX router complex
ipxon	Turn on the IPX router complex
ipxflush	Flush IPX router RIP and/or SAP tables
ipxping	IPX Ping a system
ipxfilter	Add/delete an IPX RIP/SAP filter
ipxf	Display IPX RIP/SAP filters
ipxserialf	Enable/Disable IPX Serialization Packet Filtering
ipxspooof	Enable/Disable IPX Watchdog Spoofing
spxspooof	Enable/Disable SPX Keepalive Spoofing
ipxtype20	Turn on/off forwarding of IPX Type 20 packets
ipxtimer	Add/Delete SAP and RIP timers
ipxt	Display SAP and RIP timers
ipxdrt	Turn on/off a default route for IPX
ipxext	Turn on/off extended IPX RIP and SAP packets

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

This chapter describes all of the above commands. The remaining sections of this chapter cover each of the above commands in the order in which they appear in the IPX submenu.

Viewing the IPX Routing Table

The `ipxr` command is used to display the IPX Routing Table. The entries in the table show the routes entered by the IPX RIP protocol and the static routes that you may have entered manually. All entries in the table are sorted by destination network. The IPX Routing Table can contain a maximum of 2,010 routes.

Displaying All Entries in the IPX Routing Table

To display all entries in the IPX Routing Table, enter the following command:

```
ipxr
```

A screen similar to the following displays:

Displaying all (4) routes:

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL
3333	e8024.0000c021a5b8	1	3	N	Y	N	N	N	7:1
5555	5555.Direct	0	1	N	N	N	N	Y	4:1
e8024	e8024.Direct	0	1	N	N	N	N	Y	7:1
3041c204	e8024.0000c021a5b8	1	3	N	Y	N	N	N	7:1

The fields on this screen have the following meanings:

Dest Net

The destination network IPX address.

Router

The IPX address (network.node) of the next hop router to reach the destination network.

Hops

The number of routers between this node and the destination network.

Delay

The number of “ticks” between this node and the destination network. A “tick” is about 1/18th of a second.

Static

Whether this route was statically defined (see the `aipxsr` command).

Aged

Indicates if this route has timed out. Once a route times out it is kept in the routing table for 10 “ticks.” Once the 10 “ticks” expire, the route is deleted.

Redir

Indicates that a route to an IPX network that was formerly reachable via a direct interface has been replaced by an alternate route.

Viewing the IPX Routing Table

Chg

The information in this route has recently been updated, but the new information has not yet been forwarded to neighbor routers.

Dir

Indicates that this is a local interface (direct route) as opposed to a route to a destination network.

GP:VL

The first number is the Group associated with this entry; the second number is the VLAN associated with this entry. This identifies the interface used when sending traffic to the destination network.

Using IPXR with Frame Relay or ISDN Boards

The following additional column heading appears in the `ipxr` display when a Frame Relay or ISDN board is installed in the switch:

s/p/vc or Peer ID

The Slot, Port and Virtual Connection (i.e., DLCI) identifiers or the PPP Peer ID of the interface on which the routing information was received.

Here is an example of a display generated by the `ipxr` command in this situation:

Displaying all (12) routes:

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL	s/p/vc Peer ID
100	100.Direct	0	1	N	N	N	N	Y	3:1	
120	120.Direct	0	1	N	N	N	N	Y	4:1	
5000	120.0020da092ef5	1	2	N	N	N	N	N	4:1	5/3/100
5556	8484.0020da2200f4	1	2	N	N	N	N	N	6:1	P1
8484	8484.Direct	0	1	N	N	N	N	Y	6:1	
26dc012a	120.0020da092ef5	1	2	N	N	N	N	N	4:1	5/3/220
55555555	120.0020da092ef5	1	2	N	N	N	N	N	4:1	5/3/100
66666666	66666666.Direct	0	1	N	N	N	N	Y	5:1	
95000095	120.0020da092ef5	2	3	N	N	N	N	N	4:1	5/3/100

In this example, traffic destined for Network 5000 will go through Slot 5, Port 3, DLCI 100 which is associated with the interface on Group 4.

Displaying a List of Specific IPX Routes

You can limit the number of routes that are displayed by the `ipxr` command by using an extra argument along with the command. To find out if a route to a particular destination network is known, simply include the network number on the command line. (The examples shown below came from a switch that contained a Frame Relay board and an ISDN board.)

Here is an example for destination network 5000 (the command used is: `ipxr 5000`):

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL	s/p/vc
5000	120.0020da092ef5	1	2	N	N	N	N	N	4:1	5/3/100

To display only those routes learned from a particular interface, you can specify the interface number on the command line. You can also further specify the slot/port/vc or PPP Peer ID.

This is an example for Interface 3:1 (the command used was: `ipxr 3:1`):

Displaying routes for interface 3:1

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL	s/p/vc
100	100.Direct	0	1	N	N	N	N	Y	3:1	

This is an example for Interface 4:1 5/3/100 (the command used was: `ipxr 4:1 5/3/100`):

Displaying routes for interface 4:1

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL	s/p/vc
5000	120.0020da092ef5	1	2	N	N	N	N	N	4:1	5/3/100
55555555	120.0020da092ef5	1	2	N	N	N	N	N	4:1	5/3/100
95000095	120.0020da092ef5	2	3	N	N	N	N	N	4:1	5/3/100

This is an example for Interface 6:1 P1 (the command used was: `ipxr 6:1 P1`):

Displaying routes for interface 6:1

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL	s/p/vc
5556	8484.0020da2200f4	1	2	N	N	N	N	N	6:1	P1

Viewing IPX Statistics

The **ipxs** command is used to display data on IPX statistics and errors.

To display information about IPX statistics and errors, enter the following command:

```
ipxs
```

A screen similar to the following displays:

IPX Statistics and Errors:

IPX is ON

IPX Input Statistics:

```
pkts rcvd           = 3280
pkts delivered locally = 3161
pkts discarded      = 0
input header errors  = 0
```

IPX Output Statistics:

```
pkts sent           = 4731
pkts generated locally = 4681
pkts discarded      = 0
pkts with no route found = 1
HRE pkts sent       = 0
```

There are 2 IPX interfaces defined.

Stats for IPX Router Interface on (Group:VLAN) 3:1, Net address 3333

Interface name is IPX Router 3333

```
state           = ON           status      = UP
state changes   = 1500         type        = BROADCAST
rtr encapsulation = FD
```

RIP is ON: sent = 1527, rcvd = 1568, update interval = 60 secs.

SAP is ON: sent = 1, rcvd = 1568, update interval = 60 secs.

Stats for IPX Router Interface on (Group:VLAN) 4:1, Net address 5555

Interface name is IPX Router 5555

```
state           = ON           status      = UP
state changes   = 1500         type        = BROADCAST
rtr encapsulation = EN
```

RIP is ON: sent = 1571, rcvd = 1, update interval = 60 secs.

SAP is ON: sent = 1533, rcvd = 1, update interval = 60 secs.

The fields (and the subfields) on this screen have the following meanings:

IPX

Indicates whether IPX routing is “ON” or “OFF.”

IPX Input Statistics

pkts rcvd: The number of packets received.

pkts delivered locally: The number of received packets delivered to local IPX applications (RIP and SAP).

pkts discarded: The number of discarded packets.

input header errors: The number of packets discarded due to IPX packet header errors.

IPX Output Statistics

pkts sent: The number of packets forwarded (not including fast path routed packets).

pkts generated locally: The number of packets forwarded that were generated by local IPX applications (RIP and SAP).

pkts discarded: The number of discarded packets.

pkts with no route found: The number of packets that could not be forwarded because a route to the destination IPX network could not be found.

Stats for IPX Router Interface

state: State of the IPX router for this interface (ON or OFF).

status: Status of the interface (UP or DOWN).

type: The type of interface (BROADCAST or POINT-TO-POINT).

rtr encapsulation: Router port encapsulation used for this interface (EN=Ethernet, FD=FDDI, TR=Token Ring).

state changes: The number of state changes that have occurred on this interface (up to down, down to up).

RIP

sent: The number of RIP packets sent.

received: The number of RIP packets received.

update interval: The RIP update timer interval for this interface. If a WAN interface is configured as a Triggered RIP/SAP interface, this field will contain the word "triggered." Triggered interfaces transmit information only once, when the change occurs.

SAP

sent: The number of SAP packets sent.

received: The number of SAP packets received.

update interval: The SAP update timer interval for this interface. If a WAN interface is configured as a Triggered RIP/SAP interface, this field will contain the word "triggered." Triggered interfaces transmit information only once, when the change occurs.

Viewing the IPX SAP Bindery

The **ipxsap** command is used to display a listing of the servers in the SAP Bindery, sorted by server name.

To display a list of SAP servers, enter the following command:

```
ipxsap
```

A screen similar to the following displays:

Displaying all (3) entries in the SAP bindery:

Server Name	Type	Address	Hp	Sckt	GP:VL
Develop	0004	67.000000000001	1	0451	3:1
Finance	026b	67.000000000001	1	0005	2:1
Marketing	0278	67.000000000001	1	4006	2:1

The fields on this screen have the following meanings:

Server Name

The name of the server offering this service.

Type

The service type being offered (as defined by Novell).

Address

The IPX address of this server (network.node).

Hp

The number of networks between this node and the server.

Sckt

The Novell socket number to which this service is attached.

GP:VL

The first number is the Group associated with this entry, and the second number is the VLAN associated with this entry.

Using IPXSAP with Frame Relay or ISDN Boards

The following additional column heading appears in the **ipxsap** display when a Frame Relay or ISDN board is installed in the switch.

s/p/vc or Peer ID

The Slot, Port and Virtual Connection (i.e., DLCI) identifiers or the PPP Peer ID of the interface on which the server information was received.

Here is an example of a display generated by the **ipxsap** command in this situation:

Displaying all (3) entries in the SAP bindery:

Server Name	Type	Address	Hp	Sckt	GP:VL	s/p/vc Peer ID
HR	0004	200.000000000022	1	0451	3:1	5/3/100
Sales	026b	200.000000000022	1	0005	2:1	5/3/220
Support	0278	200.000000000022	1	4006	2:1	5/3/220

Displaying a List of Specific SAP Servers

You can limit the number of SAP server names that is displayed by the **ipxsap** command by using an extra argument with the command.

To display only those servers from a specific interface, simply include the interface number on the command line. The following is an example for Interface 2:1 (the command used was **ipxsap 2:1**):

Displaying all SAPs for interface 2:1:

Server Name	Type	Address	Hp	Sckt	GP:VL
Finance	026b	67.000000000001	1	0005	2:1
Marketing	0278	67.000000000001	1	4006	2:1

To display a specific type of server, simply include a Server Type value (in hex) on the command line. The following is an example for 26b (the command used was **ipxsap 26b**):

Displaying SAP entries of type 0x26b:

Server Name	Type	Address	Hp	Sckt	GP:VL
Finance	026b	67.000000000001	1	0005	2:1

To find out if a particular server is known, simply include all, or just a portion of, the server name on the command line. The server name (or portion thereof) must be entered inside of quotation marks. The following is an example for an entry of "nance" (the command used was **ipxsap "nance"**):

Displaying SAP entries whose names contain the substring "nance":

Server Name	Type	Address	Hp	Sckt	GP:VL
Finance	026b	67.000000000001	1	0005	2:1

Adding an IPX Static Route

The **aipxsr** command is used to add IPX static routes to the switch's IPX Routing Table. You might want to add a static route to send traffic from a node in an Omni Switch/Router VLAN to an external IPX network address (such as an address reached through an external network router attached to the switch).

In order to add a static route, you will need to know the host/net and the gateway which will be used to route traffic there.

Follow the steps below to add an IPX static route.

1. Enter **aipxsr**.

A screen similar to the following displays:

Do you want to see the current route table? (y or n) (y) :

2. Enter **y** at this prompt (or press **<Enter>**) to display the current routing table.

A screen similar to the following displays:

Displaying all (4) routes:

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL
3333	e8024.0000c021a5b8	1	3	N	Y	N	N	N	7:1
5555	5555.Direct	0	1	N	N	N	N	Y	4:1
e8024	e8024.Direct	0	1	N	N	N	N	Y	7:1
3041c204	e8024.0000c021a5b8	1	3	N	Y	N	N	N	7:1

Destination IPX network :

Enter the IPX address of the network to which you are setting up a route.

3. The following prompt displays:

IPX network of next hop :

Enter the IPX network address of the next hop. This is the number that appears before the dot under the "Router" heading in the IPX Route Table.

4. The following prompt displays:

IPX node address of next hop (format - xx:xx:xx:xx:xx:xx)

Enter the IPX node address of the next hop.

5. A message will confirm the addition of the static route:

Route successfully added

Removing an IPX Static Route

The **ripksr** command is used to remove IPX static routes from the switch's IPX Routing Table. Follow the steps below to remove an IPX static route.

1. Enter **ripksr**.

A screen similar to the following displays:

```
Do you want to see the current route table?
(y or n) (y) : y
```

2. Enter **y** at this prompt (or press **<Enter>**) to display the current routing table.

A screen similar to the following displays:

```
Displaying all (4) routes:
```

Dest Net	Router	Hops	Delay	Static	Aged	Redir	Chg	Dir	GP:VL
3333	e8024.0000c021a5b8	1	3	N	Y	N	N	N	7:1
5555	5555.Direct	0	1	N	N	N	N	Y	4:1
e8024	e8024.Direct	0	1	N	N	N	N	Y	7:1
3041c204	e8024.0000c021a5b8	1	3	N	Y	N	N	N	7:1
aaaaaa	304.0020da05f694	1	1	Y	N	N	Y	N	7:1

```
Destination IPX network :
```

3. Enter the name of the destination IPX network you want to remove.

A message will confirm the deletion of the static route:

```
Route successfully deleted.
```

Turning the IPX Router Complex On and Off

The **ipxoff** command is used to turn off the IPX Router Complex, which disables IPX routing on the switch.

To turn off IPX routing, enter the following command:

ipxoff

A screen similar to the following displays:

IPX turned off.

The **ipxon** command is used to turn on the IPX Router Complex, which enables IPX routing on the switch.

To turn on IPX routing, enter the following command:

ipxon

A screen similar to the following displays:

IPX turned on.

Flushing the IPX RIP/SAP Tables

The **ipxflush** command is used to flush the IPX RIP Routing and SAP Bindery Tables.

Follow the steps below to flush both the IPX tables.

1. Enter **ipxflush**.

A screen similar to the following displays:

```
Flush tables (RIP routing and SAP bindery) in:  
{ RIP and SAP(b),  
  RIP only(r),  
  SAP only(s)} (b) :
```

2. Enter **b** (or just press Enter) to flush both tables. Enter **r** to flush just the Routing Table.
Enter **s** to flush just the SAP Bindery Table.

You will be returned to the system prompt.

Using the IPXPING Command

The **ipxping** command is used to test the reachability of certain types of IPX nodes. The software supports two different types of IPX pings:

- Novell-defined, which can test the reachability of NetWare servers currently running the NetWare Loadable Module called IPXRTR.NLM. This type *cannot* be used to reach NetWare workstations running IPXODI. Novell uses a unique type of ping for this purpose (implemented by their IPXPNG.EXE program) which is not currently supported by the switch software. Other vendors' switches may respond to this type of ping.
- Alcatel-proprietary, which can test the reachability of Omni Switch/Routers on which IPX routing has been enabled.

Network devices that do not recognize the specific type of IPX ping request sent from the switch will not respond at all. The lack of a response does not necessarily mean that a specific network device is inactive or missing. Therefore, you might want to try using both types before concluding that the network device is "unreachable."

◆ Note ◆

The **ipxping** command does not work over FDDI trunking with Token Ring SNAP or LLC encapsulation. It does work with Token Ring SNAP or LLC encapsulation over other media types.

Follow the steps below to issue an IPX ping request.

1. Enter **ipxping**.

A screen similar to the following displays:

Dest Net () : 304

Enter the Destination Network of the node that you want to ping.

2. The following prompt displays:

Dest Node (format - xx:xx:xx:xx:xx:xx) () : 00:20:da:05:f6:94

Enter the Destination Node that you want to ping.

◆ Note ◆

If you are attempting to ping an interface that is specified with a noncanonical address, you must specify a noncanonical address for the ping.

3. The following prompt displays:

Count (0 for infinite) (1) : 245

Enter a number to indicate the number of packets to be sent out. An entry of 0 (zero) will create an infinite count (press **<Enter>** to cancel). The default count is 1 (one).

4. The following prompt displays:

Size (64) :

Enter a number to indicate the number of data bytes included in the packet. The default size is 64.

5. The following prompt displays:

Timeout (1) :

Enter the number of seconds to wait for a response. The default timeout is 1.

6. The following prompt displays:

Type (n for Novell, x for Xylan) (n) :

Enter the type of IPX ping to be issued. The default is the Novell type.

7. After answering the previous prompt, a message similar to the following displays:

**IPX Ping starting, hit <RETURN> to stop
PING 304.00:20:da:05:f6:94: 64 data bytes**

```
[0      ] .....  
[50     ] .....  
[100    ] .....  
[150    ] .....  
[200    ] .....
```

**---304.00:20:da:05:f6:94 IPXPING Statistics---
245 packets transmitted, 245 packets received, 0% packet loss**

You may also elect to bypass the above prompts. To do so, simply include the options on the command line in the exact order in which they appear in the prompts. You will be prompted for any options you leave out. Therefore, the syntax for the command is:

ipxping [destnet] [destnode] [count] [size] [timeout] [type]

For example, the following command string will send 100 Novell-type pings, using 64 data bytes per packet with a timeout of 1 second, to an IPX server with MAC address of 00:00:c0:21:a5:b8 on IPX network e8024:

ipxping e8024 00:00:c0:21:a5:b8 100 64 1 n

Configuring IPX RIP/SAP Filtering

The `ipxfilter` command is used to add or delete an IPX RIP or SAP Output or Input filter. The IPX RIP/SAP Filtering feature give you a means of controlling the operation of the IPX RIP/SAP protocols. By using IPX RIP/SAP filters, you can minimize the number of entries put in the IPX RIP Routing and SAP Bindery Tables, improve overall network performance by eliminating unnecessary traffic, and control users' access to NetWare services.

Five types of IPX RIP/SAP filters are available:

1. **RIP Input** filters control which networks are allowed into the routing table when IPX RIPs are received.
2. **RIP Output** filters control the list of networks included in routing updates sent out an interface. These filters control which networks the router advertises in its IPX RIP updates.
3. **SAP Input** filters control the SAPs received by the router prior to a router accepting information about a service. The router will filter all incoming service advertisements received before accepting information about a service.
4. **SAP Output** filters control which services are included in SAP updates sent by the router. The router applies the SAP output filters prior to sending SAP packets.
5. **GNS Output** filters control which servers are included in the GNS responses sent by the router.

Here are some example uses of IPX RIP/SAP filters:

- RIP Input and Output filters can be used to isolate entire network segments (and/or routers) in order to make the network appear differently to the different segments.
- RIP Input and Output filters can be used to reduce the amount of WAN traffic needed to advertise routes that shouldn't be used by a particular network segment.
- SAP Input and Output filters can be used to improve the performance of IPX in a WAN environment by limiting the amount of SAP traffic. For example, because printing is generally a local operation, there's no need to advertise print servers to remote networks. A SAP filter can be used in this case to restrict "Print Server Advertisement" SAPs.

◆ Important Note ◆

All types of IPX Filters can be configured either to *allow* or to *block* traffic. The default setting for all filters is to allow traffic. Therefore, you will typically only have to define a filter to block traffic. However, defining a filter to allow certain traffic may be useful in situations where a more generic filter has been defined to block the majority of the traffic. For example, you could use a filter to allow traffic from a specific host on a network where all other traffic has been blocked. A discussion of the precedence of "Allow" filters appears later in this section. Keep in mind that precedence applies only to "allow" filters, *not* to "block" filters.

You can apply filters to *all* router interfaces by defining a "global" filter, or you can limit the filter to *specific* interfaces. In addition, for WAN networks, you can apply filters to a specific Frame Relay virtual circuit (DLCI) or PPP Peer. Each of these options is described under individual heading in this section.

Adding a “Global” IPX RIP/SAP Filter

Follow the steps below to add a “global” IPX RIP or SAP filter.

1. Enter **ipxfilter**.

A screen similar to the following displays:

Selecting global IPX filter:

Add or delete entry {add(a), delete(d)} (a) :

Enter **a** (or just press **<Enter>**) to select to add a filter.

2. The following prompt displays:

**Filter type {SAP output(so),
SAP input(si),
RIP Output(ro),
RIP Input(ri),
GNS output(go)} (so) :**

Enter **so** (or just press **<Enter>**) to add a SAP Output filter. Enter **si** to add a SAP Input filter. Enter **ro** to add a RIP Output filter. Enter **ri** to add a RIP Input filter. Enter **go** to add a GNS Output filter.

3. The following prompt displays:

Filter action {block(b), allow(a)} (a) :

Enter **a** (or press **<Enter>**) to define the filter to allow traffic. Enter **b** to define the filter to block traffic.

4. The following prompt displays:

IPX network (default: all networks):

Enter the IPX network address (in hexadecimal format) that is to be used (or press **<Enter>** to use the default of “all networks”).

5. The following prompt displays:

IPX network mask (default: FFFFFFFF) :

Enter the IPX network mask (in hexadecimal format) to be used (or press **<Enter>** to use the default mask of FFFFFFFF). *If you selected the default of “all networks” in the previous step, this step is skipped.*

6. The following prompt displays:

IPX node address (default: all nodes):

Enter the IPX node address (in hexadecimal format) to be used (or press **<Enter>** to use the default of “all nodes”).

7. The following prompt displays:

IPX node mask (default: all F's) :

Enter the IPX node mask (in hexadecimal format) to be used (or just press **<Enter>** to use the default mask of all F's). *If you selected the default of “all nodes” in the previous step, this step is skipped.*

- The following prompt displays:

SAP service type (default: all services) :

Enter the SAP service type (in hexadecimal format) as defined by NetWare (or press **<Enter>** to use the default of all services).

- A message will confirm the addition of the filter:

ipxfilter successfully added

Adding an IPX RIP/SAP Filter for a Specific Group or VLAN

Follow the steps below to add an IPX RIP or SAP Output or Input filter for a specific Group or VLAN.

- Enter the Group and VLAN numbers after the command like this: **ipxfilter 1:1**.

A screen similar to the following displays:

Selecting IPX filter for interface 1:1:

Add or delete entry {add(a), delete(d)} (a) :

Enter **a** (or press **<Enter>**) to select to add a filter.

- The following prompt displays:

**Filter type {SAP output(so),
SAP input(si),
RIP Output(ro),
RIP Input(ri),
GNS output(go)} (so) :**

Enter **so** (or press **<Enter>**) to add a SAP Output filter. Enter **si** to add a SAP Input filter. Enter **ro** to add a RIP Output filter. Enter **ri** to add a RIP Input filter. Enter **go** to add a GNS Output filter.

- The following prompt displays:

Filter action {block(b), allow(a)} (a) :

Enter **a** (or press **<Enter>**) to define the filter to allow traffic. Enter **b** to define the filter to block traffic.

- The following prompt displays:

IPX network (default: all networks):

Enter the IPX network address (in hexadecimal format) that is to be used (or press **<Enter>** to use the default of "all networks").

- The following prompt displays:

IPX network mask (default: FFFFFFFF) :

Enter the IPX network mask (in hexadecimal format) to be used (or just press **<Enter>** to use the default mask of FFFFFFFF). *If you selected the default of "all networks" in the previous step, this step is skipped.*

- The following prompt displays:

IPX node address (default: all nodes):

Enter the IPX node address (in hexadecimal format) to be used (or just press **<Enter>** to use the default of "all nodes").

- The following prompt displays:

IPX node mask (default: all F's) :

Enter the IPX node mask (in hexadecimal format) to be used (or just press **<Enter>** to use the default mask of all F's). *If you selected the default of "all nodes" in the previous step, this step is skipped.*

- The following prompt displays:

SAP service type (default: all services) :

Enter the SAP service type (in hexadecimal format) as defined by NetWare (or just press **<Enter>** to use the default of all services).

- A message will confirm the addition of the filter:

ipxfilter successfully added

Using Filters with Frame Relay or ISDN Boards

If the Group or VLAN you enter (such as 1:1 used in the above example) refers to a WAN interface like Frame Relay or PPP, you'll be asked if you want the filter applied to a specific WAN endpoint.

- This prompt appears after the previous prompt for "SAP Service Type":

Do you wish to apply this filter to a specific WAN endpoint? (n):

Enter **y** to select to apply this filter to a specific WAN endpoint.

- The following prompt displays:

Frame Relay VC or PPP Peer {vc(v), peer(p)} (v):

Enter **v** (or just press **<Enter>**) to apply this filter to a Frame Relay Virtual Circuit. Proceed to the next step.

Enter **p** if you want to apply this filter to a PPP Peer. Skip to the last step.

- If you chose to apply a filter to a Frame Relay VC, this prompt displays:

Slot/port:

Enter the slot and port to which you want to apply this filter (for example, **3/1**).

Enter the VC to which you want to apply this filter.

- If you chose to apply a filter to a PPP Peer, this prompt displays:

Peer ID:1

Enter the Peer ID to which you want to apply this filter (for example, **1**).

Deleting an IPX RIP/SAP Filter

Follow the steps below to delete an existing IPX RIP or SAP filter.

1. Enter `ipxfilter`.

A screen similar to the following displays:

Selecting global IPX filter:

Add or delete entry {add(a), delete(d)} (a) :

Enter `d` to select to delete a filter.

2. A screen similar to the following displays:

Displaying all filters:

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL (s/p/vc) (Peer ID)
1	SAP OUT	67/ffffff	00000000001/ffffffff	ALL	B	global
2	SAP IN	67/ffffff	00000000001/ffffffff	0278	B	1:1
3	RIP IN	67/ffffff			B	global

Entry number to delete? (default: none) : 1

This screen contains a list of the existing IPX RIP/SAP filters. The fields on this screen are described in the next section (see *Displaying IPX RIP/SAP Filters* on page 27-23).

3. Enter the index number of the filter you want to delete. If you decide at this point that you want to abort out of the deletion process, simply press `<Enter>` to accept the default of "none."
4. A message will confirm the deletion of the filter:

ipxfilter successfully deleted.

Displaying IPX RIP/SAP Filters

The **ipxf** command is used to display a list of all existing IPX RIP and SAP filters. See *Adding a "Global" IPX RIP/SAP Filter* on page 27-19 for complete information on creating these filters. You can enter optional parameters with the **ipxf** command to display specific filters.

Displaying a List of All IPX Filters

To display a listing of all existing IPX RIP and SAP filters, enter the following command:

```
ipxf
```

A screen similar to the following displays:

Displaying all filters:

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL (s/p/vc) (Peer ID)
1	SAP OUT	67/ffffff	00000000001/ffffffff	ALL	B	global
2	SAP IN	67/ffffff	00000000001/ffffffff	0278	B	1:1
3	RIP IN	67/ffffff			B	global
4	SAP IN	All Networks	All Nodes	ALL	B	3:1 (P1)

This screen contains a list of the existing IPX RIP and SAP filters. The fields on this screen have the following meanings.

#

The index number assigned to identify each filter.

Type

The type of filter. The five types are: RIP IN, RIP OUT, SAP IN, SAP OUT, and GNS OUT.

Net/Mask

The IPX network address to be filtered ("All networks" means all networks are filtered).

Node/Mask

The IPX node address to be filtered ("All nodes" means all nodes are filtered). This field does not apply to RIP IN or RIP OUT filters.

Svc

The SAP service type (shown as a hexadecimal number) on which the filter is applied, as defined by Novell. By default, all services will be filtered. (Note: This field does not apply to RIP IN or RIP OUT filters.)

Md

The Mode of operation for the filter: A to Allow, B to Block.

GP:VL (s/p/vc) or (Peer ID)

The first number (**GP**) is the Group associated with this entry. The second number (**VL**) is the VLAN associated with this entry. When a filter applies to all interfaces, this field will say "global." If an entry refers to a Frame Relay interface, column headings for slot, port, and virtual circuit (**s/p/vc**) may be displayed when the filter is applied to a particular virtual circuit rather than to the entire VLAN. If an entry refers to a PPP interface, the Peer ID (**Peer ID**) may be displayed when the filter is applied to a particular PPP Peer.

Displaying a List of "Global" IPX Filters

To display a listing of just the global IPX filters, enter the following command:

```
ipxf global
```

A screen similar to the following displays:

Displaying global filters:

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL (s/p/vc) (Peer ID)
1	SAP OUT	67/ffffff	000000000001/ffffffff	ALL	B	global
3	RIP IN	67/ffffff			B	global

Displaying a List of Specific IPX Filters

To display a listing of IPX RIP or SAP filters for a specific interface, you can specify other parameters along with the **ipxf** command. The format for the command in this case is:

```
ipxf <type> <GP:VL>
```

The type is one of these codes:

ri	for RIP INput
ro	for RIP OUTput
si	for SAP INput
so	for SAP OUTput
go	for GNS OUTput

For example, to display a list of the filters defined for Group 1, VLAN 1, you would enter:

```
ipxf 1:1
```

A screen similar to the following displays:

Displaying filters for interface 1:1:

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL (s/p/vc) (Peer ID)
2	SAP IN	67/ffffff	000000000001/ffffffff	0278	B	1:1

As another example, to display a list of all global RIP Input filters, you would enter:

```
ipxf ri global
```

A screen similar to the following displays:

Displaying all global RIP INPUT filters:

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL (s/p/vc) (Peer ID)
3	RIP IN	67/ffffff			B	global

IPX RIP/SAP Filter Precedence

Whenever you use multiple “allow” filters you must first define a filter to block all RIPs or SAPs. Then, all of the seceding “allow” filters of the same type must be *at least* as specific in all areas in order for the filters to work. Note that filtering precedence is related only to “allow” filters. Multiple “block” filters can be defined with varying specificity in each of the areas of the filter. The filtering done by the configurable parameters (Net/Mask, Node/Mask, Service/Mode) in the “allow” filter must be at least as specific as the filtering defined in the “block” filter.

As an example, consider a switch that knows of multiple Type 4 SAPs on various networks, including a network with an address of “40.” The switch also knows of various types of SAPs on Network 40. For this example, you want to block all SAPs coming from Network 40, but you want to allow all Type 4 SAPs, including the ones that come from Network 40.

To meet these objectives, you must configure the filters like this:

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL
1	SAP IN	40/ffffff	all nodes	ALL	B	global
2	SAP IN	40/ffffff	all nodes	4	A	global

The filters shown below will *not* work for our example because in Filter 2 the type of service is *less* specific than the type defined in Filter 1. All Type 4 SAPs will be blocked by the filter.

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL
1	SAP IN	All networks	all nodes	4	B	global
2	SAP IN	40/ffffff	all nodes	ALL	A	global

The following filters will also *not* work because in Filter 2 the network and netmask are *less* specific than the network and netmask defined in Filter 1. All SAPs from Network 40 will be blocked by the filter.

#	Type	Net/Mask	Node/Mask	Svc	Md	GP:VL
1	SAP IN	40/ffffff	all nodes	ALL	B	global
2	SAP IN	All networks	all nodes	4	A	global

Configuring IPX Serialization Packet Filtering

The **ipxserialf** command is used to enable and disable IPX Serialization Packet filtering on any or all WAN routing services. This feature can be used to reduce traffic on WAN links by preventing the transmission of NetWare serialization packets.

Novell uses a serialization mechanism to make sure that licensed copies of NetWare are not improperly copied to multiple servers. NetWare's built-in copy protection scheme transmits serialization packets between file servers which contain unique serialization numbers. These packets are sent out at about 66-second intervals. If a server detects duplicate serialization identifiers, it broadcasts a copyright violation message to all users and to the console log. The major problem with this protection scheme for dial-on-demand links, such as ISDN, is the generation of traffic that continuously reactivates the WAN link.

Enabling IPX Serialization Filtering

Follow the steps below to enable IPX Serialization Packet Filtering.

1. Enter **ipxserialf**.

A screen similar to the following displays:

View the current status of all WAN routing services? (y or n) (n) :

Enter **y** if you do want to see the current filtering status. Enter **n** (or press **<Enter>**) if you entered this command by mistake or if you don't need to see the current status.

2. A screen similar to the following displays:

<u>Group</u>	<u>IPX Serialization Filtering</u>
3	Disabled
4	Disabled

Enter Group (default: all WAN) :

This screen shows the WAN routing Groups that exist in the switch and the current status of the IPX Serialization packet filtering for these groups.

Enter a Group number to proceed to enable IPX Serialization filtering for that Group.

Or, press **<Enter>** to select to enable filtering for *all* WAN routing services.

3. The following prompt displays:

Enable IPX Serialization Filtering? (y or n) (n) :

Enter **y** to select to enable IPX Serialization Filtering.

Enter **n** (or press **<Enter>**) if you do *not* want to enable Serialization Filtering.

4. The following prompt displays:

Enable IPX Serialization Filtering on all WAN routing services? (y or n) (n) :

This prompt requires you to verify that you want to enable filtering in order to avoid the situation of accidental filtering of IPX Serialization packets. This example prompt asks if you want to disable filtering on all WAN routing services. If you had entered a specific Group number, the prompt would refer to that particular Group.

Enter **y** to enable IPX Serialization Filtering.

5. Filtering will then become active. A message will appear indicating that IPX Serialization Filtering is enabled, either on all WAN routing services or for a specific Group:

IPX Serialization Filtering is now enabled on all WAN routing services

Disabling IPX Serialization Filtering

Follow the steps below to disable IPX Serialization Packet Filtering.

1. Enter **ipxserialf**.

A screen similar to the following displays:

View the current status of all WAN routing services? (y or n) (n) :

Enter **y** if you do want to see the current filtering status. Enter **n** (or press **<Enter>**) if you entered this command by mistake or if you don't need to see the current status.

2. A screen similar to the following displays:

Group	IPX Serialization Filtering
3	Enabled
4	Enabled

Enter Group (default: all WAN) :

This screen shows the WAN routing Groups that exist in the switch and the current status of the IPX Serialization packet filtering for these groups.

Enter a Group number to proceed to disable IPX Serialization filtering for that Group.

Or, just press **Enter** to select to proceed to disable filtering for *all* WAN routing services.

3. The following prompt displays:

Enable IPX Serialization Filtering? (y or n) (n) :

Enter **n** (or press **<Enter>**) to select to disable Serialization Filtering.

4. The following prompt displays:

Disable IPX Serialization Filtering on all WAN routing services? (y or n) (n) :

This prompt requires you to verify that you want to disable filtering. This example prompt asks if you want to disable filtering on all WAN routing services. If you had entered a specific Group number, the prompt would refer to that particular Group.

Enter **y** to disable IPX Serialization Filtering.

5. A message will appear indicating that IPX Serialization Filtering is disabled, either on all WAN routing services or for a specific Group:

IPX Serialization Filtering is now disabled on all WAN routing services

Configuring IPX Watchdog Spoofing

The **ipxspoo** command is used to enable and disable IPX Watchdog Spoofing on any or all WAN routing services. The use of this feature is explained below:

Novell's IPX Watchdog Protocol, which is used by NetWare to maintain network node and server connections, can consume significant network bandwidth and thereby incur costs on expensive dial-on-demand, pay-per-packet WAN links. The Omni Switch/Router provides an IPX Watchdog Spoofing feature to prevent Watchdog packets from initiating connections on WAN links in situations where no other data is ready to be transferred.

The IPX Watchdog Spoofing feature enables the switch to respond to a NetWare server's Watchdog "Query" requests on behalf of a remote client, thus spoofing the requests. The spoofing action occurs when the switch "sees" an incoming Watchdog packet destined for an interface on which spoofing has been enabled. The switch responds to the server by sending out a valid Watchdog response. Spoofing thus maintains the required Watchdog function while avoiding the cost of making and maintaining a WAN link.

In some situations, the use of the IPX Watchdog Spoofing feature can make a NetWare server "believe" that an inactive session is still active. This occurrence can cause connectivity problems by denying login rights to legitimate users. Therefore, if you use the spoofing feature on networks that also limit the number of IPX or SPX sessions, you should utilize NetWare's "auto-logoff" function to minimize inappropriate denials of legitimate logins.

Enabling IPX Watchdog Spoofing

Follow the steps below to enable IPX Watchdog Spoofing.

1. Enter **ipxspoo**.

A screen similar to the following displays:

View the current spoofing status of all WAN routing services? (y or n) (n) :

Enter **y** if you do want to see the current IPX spoofing status. Enter **n** (or just press **Enter**) if you entered this command by mistake or if you don't need to see the current status.

2. A screen similar to the following displays:

Group	IPX Spoofing
3	Disabled
4	Disabled

Enter Group (default: all WAN) :

Enter a Group number to proceed to enable IPX spoofing for that particular Group.

Or, just press **Enter** to proceed to enable IPX spoofing for *all* WAN routing services.

3. The following prompt displays:

Enable Spoofing? (y or n) (n) :

Enter **y** to proceed to enable IPX spoofing.

4. The following prompt displays:

Enable IPX Spoofing on all WAN routing services? (y or n) (n) : y

This prompt requires you to verify that you want to enable spoofing in order to avoid the situation of accidental spoofing of IPX packets.

This example prompt asks if you want to enable spoofing on all WAN routing services. If you had entered a specific Group number, the prompt would refer to that particular Group.

Enter **y** to enable IPX Watchdog Spoofing.

- IPX Spoofing will then become active. A message will appear indicating that IPX Watchdog Spoofing is enabled, either on all WAN routing services, or for a specific Group:

IPX Spoofing is now enabled on all WAN routing services

Disabling IPX Watchdog Spoofing

Follow the steps below to disable IPX Watchdog Spoofing.

- Enter **ipxspoo**.

A screen similar to the following displays:

View the current spoofing status of all WAN routing services? (y or n) (n) :

Enter **y** if you do want to see the current IPX spoofing status. Enter **n** (or just press **<Enter>**) if you entered this command by mistake or if you don't need to see the current status.

- A screen similar to the following displays:

Group	IPX Spoofing
3	Enabled
4	Enabled

Enter Group (default: all WAN) :

Enter a Group number if you want to disable IPX spoofing for that particular Group.

Or, press **<Enter>** to disable IPX spoofing for *all* WAN routing services.

- The following prompt displays:

Enable Spoofing? (y or n) (n) :

Enter **n** (or just press **<Enter>**) to proceed to disable IPX spoofing.

- The following prompt displays:

Disable IPX Spoofing on all WAN routing services? (y or n) (n) : y

This prompt requires you to verify that you want to disable spoofing. This example prompt asks if you want to disable spoofing on all WAN routing services. If you had entered a specific Group number, the prompt would refer to that particular Group.

Enter **y** to disable IPX Watchdog Spoofing.

- IPX Spoofing will then become inactive. A message will appear indicating that IPX Watchdog Spoofing is disabled, either on all WAN routing services, or for a specific Group:

IPX Spoofing is now disabled on all WAN routing services

Configuring SPX Keepalive Spoofing

The `spxspooof` command is used to enable and disable SPX Keepalive Spoofing on any or all WAN routing services. The use of this feature is explained below:

Novell's SPX Keepalive Protocol, which is used by NetWare to maintain SPX connections between end nodes, can also consume significant network bandwidth and thereby incur unnecessary costs on expensive dial-on-demand, pay-per-packet WAN links. The Omni Switch/Router provides a SPX Keepalive Spoofing feature to prevent keepalive packets from keeping WAN links active when they are not otherwise needed for data transmissions.

The SPX Spoofing feature enables the switch to respond to client/server keepalive packets on the behalf of the remote clients/servers. SPX spoofing thereby effectively stops keepalive packets from crossing a WAN link while maintaining existing SPX connections.

SPX-Packet Tolerance Counting

NetWare's SPX and SPXII watchdog and keepalive packets unfortunately are not labeled with a unique packet type. Therefore, valid acknowledge packets or window-update packets could be mistaken for keepalive packets. To prevent blocking of critical packets, a packet tolerance counting mechanism is employed by the Spoofing feature to count SPX packets.

When active, the Spoofing feature observes all watchdog and keepalive packets as they go between network endpoints. If successive packets are found to have the same sequence number, acknowledge number, and "alloc" number, spoofing will not begin until the specified SPX-packet tolerance count has been reached. Only watchdog packets which have the ACK_REQUESTED bit set will have an effect on the SPX-packet tolerance counter.

Once the specified tolerance count has been reached, spoofing of watchdog packets will begin and all keepalive packets will be dropped. Refer to *Controlling IPX Type 20 Packet Forwarding* on page 27-32 for help on using NetWare's configurable parameters to change the frequency and number of keepalive/watchdog packets sent.

Enabling SPX Keepalive Spoofing

Follow the steps below to enable SPX Keepalive Spoofing.

1. Enter `spxspooof`.

A screen similar to the following displays:

View the current spoofing status of all WAN routing services? (y or n) (n) :

Enter **y** if you do want to see the current SPX spoofing status. Enter **n** (or press **<Enter>**) if you entered this command by mistake or if you don't need to see the current status.

2. A screen similar to the following displays:

<u>Group</u>	<u>SPX Spoofing</u>
3	Disabled
4	Disabled

Enter Group (default: all WAN) :

Enter a Group number to proceed to enable SPX spoofing for that particular Group.

Or, press **<Enter>** to proceed to enable SPX spoofing for *all* WAN routing services.

3. The following prompt displays:

Enable Spoofing? (y or n) (n) :

Enter **y** to proceed to enable spoofing.

- The following prompt displays:

Enable SPX Spoofing on all WAN routing services? (y or n) (n) : y

This prompt requires you to verify that you want to enable spoofing in order to avoid the situation of accidental spoofing of SPX packets. This example prompt asks if you want to enable SPX spoofing on all WAN routing services. If you had entered a specific Group number, the prompt would refer to that particular Group.

Enter **y** to enable spoofing.

- SPX Spoofing will then become active. A message will appear indicating that SPX Spoofing is enabled, either on all WAN routing services, or for a specific Group:

SPX Spoofing is now enabled on all WAN routing services

Disabling SPX Keepalive Spoofing

Follow the steps below to disable SPX Keepalive Spoofing.

- Enter **spxspoof**.

A screen similar to the following displays:

View the current spoofing status of all WAN routing services? (y or n) (n) :

Enter **y** if you do want to see the current SPX spoofing status. Enter **n** (or press **<Enter>**) if you entered this command by mistake or if you don't need to see the current status.

- A screen similar to the following displays:

Group	SPX Spoofing
-----	-----
3	Enabled
4	Enabled

Enter Group (default: all WAN) :

Enter a Group number to proceed to disable SPX spoofing for that particular Group.

Or, just press **<Enter>** to proceed to disable SPX spoofing for *all* WAN routing services.

- The following prompt displays:

Enable Spoofing? (y or n) (n) :

Enter **n** to proceed to disable spoofing.

- The following prompt displays:

Disable SPX Spoofing on all WAN routing services? (y or n) (n) : y

This prompt requires you to verify that you want to disable spoofing in order to avoid the situation of accidental spoofing of SPX packets. This example prompt asks if you want to disable SPX spoofing on all WAN routing services. If you had entered a specific Group number, the prompt would refer to that particular Group.

Enter **y** to disable spoofing.

- SPX Spoofing will then become inactive. A message will appear indicating that SPX Spoofing is disabled, either on all WAN routing services, or for a specific Group:

SPX Spoofing is now disabled on all WAN routing services

Controlling IPX Type 20 Packet Forwarding

The `ipxtype20` command is used to control the forwarding of IPX Type 20 packets. The default setting is to *not* forward IPX Type 20 packets. You can use the `ipxtype20` command to explicitly enable the forwarding of Type 20 packets for individual interfaces routing IPX traffic.

Type 20 packets contain the value 20 (14 hex) in the “packet type” field of the IPX header. Novell has defined the use of these packets to support certain protocol implementations, such as NetBIOS. As these packets are broadcasted and propagated across networks, the addresses of those networks (up to 8) are stored in the packet’s data area.

The reason why forwarding of Type 20 packets is normally “off” is that they can cause problems in highly redundant IPX networks by causing what appears to be a broadcast storm. This problem is aggravated whenever misconfigured PCs are added to a network.

Follow the steps below to enable IPX Type 20 packet forwarding on a given interface.

1. Enter `ipxtype20`.

A screen similar to the following displays:

```
Do you want to see the status of IPX Type 20 packet forwarding?  
(y or n) (y) :
```

2. Enter a **y** at this prompt (or press **<Enter>**) to display the current handling of IPX Type 20 packets on all configured IPX interfaces.

A screen similar to the following displays:

```
GP:VL   Type20 Packet Forwarding  
-----  
3:1     off  
4:1     off
```

```
group:vlan () :
```

3. Enter the Group and VLAN numbers associated with the IPX interface for which you wish to enable Type 20 packet forwarding. For example, you could enter **3:1**.

A screen similar to the following displays:

```
Currently, Group 3:Vlan 1 has IPX Type 20 packet forwarding off.  
“on” or “off” (off) :
```

4. Enter **on** to turn IPX Type 20 packet forwarding “on” for this interface. The default is “off”.

A screen similar to the following displays:

```
IPX Type 20 packet forwarding on 3:1 has been changed to on.
```

You may also elect to bypass the above prompts. To do so, simply include the Group and/or VLAN number and the word “on” (or “off”) as part of the command line.

For example, to turn forwarding “on” for Group 4, VLAN 1, enter `ipxtype20 4 on`.

A screen similar to the following displays:

```
IPX Type 20 packet forwarding on 4:1 has been changed to on.
```

If you enter the `ipxtype20` command with options for an interface that is not configured for IPX, a message similar to the following will appear:

```
Group 1:Vlan 1 isn't configured for IPX.
```

```
Usage: ipxtype20 [group:vlan] [on | off]
```

Configuring NetWare to Minimize WAN Connections

If you have access to NetWare's control parameters, you can "fine-tune" your network to minimize traffic on WAN links such as ISDN connections or Frame Relay lines. Doing so will reduce the costs associated with each connection that is made. Some suggested approaches are described below.

1. NetWare Directory Services (NDS), included in NetWare 4.x, includes a time synchronization protocol. By default, NetWare servers send time synchronization packets every 10 minutes. To help cut down on unnecessary connections that result from the time synchronization protocol, you could load the NLM (NetWare Loadable Module) named TIME-SYNC.NLM onto your NetWare time servers. This NLM will allow you to modify the update interval of the time synchronization packets.
2. NDS also introduces more traffic in order to maintain replicas of NDS partitions. The NLMs named DSFILTER.NLM and PINGFILT.NLM can be used to modify NDS synchronization updates.
3. NetWare's IPX Watchdog protocol monitors the connection status of NetWare clients and transmits reports when a connection fails to respond. You could modify the following three Watchdog parameters on your NetWare file servers to help cut down the costs associated with the IPX protocol:
 - SET NUMBER OF WATCHDOG PACKETS (the default is 10, range is 5 to 100 packets).
 - SET DELAY BETWEEN WATCHDOG PACKETS (the default is 59.3 seconds, range is 9.9 seconds to 10 minutes and 26.2 seconds).
 - SET DELAY BEFORE FIRST WATCHDOG PACKET (the default is 4 minutes 56.6 seconds, range is 15.7 seconds to 20 minutes and 52.3 seconds).
4. There are two basic categories of timeouts which can cause extra network traffic and/or loss of SPX connections:
 - If a data packet goes unacknowledged, it is re-transmitted a certain number of times before the connection is aborted.
 - When a connection is idle and the SPX Watchdog is enabled, system packets are sent periodically, and if not eventually acknowledged, the connection is aborted.
5. The following parameters can be modified in the NET.CFG file to determine when packets should be resent or when connections should be aborted:
 - MINIMUM SPX RETRIES determines how many unacknowledged transmit requests are allowed before assuming the connection has failed.
 - SPX VERIFY TIMEOUT determines how often (in ticks) the SPX protocol sends a packet to the other side of a connection to indicate that it is still alive.
 - SPX LISTEN TIMEOUT specifies how long (in ticks) the SPX protocol waits without receiving a packet from the other side of the connection before it requests the other side to send a packet to ascertain whether the connection is still valid.
 - SPX ABORT TIMEOUT specifies how long (in ticks) the SPX protocol waits without receiving any response from the other side of the connection before it terminates the session.

6. Novell has developed a workaround that can be used to disable the SPX Watchdog mechanism. This workaround could be used instead of enabling the SPX Spoofing feature on your switch. SPWXDOG.NLM is a patch that is used to disable NetWare's SPX Watchdog mechanism on 3.x and 4.x servers. The patch adds the following file server set parameter:

“set spx watchdogs=ON/OFF” (The default is ON.)

To fully disable SPX Watchdog packets, the remote client/server should also disable Watchdogs. IPXODI v3.02 and IPX.NLM support a NET.CFG parameter to disable SPX Watchdogs (“spx watchdog=off”).

Configuring RIP and SAP Timers

The standard time between broadcasts of RIP and SAP messages is 60 seconds. This default may be modified in order to alleviate network congestion or facilitate the discovery of network resources.

Adding a RIP and SAP Timer

1. To adjust the time between RIP and SAP messages, enter the following command at the system prompt:

```
ipxtimer
```

The following prompt displays:

```
Add or delete entry {add(a), delete(d)} (a) :
```

2. Enter **a** and the following prompt displays:

```
Group: (global) : 1
```

3. Enter the group number or leave the field blank and press Enter. If you do not enter a group number, the SAP and RIP timers will be adjusted for all groups on the switch.

The following prompt displays:

```
RIP timer (1..180 secs): (60) :
```

4. Enter the desired value or press **<Enter>** to configure the default value, which is 60 seconds. The following prompt displays:

```
SAP timer (1..180 secs): (60) :
```

5. Enter the desired value or press **<Enter>** to configure the default value, which is 60 seconds. The following message displays:

```
ipxtimer successfully added
```

Viewing RIP and SAP Timers

To view the RIP and SAP timers that have been configured through the **ipxtimer** command, enter the following command:

```
ipxt
```

A screen similar to the following displays:

<u>#</u>	<u>Group</u>	<u>RIP Timer (secs)</u>	<u>SAP Timer (secs)</u>
1	1	30	15
2	global	45	45

The fields are defined as follows:

Group

Displays the group number or **global** to indicate all groups.

RIP Timer (secs)

Displays the RIP timer configured for the group using the **ipxtimer** command.

SAP Timer (secs)

Displays the SAP timer configured for the group using the **ipxtimer** command.

Configuring Extended RIP and SAP Packets

Larger RIP and SAP packets may be transmitted so that congestion in the network is reduced. Other switches and routers in the network must support larger packet size if this feature is configured on the switch.

Use the **ipxext** command to enable or disable extended packets or to view the current status of extended packet transmission.

Enabling or Disabling Extended RIP and SAP Packets

To enable larger RIP and SAP packets, enter the following command:

```
ipxext on
```

To disable larger RIP and SAP packets, enter the following command:

```
ipxext off
```

Viewing the Current Status of Extended Packets

To display the current status of this feature, enter the following command:

```
ipxext
```

When the feature is disabled (the default), the following message displays:

```
IPX extended RIPs and SAPs off
```

When the feature is enabled, the following message displays:

```
IPX extended RIPs and SAPs on
```

Configuring an IPX Default Route

A default IPX route may be configured for packets destined for networks unknown to the switch. If RIP messages are disabled, packets can still be forwarded to a router that knows where to send them. Use the **ipxdrtr** command to add a default route, view the status of a default route, or disable the default route.

Adding an IPX Default Route

To configure a default route, use the **ipxdrtr** command with the relevant network ID. For example:

```
ipxdrtr 222
```

If the network ID indicates a direct network on the switch, the MAC address must also be specified, and the following prompt will display:

```
IPX node address of next hop (format - xx:xx:xx:xx:xx:xx) :
```

Enter the relevant address.

Viewing the Status of an IPX Default Route

To view the status of the default route, enter the **ipxdrtr** command. A message similar to the following displays:

```
IPX default route: 00000222 00:20:da:99:88:77
```

Disabling an IPX Default Route

To disable the default route, enter the following:

```
ipxdrtr off
```

If you enter the **ipxdrtr** command again, the following message displays:

```
IPX default route is disabled
```

28 Managing WAN Switching Modules

Introduction

The Omni Switch/Router WAN Switching Modules (WSXs) are a family of modules that enable the creation of WANs by providing connectivity between geographically-distanced LANs. These modules support a variety of protocols, including Frame Relay, synchronous Point to Point Protocol (PPP), and Integrated Services Digital Network (ISDN).

WSXs extend the power and flexibility of LAN switching over greater geographic distances using either a Frame Relay network, ISDN network or leased line connection, such as T1. In a Frame Relay network configuration, WSXs provide a cost-effective link that is capable of supporting multiple virtual circuits. In a leased line configuration, WSXs provide dedicated bandwidth to a single remote site. In an ISDN line configuration, the WSX supports both inbound and outbound call circuits for interconnection to remote WAN Switching Modules or other devices that support standard PPP over ISDN. In addition, an ISDN configuration supports bandwidth on demand and backup of failed lines.

The family of WSX modules provides either 2, 4, or 8 ports, which provide a range of access rates from 9.6 kbps to 2 Mbps. Management, data handling, compression, and multi-protocol encapsulation are compatible with the current Frame Relay and PPP standards.

VLAN architectures are preserved and consistent on both sides of a WAN link. WSXs support Alcatel Frame Relay trunking. As a result, VLAN groups on one side of a Frame Relay link are compatible with those on the other side. In addition, the WSX is capable of both Frame Relay and PPP transparent bridging, and IP and IPX routing.

VLAN architectures are preserved and consistent on both sides of a WAN link. The WSX supports standard RFC 1490 multiprotocol over Frame Relay and synchronous PPP for bridging and routing interoperability with numerous other WAN networking devices. In addition, the WSX supports Alcatel Frame Relay trunking, so multiple VLAN groups on one side of a Frame Relay link can be transported across the WAN.

Type of Service (ToS)

The Type of Service (ToS) settings allow you to prioritize voice data and voice signaling data. Since voice data is time critical, and requires steady throughput, it should be given higher priority than other forms of data. This can be done by assigning a priority value for the Voice Data and Voice Signaling Data fields.

There are two methods of specifying the ToS priority: IP Precedence and Differentiated Services Code Point (DSCP). Both of these methods use a binary value to indicate priority. IP Precedence uses three bits, and DSCP uses six bits; therefore the values for IP Precedence range from 0 to 7, and the values for DSCP range from 0 to 63. The higher the number, the higher the priority of the traffic. IP Precedence uses the most significant (upper) 3 bits, and DSCP uses the most significant (upper) 6 bits.

◆ Important Note ◆

The ToS setting is a feature of Quality of Service (QoS). It was set into the UI to make Voice over IP (VoIP) modules compatible with an Omni Switch/Router that does not have QoS installed. If the QoS image is installed (**qos.img**) on the switch, the UI ToS commands do not function. They are overridden by standard QoS functionality.

Currently, the defaults set for voice data and signaling data are the setting recommended by both Alcatel and Cisco. The default values for switches use hexadecimal forms of IP Precedence; the default value for voice data is 5 decimal, and the default value for signaling is 3 decimal.

Below is a table that shows the relation of IP Precedence levels and DSCP levels.

Relation of IP Precedence, DSCP, and Level of Priority

Priority Level	IP Precedence Value (Decimal)	DSCP Value Range (Decimal)
Routine	0	0 - 7
Priority	1	8 - 15
Immediate	2	16 - 23
Flash	3	24 - 31 (AF31)
Flash-Override	4	32 - 39
Critical	5	40 - 47 (EF)
Internet	6	48 - 55
Network	7	56 - 63

If you feel that changing the default values is imperative to the working of the network, the following table is provided to give the hexadecimal values for various settings:

Hexadecimal Settings

IP Precedence Value	Hexadecimal Value	DSCP Value	Hexadecimal Value
0	0	0	0
1	20	10 (AF11)	28
2	40	18 (AF21)	48
3*	60	26 (AF31)*	68
4**	80	34 (AF41)**	88
5***	a0	46 (EF)***	b8
6	c0	54	d8
7	e0	62	f8

*Default settings for signalling data.

**Cisco suggested default settings for video data.

***Default settings for voice data.

A bit mask is also set with the UI in hexadecimal form. The mask is used during the lookup phase of ToS and screens out the insignificant bits. For IP precedence, the mask should be set to **e0** (this is the default value). For DSCP, the mask is **fc**.

◆ Important Note ◆

These values are set to work with the Alcatel VoIP modules. DO NOT attempt to change them unless you are an advanced user with detailed knowledge of Alcatel products and how they interact.

ToS and QoS Interaction

On the Omni Switch/Router and OmniSwitch, ToS policies may only be configured through WAN commands. The WAN UI/CLI commands allow a higher priority of service for voice and voice signalling data.

WAN ToS policies are supported when bridging or routing; ToS policies configured through the QoS Manager are only supported for routing.

With the WAN ToS feature, the switch can examine an egress IP packet (either bridged or routed) on a WAN interface and compare the value in the ToS/DSCP header with a configured mask and value for voice data. If it does not match, the value is then compared to the ToS value for voice signaling. (the default values correspond to compatible values for Alcatel Voice over IP modules).

On the Omni Switch/Router or OmniSwitch, when the frame matches either the voice data or voice signaling value the frame is forwarded to the high priority queue without any bandwidth monitoring or restrictions.

The WAN ToS feature should be used for ToS/DSCP flows in trusted networks when no other QoS mechanism is required, such as a voice/data converged network with private WAN infrastructure. Typically the default WAN ToS settings are sufficient. For more information about WAN modules and WAN ToS commands, see your switch user manual.

Use the QoS Manager instead of the WAN ToS feature if you need to give priority to other types of traffic or if you want to classify or filter traffic. (The QoS Manager and WAN ToS features cannot be used together.) For example, if a WAN network is not used for voice/data convergence but prioritization is required for data traffic, use the QoS Manager to create a rule to classify the ToS/DSCP flow with the Voice over IP (VoIP) gateway address and give priority to the flow.

DTR Dial Backup

Currently, a feature is available to use a dynamic ISDN call as a backup WAN connection for a primary WAN connection. The primary WAN connection is a permanent virtual circuit (PVC) with upper layer protocol of Frame Relay or PPP. The DTR Dial Backup feature will allow another synchronous serial interface to be used for the backup purpose.

The process is analogous to the ISDN backup feature. When the primary circuit fails over a configurable elapsed time, the initiating side of the backup connection asserts the DTR line of the designated backup port, causing the attached modem to dial a pre-configured phone number (i.e. the phone number to call is stored within the modem itself). The modem shall assert the DSR, CD, and the CTS signals of the serial interface when the call is established and the modems have fully trained and synchronized. Sensing the call having completed successfully (via a port-up event), the PPP negotiation commences. Upon a successful PPP negotiation, the backup virtual circuit, the virtual port and the associated router interface are activated. The receiving side of the backup connection shall wait for the attached modem to signal an incoming call (again by looking for all three primary modem control signals: DSR, CD and CTS to be active and the resulting port-up event). The backup virtual circuit, virtual port, and router interface are activated after the ensuing PPP negotiation completes successfully.

After the primary circuit is restored over a configurable elapsed time, the call initiation side shall turn off the DTR signal causing the modem to terminate the call. It will then also deactivate the backup circuit at this time. The call receiving side shall receive a indication from the attached modem on call termination and proceed to deactivate the backup connection.

This feature only works in conjunction with modems configured for DTR dialing and synchronous mode operation.

The DTR dial backup feature is only valid on serial ports using Point to Point Protocol (PPP).

Supported Physical Interfaces

The WSX family of products support numerous physical interface (port) types. The port types available with the WSX family are:

Universal Serial Port

The Universal Serial Port (USP) provides connectivity to legacy synchronous serial port devices. With the addition of an adapter cable, it supports RS-232, RS-449, RS-530, V.35 and X.21 Data Terminal Equipment (DTE) and Data Carrier Equipment (DCE) interfaces at speeds up to 2.048 Mbps. USPs support access via Frame Relay or synchronous PPP. The WSX automatically detects the cable type connected and will configure the correct physical interface to use.

ISDN Basic Rate Interface Port

The ISDN Basic Rate Interface (BRI) port supports either a U or S/T interface (jumper selectable) for interfacing to public or private ISDN networks. Synchronous PPP is supported on the two bearer (B) channels. Multiple ISDN switch protocol variations are supported on the delta (D) channel (used for signaling). Each B channel runs at 64 kbps, and the D channel runs at 16 kbps.

Fractional T1 Port

The fractional T1 port connects directly to North American and Japanese circuit switch digital data public or private networks without requiring an external Digital Service Unit/Channel Service Unit (DSU/CSU). The port provides an integral DSU/CSU function with both short-haul (i.e., short distance) and long-haul (i.e., long distance) capabilities. The port allows the user to configure a range of time slots from 1 to 24 time slots used to allow for full T1 (all 24 time slots used) or a fractional T1 (less than 24 time slots) service. The fractional T1 port can support access via Frame Relay or synchronous PPP.

◆ Note ◆

For public digital networks, check with your service provider. They may allow only connections that use a configured short-haul interface via a network-provided Channel Service Unit (CSU).

Fractional E1 Port

The fractional E1 port connects directly to ITU-T standard circuit switch digital data public or private networks without requiring an external DSU/CSU. The port provides an integral DSU/CSU function with both short-haul (i.e., short distance) and long-haul (i.e., long distance) capabilities. The port allows you to configure for full E1 (all 30 or 31 time slots used) or fractional E1 (1-29 time slots) service. The fractional E1 port supports access via either Frame Relay or synchronous PPP.

◆ Note ◆

For public digital networks, check with your service provider. They may allow only connections that use a configured short-haul interface via a network-provided Channel Service Unit (CSU).

Supported Protocols

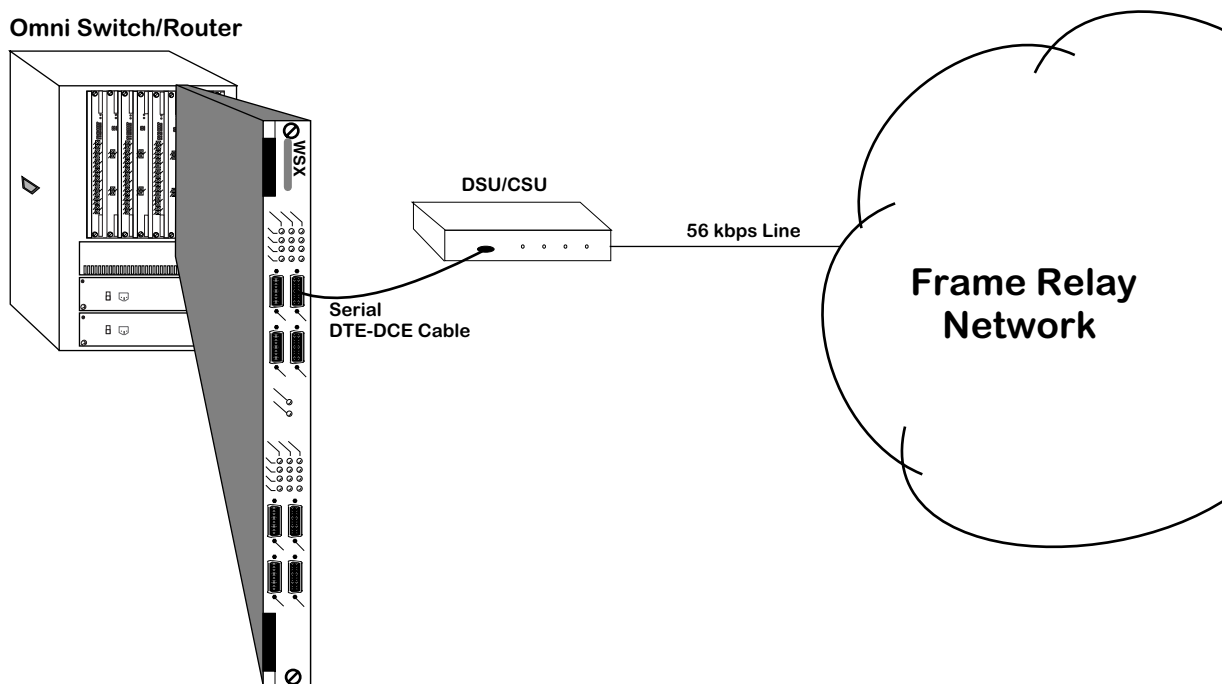
The WAN switching modules support both Frame Relay and synchronous Point-To-Point Protocol (PPP). For ISDN signalling protocols, the modules support D-channel signalling (see Chapter 32, “Managing ISDN Ports.” For more details on implementing these protocols, see Chapter 29, “Managing Frame Relay,” and chapter 30, “Point-to-Point Protocol.”

Application Examples

This section provides several examples of the types of WAN networking possible using WAN switching modules.

Frame Relay WSX Using Serial Ports

In a typical configuration, the WSX occupies either a slot in a switch chassis or a submodule in an OmniAccess 512. Because it is compatible with Omni Switch/Router any-to-any switching and VLAN architecture, you can switch other topologies in the LAN to Frame Relay or PPP. The WSX connects to a DSU/CSU or T1 multiplexer through a serial cable. The following diagram shows a typical WSX setup using a 56 kbps Frame Relay line (up to 2 Mbps access rates are supported).




WSX Frame Relay Configuration Using Serial Ports

For serial ports, the WSX supports automatic detection of cable types. It also supports internal, external, and split clocking.

Software in the switch allows you to configure access rate, clocking and protocol-related parameters. Additional software commands allow you to view status at the WSX board, port, or protocol level. Extensive statistics are provided at each level, including a breakdown of traffic by frame type (Ethernet, IP, IPX, or BPDU) at the virtual circuit or PPP connection level.

Back-to-Back WSX Using T1 Ports

WAN switching modules may be connected “back-to-back” without an intervening Frame Relay network or switch. Because the T1 port internally provides a DSU/CSU function, an external DSU/CSU is not required. Such connections are made by using private leased lines, such as T1 lines, instead of public Frame Relay networks, usually over large geographic distances.




Back-to-Back Configuration Using Fractional T1 Ports

Combined Frame Relay with ISDN Backup

In a typical configuration, the WSX occupies either a slot in a switch chassis or a submodule on an OmniAccess 512. Because it is compatible with Omni Switch/Router any-to-any switching and VLAN architecture, you can switch other topologies in the LAN to Frame Relay or PPP. The WSX connects to a DSU/CSU or T1 multiplexer through a serial cable. The following diagram shows a typical WSX setup using a 56 kbps Frame Relay line (up to 2 Mbps access rates are supported)

Refer to the Chapter 29, “Managing Frame Relay,” and Chapter 56, “Backup Services,” for details on how to implement this configuration.³⁴



Omni Switch/Router WAN Modules

The Omni Switch/Router currently supports four Wide Area Network modules:

- **WSX-S-2W** Provides two serial ports that support the frame relay or PPP protocol.
- **WSX-SC** Provides 4 or 8 serial ports that support the frame relay or PPP protocol. In addition, hardware compression is also supported.
- **WSX-FT1/E1-SC** Provides one or two T1/E1 ports and one or two serial ports that support the frame relay or PPP protocol
- **WSX-BRI-SC** Provides one or two UPS (Universal Serial Port) and 1 or 2 ISDN-BRI ports that support Frame Relay or PPP

The WSX-S-2W, WSX-SC, WSX-FT1/E1-SC, and WSX-BRI-SC modules for the Omni Switch/Router are described in Chapter 3, “Omni Switch/Router Switching Modules.”

Cable Interfaces for Universal Serial Ports

The WSX automatically senses the cable type that you plug into one of its Universal Serial Ports. It can sense whether the cable type is DCE or DTE and whether it is one of the following interfaces:

- RS-232
- RS-449
- RS-530
- V.35
- X.21 (European)

All cable types (except RS-232) are capable of access rates from 9.6 kbps to 2 Mbps. The RS-232 cable is not compatible with speeds greater than 64 kbps. Each cable type is illustrated and described in Appendix B, "Custom Cables."

The WSX port is normally considered a physical DTE device. It is possible to turn it into a physical DCE device simply by plugging in a DCE cable. The WSX board internally senses whether a DCE or DTE cable is connected.

DTE/DCE Type and Transmit/Receive Pins

The RS-232 protocol, which is employed at the physical level for all cable types, always defines Transmit and Receive pins in relation to the DTE. So, the type of cable you attach (DCE or DTE) determines the direction of data flow on your connector's Transmit and Receive pins.

If the WSX port is a physical DTE, which is probably the most common configuration, then data is received on Receive pins and transmitted on Transmit pins. If you are using a WSX port as a physical DCE, then data is transmitted on the receive pins and received on the transmit pins.

Data Compression

Data compression allows you to get more data through the Frame Relay pipeline, further enhancing cost benefits. A typical data compression ratio on the WSX board at the hardware level is 4:1. In addition, the compression processor (STAC 9705) has its own memory (DRAM) that can store up to 100 compression histories (on a 4-port WSX) without degrading performance. An 8-port WSX can store up to 200 compression histories without performance degradation.

The WSX will only compress data if you enable compression through software and the bridge/router on the other end of the connection supports standard Frame Relay (FRF.9) or PPP (STAC-LZS) compression. (An Omni Switch/Router-to-Omni Switch/Router connection would support compression.) Negotiation is necessary because if compressed data is sent to a bridge/router that does not support compression, this bridge/router will not recognize the data and will automatically drop the unrecognizable frames.

If you enable compression, the WSX will query the Frame Relay or PPP device on the other end of the circuit to determine whether it supports compression. If it does, the WSX will compress all data except frame DLCMI (management) data and PPP control messages. If it does not support compression, data on that connection will be sent uncompressed. Refer to either Chapter 29, "Managing Frame Relay," or Chapter 30, "Point-To-Point Protocol," for information on enabling compression.

Note

Compression is not supported on the Omni Switch/Router WSX-S-2W modules.

Loopback Detection

Loopback Detection is a common method for Carrier Service Providers to test clients' circuits in the event of suspected line transmission problems. For both Frame Relay and PPP, loopback detection involves periodically transmitting a message and looking for that message to be received. When implementing Loopback Detection, it is important to keep two issues in mind: the message must not violate any standards; the message must be unique in such a way that it can be differentiated from a message sent by a remote node.

The messages are transmitted in one of two fixed intervals. When the port is in normal mode, the message is transmitted once every second. When two consecutive messages are received that match the transmitted message, the port is considered to be in loopback. Once in loopback mode, the message is transmitted once every 100 milliseconds. After ten consecutive messages are transmitted without receiving a match, the port is returned to normal mode. Consequently, it takes up to 2 seconds to detect the loopback condition and an additional second to exit it.

The message sent on a Frame Relay port uses standard 1490 encapsulation with a SNAP header. The OUI (Organizationally Unique Identifier) of the SNAP header is the Alcatel OUI, so encapsulation is standard, but the message is proprietary. The message is transmitted using the lowest available DLCI, or 32 if there are no DLCI's operating on the port. Because the message is merely attempting to determine the state of the physical port, the state of the DLCI, whether active, inactive or non-existent, is not important; the Frame Relay switch will discard any data for non-existent or inactive DLCIs.

The message sent on a PPP port uses the standard LCP Echo message.

Uniqueness of messages is accomplished by including a word in the message that is based upon the configuration of the port and a free-running timer. For PPP, uniqueness is enhanced by negotiating the LCP magic number option.

The WAN Port Software Menu

User interface commands for the WSX board are on a separate menu that is accessed through the **wan** command. The WAN Port menu is a submenu of the Interface menu. Typing **wan** at any system prompt displays the following menu:

Command	Wide Area Networking Menu				
wpmodify	Modify a given WAN port's parameters				
wpdelete	Delete a given port's parameters, and restore defaults				
wpview	View WAN port parameters for a given slot and port				
wpstatus	View WAN port status of entire chassis, slot, or individual port				
fr	Enter the Frame Relay submenu				
ppp	Enter the PPP submenu				
isdn	Enter the ISDN-specific submenu				
link	Enter the link-specific submenu				
Main	File	Summary	VLAN	Networking	
Interface	Security	System	Services	Help	

◆ Note ◆

The ISDN menu will only appear on systems with a least one WSX-BRI module installed.

You can start most of the commands by typing the first three (3) letters of the command name. For example, to use the **wpview** command, type **wpv**.

The following sections describe the use of commands on the WAN Port menu.

Setting Configuration Parameters

When you plug in a WSX board, it is automatically configured to the default settings. By default, the WSX uses Frame Relay protocol. In addition, the access rate for serial ports defaults to 64 kbps for RS-232 cables. The access rate for other cable types defaults to 2 Mbps. You can change these settings, as well as several other settings, such as clocking and protocol type, with the **wpmodify** command.

Modifying a Port

Use the **wpmodify** command to modify a port, as shown below:

```
wpmodify <slot>/<port>
```

in which **<slot>** is the slot number where the WSX board is located, and **<port>** is the port number on the WSX board that you want to modify. When this command is entered, the system automatically senses what type of port is being configured, and displays the appropriate screen for that type of port. The screen is different depending on the type of encapsulation used, either Frame-Relay or Point-to-Point Protocol.

Make changes by entering the line number for the option you want to change, an equal sign (=), and the value for the new parameter. When you have finished entering the new values, type **save** at the prompt to save the new parameters. The following sections describe the options you can alter through this menu. The following three examples show a typical setup screen for a serial port, an ISDN-BRI port, and a fractional T1 port, respectively.

Serial Port Example

In this example, port 1 on slot 3 is a serial port, using Frame-Relay. To modify serial port 3/1, enter:

```
wpm 3/1
```

A screen similar to following displays:

```

1) Admin Status ..... UP
   {(U)p, (D)own}
2) Speed in BPS ..... 2048000
   {9600, 19200, 56000, 64000, 128000, 256000, 512000, 768000}
   {1024000, 1544000, 2048000}
3) Clocking ..... External
   {(I)nternal, (E)xternal, (S)plit}
4) Protocol Type ..... Frame Relay
   {(F)rame Relay, (P)PP(Point to Point)}
7) Receive Clock ..... Normal
   {(N)ormal, (I)nverted}
8) TOS for Voice Data ..... a0
   TOS Value in Hex for Voice Data, 0 to disable TOS
9) TOS for Voice Signaling Data ..... 60
   TOS Value in Hex for Voice Signaling Data, 0 to disable TOS
10) TOS Mask for both TOS Value ..... e0
   TOS Mask in Hex for Type of Data
11) Signaling IP Address ..... 0.0.0.0
   IP Address range
12) Signaling IP Mask ..... 255.255.255.255
   IP Mask range
15) Loopback Timeout ..... 10
   {Timeout (0..255)}

```

If the interface was using PPP, the following screen would display:

```
1) Admin Status ..... UP
   {(U)p, (D)own}
2) Speed in BPS ..... 2048000
   {9600, 19200, 56000, 64000, 128000, 256000, 512000, 768000}
   {1024000, 1544000, 2048000}
3) Clocking ..... External
   {(I)nternal, (E)xternal, (S)plit}
4) Protocol Type ..... Frame Relay
   {(F)rame Relay, (P)PP(Point to Point)}
7) Receive Clock ..... Normal
   {(N)ormal, (I)nverted}
8) TOS for Voice Data ..... a0
   TOS Value in Hex for Voice Data, 0 to disable TOS
9) TOS for Voice Signaling Data ..... 60
   TOS Value in Hex for Voice Signaling Data, 0 to disable TOS
10) TOS Mask for both TOS Value ..... e0
   TOS Mask in Hex for Type of Data
11) Signaling IP Address ..... 0.0.0.0
   IP Address range
12) Signaling IP Mask ..... 255.255.255.255
   IP Mask range
13) KeepAlive Up Count ..... 0
   {Up Count (0..255)}
14) KeepAlive Down Count ..... 0
   {Down Count (0..255)}
15) KeepAlive Timeout ..... 10
   {Timeout (0..255)}
16) DTR Pulse Width ..... 0
   {Pulse Width (0..255)}
17) DTR Pulse Count ..... 0
   {Pulse Count (0..255)}
20) Connection Function ..... Dedicated
```

Admin Status

The options for the Admin Status are **UP** and **DN**. If **UP**, the port has been enabled and can transmit data as long as its Operational Status is also **UP**. If set to **DN**, the port will not pass data even, if its physical connection is good.

Speed in BPS

This option specifies the access rate for the Frame Relay or PPP line to the service provider. This parameter is the speed of the entire connection, not an individual virtual circuit. For example, if you have a 56 kbps line to your service provider, this field should be set to 56000. A full T1 line would have an access rate of 1,544,000 bps, and a full E1 line would have an access rate of 2,048,000 bps. For either T1 or E1, you can also have a fractional service with an access rate that is a multiple of 64 kbps. Enter a value that is the same as one of the values displayed below this field.

Note

If the port you are configuring is a physical DCE port (i.e., DCE cable plugged into the submodule port) that can control the access rate and clocking, always enter a value for this field. This value will be used in computing congestion control parameters, such as the Committed Information Rate (CIR). If the port is a DTE, this setting will have no effect, except for informational purposes.

Clocking

This field sets the type of clocking used to clock transmit and receive data on the serial port. If the clock goes out-of-phase, you will receive errors.

Note

The clocking value is only relevant if the port is a physical DCE port (i.e., DCE cable plugged into the submodule port). If the submodule port is a physical DTE port, clocking will default to External.

External Clocking

If you set this value to External, clocking will be controlled by the external DCE (a DSU or other DCE device on the other end of the cable from the submodule port). External clocking is the default option when the submodule is a physical DTE device (i.e., controlled by an external DCE device).

When the submodule is acting as a physical DTE and

- the speed is greater than 256 kbps, or
- excessive FCS errors or Aborts are being detected coming from the submodule at the remote port or line monitor

then it is recommended that the external DCE (usually a CSU/DSU) be set to take a transmit data clock from the external DTE transmit clock (TXCE).

You can set up the external DCE this way by configuring its DTE, or dataport, configuration options. Set the "Transmit Clock Source" to "External." In this mode of operation, the transmit clock is output by the DCE device and the submodule turns it around on the external transmit clock back to the DCE, eliminating any phase misalignment between transmit clock and transmit data.

If the external DCE does not provide a DTE configuration option for the transmit clock source, then try setting the "Transmit Clock *Polarity*" to "Invert." Note that Invert is the clock polarity for Transmit (not Receive) and should only be used when excessive FCS errors or Aborts are detected.

Internal Clocking

If you set this value to Internal, clocking is controlled by the internal DCE (the submodule). Internal clocking should only be selected if the submodule is a physical DCE device and you are using an RS-232 cable. Internal clocking is the default setting when the submodule is a physical DCE device and an RS-232 DCE cable is connected to this port.

Split Clocking

Split clocking, which is also known as “loop timing,” uses an additional control signal (TXCE) to keep the submodule and external DTE clocking synchronized. In split clocking, the external DTE takes the incoming transmit clock from the submodule and loops it back to TXCE. The submodule then uses this signal to clock in data from the external DTE device. Split clocking should only be used if the submodule is a physical DCE device and you are using a non-RS-232 cable, such as V.35.

◆ Important Note ◆

Split clocking is required if the access rate of the submodule port is greater than 256 kbps and it is acting as a DCE device. If split clocking is not used at these data rates, data out-of-phase errors, aborts, or CRC errors may occur.

Split clocking is the default when the submodule port is a physical DCE device and a non-RS-232 DCE cable is connected to the port.

Protocol Type

The protocol type can be set to either Frame Relay or Point to Point Protocol (PPP). The default setting is Frame Relay.

◆ Important Note ◆

A port must be set to either Frame Relay or PPP before any other protocol-related parameters can be set.

Receive Clock

Often, due to delays added to timestamps in when running through switch hardware, the receive clock time is significantly different than expected from the transmitting data source. To correct the problem, it is possible to set the receive clock to invert the delay information. The following options are available:

Normal

The port uses the internal clock time as the timestamp for receive data (timestamp information is not modified).

Inverted

The port uses an inverted timestamp for receive data.

TOS for Voice Data

Set the priority for voice data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). Enter **0** to disable this feature. See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS for Voice Signaling Data

Set the priority for voice signaling data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). Enter **0** to disable this feature. See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS Mask for both TOS Value

Set the mask bits for both voice data and signaling data. Enter **0** to disable this feature.

Signaling IP Address

When a data frame cannot be identified by the ToS voice data or ToS signaling data, the Signaling IP Address is checked. Matched frames are loaded on a High Priority Software Queue and a Nominal Hardware Queue.

Signaling IP Mask

The mask associated with the Signaling IP Address described above.

KeepAlive Up Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Up Count is the number of requests sent from the port when the port transitions from Down to Up, to verify the status of the port. The valid range is 0-255.

KeepAlive Down Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Down Count is the number of requests sent from the port when the port transitions from Up to Down, to verify the status of the port. This only displays if the port is using PPP as its encapsulation type. The valid range is 0-255.

KeepAlive Timeout

The number of 100 millisecond increments between generated echo message requests. This only displays if the port is using PPP as its encapsulation type. The valid range is 0-255.

DTR Pulse Width

A Data Terminal Ready (DTR) Pulse is sent at the hardware level to determine a port is still synchronized with its far end connection. The Pulse Width is the number of 100 millisecond increments that the pulse lasts. This only displays if the port is using PPP as its encapsulation type. The valid range is 0-255.

DTR Pulse Count

A Data Terminal Ready (DTR) Pulse is sent at the hardware level to determine a port is still synchronized with its far end connection. The Pulse Count is the number of pulses generated when a line is down. This only displays if the port is using PPP as its encapsulation type. The valid range is 0-255.

Loopback Timeout

Sets the transition time between proprietary messages sent over the link. These messages are analyzed to determine whether the link is in a loopback state. This only displays if the port is using Frame Relay as its encapsulation type. The valid range is 0-255.

Connection Function

On serial ports using PPP, it is possible to configure the port to be a DTR dial backup port. In case the primary WAN port fails, a DTR dial backup port will use a synchronous modem to dial out and reestablish the WAN connection. If this feature is disabled, the **Connection Function** field will read **Dedicated**. If it is enabled, the field will read **DTR-Dial**, and the following options will also be displayed:

200) **Connect Timeout (seconds)** 60
201) **Retry Delay (seconds)** 10

Connect Timeout

The number of seconds the switch attempts to establish a connection via this port before declaring the attempt a failure. The valid range is 10 to 2147483647. The default is 60. If the value is entered as 0, the attempt will never timeout.

Retry Delay

The number of seconds after a connection failure the switch waits before attempting a new connection. The valid range is 1 to 2147483647. The default is 1.

ISDN-BRI Port Example

In this example: port 2 on slot 3 is an ISDN-BRI port. To modify ISDN-BRI port 2/2, enter:

```
wpm 3/2
```

A screen similar to following displays:

```

1) Admin Status ..... UP
   {(U)p, (D)own}
2) Speed in BPS ..... 2048000
   {9600, 19200, 56000, 64000, 128000, 256000, 512000, 768000}
   {1024000, 1544000, 2048000}
3) Clocking ..... External
   {(I)nternal, (E)xternal, (S)plit}
4) Protocol Type ..... Frame Relay
   {(F)rame Relay, (P)PP(Point to Point)}
8) TOS for Voice Data ..... a0
   TOS Value in Hex for Voice Data, 0 to disable TOS
9) TOS for Voice Signaling Data ..... 60
   TOS Value in Hex for Voice Signaling Data, 0 to disable TOS
10) TOS Mask for both TOS Value ..... e0
   TOS Mask in Hex for Type of Data
11) Signaling IP Address ..... 0.0.0.0
   IP Address range
12) Signaling IP Mask ..... 255.255.255.255
   IP Mask range
13) KeepAlive Up Count ..... 0
   {Up Count (0..255)}
14) KeepAlive Down Count ..... 0
   {Down Count (0..255)}
15) KeepAlive Timeout ..... 10
   {Timeout (0..255)}
16) DTR Pulse Width ..... 0
   {Pulse Width (0..255)}
17) DTR Pulse Count ..... 0
   {Pulse Count (0..255)}

```

Note that the only parameters you can set for an ISDN port from this screen is the Admin Status and the ToS settings. All other parameters must be set from the ISDN, PPP, peer or WAN link menus. For more details on ISDN ports, see Chapter 32, “Managing ISDN Ports.” For more details on managing PPP ports, see Chapter 30, “Point-to-Point Protocol.” For more information on managing WAN links, see Chapter 31, “WAN Links.”

◆ Note ◆

The ISDN **wpmmodify** menu displays PPP specific line options described in the section *Modifying a Port* on page 28-14. However, they do not apply to an ISDN port, and are not described below.

Admin Status

The options for the Admin Status are **UP** and **DN**. If **UP**, the port has been enabled and can transmit data as long as its Operational Status is also **UP**. If set to **DN**, the port will not pass data even if its physical connection is good.

Speed in BPS

This option specifies the access rate for the Frame Relay or PPP line to the service provider. This parameter is the speed of the entire connection, not an individual virtual circuit. For example, if you have a 56 kbps line to your service provider, this field should be set to 56000. A full T1 line would have an access rate of 1,544,000 bps, and a full E1 line would have an access rate of 2,048,000 bps. For either T1 or E1, you can also have a fractional service with an access rate that is a multiple of 64 kbps.

Enter a value that is the same as one of the values displayed below this field.

Note

If the port you are configuring is a physical DCE port (i.e., DCE cable plugged into the submodule port) that can control the access rate and clocking, always enter a value for this field. This value will be used in computing congestion control parameters, such as the Committed Information Rate (CIR). If the port is a DTE, this setting will have no effect, except for informational purposes.

Clocking

This field sets the type of clocking used to clock transmit and receive data on the serial port. If the clock goes out-of-phase, you will receive errors.

Note

The clocking value is only relevant if the port is a physical DCE port (i.e., DCE cable plugged into the submodule port). If the submodule port is a physical DTE port, clocking will default to External.

External Clocking

If you set this value to External, clocking will be controlled by the external DCE (a DSU or other DCE device on the other end of the cable from the submodule port). External clocking is the default option when the submodule is a physical DTE device (i.e., controlled by an external DCE device).

When the submodule is acting as a physical DTE and

- the speed is greater than 256 kbps, or
- excessive FCS errors or Aborts are being detected coming from the submodule at the remote port or line monitor

then it is recommended that the external DCE (usually a CSU/DSU) be set to take a transmit data clock from the external DTE transmit clock (TXCE).

You can set up the external DCE this way by configuring its DTE, or dataport, configuration options. Set the “Transmit Clock Source” to “External.” In this mode of operation, the transmit clock is output by the DCE device and the submodule turns it around on the external transmit clock back to the DCE, eliminating any phase misalignment between transmit clock and transmit data.

If the external DCE does not provide a DTE configuration option for the transmit clock source, then try setting the “Transmit Clock *Polarity*” to “Invert.” Note that Invert is the clock polarity for Transmit (not Receive) and should only be used when excessive FCS errors or Aborts are detected.

Internal Clocking

If you set this value to Internal, clocking is controlled by the internal DCE (the submodule). Internal clocking should only be selected if the submodule is a physical DCE device and you are using an RS-232 cable. Internal clocking is the default setting when the submodule is a physical DCE device and an RS-232 DCE cable is connected to this port.

Split Clocking

Split clocking, which is also known as “loop timing,” uses an additional control signal (TXCE) to keep the submodule and external DTE clocking synchronized. In split clocking, the external DTE takes the incoming transmit clock from the submodule and loops it back to TXCE. The submodule then uses this signal to clock in data from the external DTE device. Split clocking should only be used if the submodule is a physical DCE device and you are using a non-RS-232 cable, such as V.35.

◆ Important Note ◆

Split clocking is required if the access rate of the submodule port is greater than 256 kbps and it is acting as a DCE device. If split clocking is not used at these data rates, data out-of-phase errors, aborts, or CRC errors may occur.

Split clocking is the default when the submodule port is a physical DCE device and a non-RS-232 DCE cable is connected to the port.

Protocol Type

The protocol type can be set to either Frame Relay or Point to Point Protocol (PPP). The default setting is Frame Relay.

◆ Important Note ◆

A port must be set to either Frame Relay or PPP before any other protocol-related parameters can be set.

TOS for Voice Data

Set the priority for voice data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). Enter **0** to disable this feature. See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS for Voice Signaling Data

Set the priority for voice signaling data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). Enter **0** to disable this feature. See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS Mask for both TOS Value

Set the mask bits for both voice data and signaling data. Enter **0** to disable this feature.

Signaling IP Address

When a data frame cannot be identified by the ToS voice data or ToS signaling data, the Signaling IP Address is checked. Matched frames are loaded on a High Priority Software Queue and a Nominal Hardware Queue.

Signaling IP Mask

The mask associated with the Signaling IP Address described above.

Fractional T1 Port Example

In this example: port 1 on slot 3 is a fractional T1 port using Frame Relay. To modify fractional T1 port 2/1, enter:

wpm 3/1

A screen similar to following displays:

```
1) Admin Status ..... UP
   {(U)p, (D)own}
2) Speed in BPS ..... 1544000
3) Clocking ..... Local
4) Protocol Type ..... Frame Relay
   {(F)rame Relay, (P)PP(Point to Point)}
5) T1 Starting Time Slot ..... 1
   {T1 (1..24)}
6) T1 Number of Time Slots ..... 23
   {T1 (1..24)}
8) TOS for Voice Data ..... a0
   TOS Value in Hex for Voice Data, 0 to disable TOS
9) TOS for Voice Signaling Data ..... 60
   TOS Value in Hex for Voice Signaling Data, 0 to disable TOS
10) TOS Mask for both TOS Value ..... e0
   TOS Mask in Hex for Type of Data
11) Signaling IP Address ..... 0.0.0.0
   IP Address range
12) Signaling IP Mask ..... 255.255.255.255
   IP Mask range
15) Loopback Timeout ..... 10
   {Timeout (0..255)}
(save/quit/cancel)
:
```

If the interface was using PPP, the following screen would display:

1) Admin Status	UP
{(U)p, (D)own}	
2) Speed in BPS	1544000
3) Clocking	External
{(I)nternal, (E)xternal, (S)plit}	
4) Protocol Type	PPP
{(F)rame Relay, (P)PP(Point to Point)}	
7) Receive Clock	Normal
{(N)ormal, (I)nverted}	
8) TOS for Voice Data	a0
TOS Value in Hex for Voice Data, 0 to disable TOS	
9) TOS for Voice Signaling Data	60
TOS Value in Hex for Voice Signaling Data, 0 to disable TOS	
10) TOS Mask for both TOS Value	e0
TOS Mask in Hex for Type of Data	
11) Signaling IP Address	0.0.0.0
IP Address range	
12) Signaling IP Mask	255.255.255.255
IP Mask range	
13) KeepAlive Up Count	0
{Up Count (0..255)}	
14) KeepAlive Down Count	0
{Down Count (0..255)}	
15) KeepAlive Timeout	10
{Timeout (0..255)}	
16) DTR Pulse Width	0
{Pulse Width (0..255)}	
17) DTR Pulse Count	0
{Pulse Count (0..255)}	

◆ Note ◆

The DTR Pulse settings do not apply to T1 and E1 interfaces, and are not described below.

Admin Status

The options for the Admin Status are **Enable** and **Disable**. If **Enable**, the port has been enabled and can transmit data as long as its Operational Status is also enabled. If set to **Disable**, the port will not pass data, even if its physical connection is good.

Speed in BPS

This field shows the speed for the T1/E1 port. This field is for reference only.

Clocking

This field shows the type of clocking set for the T1 port. This field is for reference only.

Protocol Type

The protocol type can be set to either Frame Relay or Point to Point Protocol (PPP).

◆ Important Note ◆

A port must be set to either Frame Relay or PPP before any other protocol-related parameters can be set.

T1/E1 Starting Time Slot

This field specifies the first time slot number to use on a T1 or E1 port. For a full T1 or E1 connection, specify time slot 1. For a fractional T1 or E1 connection, set this field to the starting time slot number as specified by your service provider.

T1/E1 Number of Time Slots

This field specifies the total number of 64 kbps time slots to use on the T1 or E1 connection. For a full T1, set this number to 24. For a full E1 connection, set this number to 30 if you are running multiframe; otherwise, set to 31. For fractional T1 or E1, you must set the number of time slots to the value specified by your service provider. For example, a 256 kbps service uses four time slots ($4 \times 64 = 256$).

TOS for Voice Data

Set the priority for voice data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). Enter **0** to disable this feature. See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS for Voice Signaling Data

Set the priority for voice signaling data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). Enter **0** to disable this feature. See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS Mask for both TOS Value

Set the mask bits for both voice data and signaling data. Enter **0** to disable this feature.

Signaling IP Address

When a data frame cannot be identified by the ToS voice data or ToS signaling data, the Signaling IP Address is checked. Matched frames are loaded on a High Priority Software Queue and a Nominal Hardware Queue.

Signaling IP Mask

The mask associated with the Signaling IP Address described above.

KeepAlive Up Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Up Count is the number of requests sent from the port when the port transitions from Down to Up, to verify the status of the port. The valid range is 0-255.

KeepAlive Down Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Down Count is the number of requests sent from the port when the port transitions from Up to Down, to verify the status of the port. This only displays if the port is using PPP as its encapsulation type. The valid range is 0-255.

KeepAlive Timeout

The number of 100 millisecond increments between generated echo message requests. This only displays if the port is using PPP as its encapsulation type. The valid range is 0-255.

Loopback Timeout

Sets the transition time between proprietary messages sent over the link. These messages are analyzed to determine whether the link is in a loopback state. This only displays if the port is using Frame Relay as its encapsulation type. The valid range is 0-255.

Viewing Configuration Parameters for the WSX

You can view all current parameters for a WSX port or an individual virtual circuit using the **wpview** command. These parameters will be either the default parameters or parameters you modified using the **wpmodify** command or network management software.

You have a choice of viewing parameters at the chassis, slot or port level. You receive different configuration choices depending upon which level you choose. The sections below describe both ways to use the **wpview** command.

Viewing Parameters for all Submodules in the Chassis

To view port parameters for all submodule boards in a chassis, enter the following command

```
wpview
```

or

```
wpv
```

A screen similar to following displays. In this example, the port parameters being displayed are for a system that contains a 2-port BRI submodule in slot 3.

Slot/Port	Port Type	Intf. Type	Admin/ Oper/ State	Protocol	Speed BPS	Clocking
=====	=====	=====	=====	=====	=====	=====
3/1	Serial	*NONE*	UP/DN	FR	0	External
3/2	ISDN	ISDN-ST	UP/UP	PPP	N/A	External

This screen lists the current values for the listed parameters.

For **Port Type**, **Intf. Type** and **Oper/State**, these parameters are the same as those set through the **wpmodify** command. For detailed information on these values, see *Modifying a Port* on page 28-14. For **Protocol**, **Speed BPS** and **Clocking**, these parameters are the same as those set through the **wpstatus** command. See *Obtaining Status and Statistical Information* on page 28-38.

Viewing Parameters for all Ports in a Single Submodule

To view port parameters for all ports on a particular submodule, enter the **wpview** command, followed by the number of the slot. In the following three examples, the port parameters are displayed for an ISDN-BRI board, a serial board, and a T1 board.

ISDN-BRI Board Example

To display the parameters for all ports on the ISDN-BRI board (in slot 3), enter:

```
wpview 3
```

or

```
wpv 3
```

A screen similar to following displays:

Port	PortType	Intf. Type	Admin/ Oper/ State	Protocol	Speed BPS	Clocking
1	ISDN	ISDN-ST	UP/UP	PPP	N/A	N/A

Serial Board Example

To display the parameters for all ports on the serial board (in slot 3), enter:

```
wpview 3
```

or

```
wpv 3
```

A screen similar to following displays:

Port	PortType	Intf. Type	Admin/ Oper/ State	Protocol	Speed BPS	Clocking
1	Serial	V35DCE	UP/UP	PPP	2048000	Split

T1 Board Example

To display the parameters for all ports on the T1 board (in slot 3), enter:

```
wpview 3
```

A screen similar to following displays:

Slot/Port	PortType	Intf. Type	Admin/ Oper/ State	Protocol	Speed BPS	Clocking
3/1	T1	T1	UP/UP	FR	1544000	Loop

◆ Note ◆

E1 boards provide a similar display, except the port type and interface type display as **E1** and speed displays as **2048000**.

Viewing Port Parameters

To view port parameters, enter the following command:

```
wpview 3/<port>
```

where **3** is the slot number for WAN uplinks, and **<port>** is the port number for which you want to view information (either **1** or **2**). The following three examples show the configuration setup screens for a fractional T1 port, a universal serial port, and an ISDN-BRI port. The display is slightly different depending upon the encapsulation type, either Frame Relay or PPP.

Fractional T1 Port Example

The following example displays the configuration view screen for a fractional T1 port (port 1) using Frame Relay. To view 3/1, enter:

```
wpview 3/1
```

or

```
wpv 3/1
```

A screen similar to following displays:

```
Configuration View for Slot 3, Port 1.
1) Admin Status ..... UP
2) Protocol Type ..... Frame Relay
3) T1/E1 Starting Time Slot ..... 1
4) T1/E1 Number of Time Slots ..... 24
8) TOS for Voice Data ..... a0
9) TOS for Voice Signaling Data ..... 60
10) TOS Mask for both TOS Value ..... e0
11) Signaling IP Address ..... 0.0.0.0
12) Signaling IP Mask ..... 255.255.255.255
15) Loopback Timeout ..... 10
```

This next example displays the configuration view screen for a fractional T1 port (port 1) using PPP. To view 3/1, enter:

```
wpview 3/1
```

or

```
wpv 3/1
```

A screen similar to following displays:

```
Configuration View for Slot 3, Port 1.
1) Admin Status ..... UP
2) Protocol Type ..... Frame Relay
3) T1/E1 Starting Time Slot ..... 1
4) T1/E1 Number of Time Slots ..... 24
8) TOS for Voice Data ..... a0
9) TOS for Voice Signaling Data ..... 60
10) TOS Mask for both TOS Value ..... e0
11) Signaling IP Address ..... 0.0.0.0
12) Signaling IP Mask ..... 255.255.255.255
13) KeepAlive Up Count ..... 0
14) KeepAlive Down Count ..... 0
15) KeepAlive Timeout ..... 10
16) DTR Pulse Width ..... 0
17) DTR Pulse Count ..... 0
```

◆ Note ◆

The DTR Pulse setting do not apply to T1 and E1 interfaces, and are not described below.

Admin Status

The options for the Admin Status are **UP** and **DN**. If **UP**, the port has been enabled and can transmit data as long as its Operational Status is also **UP**. If set to is **DN**, the port will not pass data even if its physical connection is good.

Protocol Type

The protocol type can be set to either Frame Relay or Point to Point Protocol (PPP). The default setting is Frame Relay.

T1/E1 Starting Time Slot

This field specifies the first time slot number to use on a T1 or E1 port. For a full T1 or E1 connection, specify time slot 1. For a fractional T1 or E1 connection, set this field to the starting time slot number as specified by your service provider.

T1/E1 Number of Time Slot

This field specifies the total number of 64 kbps time slots to use on the T1 or E1 connection. For a full T1, set this number to 24. For a full E1 connection, set this number to 30 if you are running multiframe, or 31 if you are not. For fractional T1 or E1, you must set the number of time slots to the value specified by your service provider. For example, a 256 kpbs service uses four time slots (4 x 64 = 256).

TOS for Voice Data

Shows the priority for voice data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS for Voice Signaling Data

Shows the priority for voice signaling data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS Mask for both TOS Value

Shows the mask bits for both voice data and signaling data.

Signaling IP Address

When a data frame cannot be identified by the ToS voice data or ToS signaling data, the Signaling IP Address is checked. Matched frames are loaded on a High Priority Software Queue and a Nominal Hardware Queue.

Signaling IP Mask

The mask associated with the Signaling IP Address described above.

KeepAlive Up Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Up Count is the number of requests sent from the port when the port transitions from Down to Up, to verify the status of the port.

KeepAlive Down Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Down Count is the number of requests sent from the port when the port transitions from Up to Down, to verify the status of the port. This only displays if the port is using PPP as its encapsulation type.

KeepAlive Timeout

The number of 100 millisecond increments between generated echo message requests. This only displays if the port is using PPP as its encapsulation type.

Loopback Timeout

Sets the transition time between proprietary messages sent over the link. These messages are analyzed to determine whether the link is in a loopback state. This only displays if the port is using Frame Relay as its encapsulation type.

Universal Serial Port Example

The following example displays the configuration view screen for a universal serial port (port 2). To view 3/2, enter:

```
wpview 3/2
```

or

```
wpv 3/2
```

If the serial port is using Frame-Relay, a screen similar to following displays:

```
Configuration View for Slot 3, Port 2.
1) Admin Status ..... UP
2) Speed in BPS ..... 2048000
3) Clocking ..... Split
4) Protocol Type ..... Frame Relay
7) Receive Clock ..... Normal
8) TOS for Voice Data ..... a0
9) TOS for Voice Signaling Data ..... 60
10) TOS Mask for both TOS Value ..... e0
11) Signaling IP Address ..... 0.0.0.0
12) Signaling IP Mask ..... 255.255.255.255
15) Loopback Timeout ..... 0
```

If the serial port is using PPP, a screen similar to following displays:

```
Configuration View for Slot 3, Port 2.
1) Admin Status ..... UP
2) Speed in BPS ..... 2048000
3) Clocking ..... Split
4) Protocol Type ..... Frame Relay
7) Receive Clock ..... Normal
8) TOS for Voice Data ..... a0
9) TOS for Voice Signaling Data ..... 60
10) TOS Mask for both TOS Value ..... e0
11) Signaling IP Address ..... 0.0.0.0
12) Signaling IP Mask ..... 255.255.255.255
13) KeepAlive Up Count ..... 0
14) KeepAlive Down Count ..... 0
15) KeepAlive Timeout ..... 10
16) DTR Pulse Width ..... 0
17) DTR Pulse Count ..... 0
20) Connection Function ..... Dedicated
```

Admin Status

The options for the Admin Status are **UP** and **DN**. If **UP**, the port has been enabled and can transmit data as long as its Operational Status is also **UP**. If set to **DN**, the port will not pass data even if its physical connection is good.

Speed in BPS

This field displays the access rate for the Frame Relay line to the service provider. This parameter is the speed of the entire connection, not an individual virtual circuit. For example, if you have a 56 kbps line to your service provider, this field should be set to 56000. A full T1 line would have an access rate of 1,544,000 bps, and a full E1 line would have an access rate of 2,048,000 bps. For either T1 or E1, you can also have a fractional service with an access rate that is a multiple of 64 kbps.

Clocking

This field displays either **External**, **Internal**, or **Split**. For a more detailed discussion of clocking, see *Clocking* under *Modifying a Port* on page 28-14.

Receive Clock

Often, due to delays added to timestamps in when running through switch hardware, the receive clock time is significantly different than expected from the transmitting data source. To correct the problem, it is possible to set the receive clock to invert the delay information. The following options are available:

Protocol Type

The protocol type can be set to either Frame Relay or Point to Point Protocol (PPP). The default setting is Frame Relay.

TOS for Voice Data

Shows the priority for voice data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS for Voice Signaling Data

Shows the priority for voice signaling data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS Mask for both TOS Value

Shows the mask bits for both voice data and signaling data.

Signaling IP Address

When a data frame cannot be identified by the ToS voice data or ToS signaling data, the Signaling IP Address is checked. Matched frames are loaded on a High Priority Software Queue and a Nominal Hardware Queue.

Signaling IP Mask

The mask associated with the Signaling IP Address described above.

KeepAlive Up Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Up Count is the number of requests sent from the port when the port transitions from Down to Up, to verify the status of the port.

KeepAlive Down Count

If the switch detects that a port may be down, it will generate echo message requests to the far end of the connection. The KeepAlive Down Count is the number of requests sent from the port when the port transitions from Up to Down, to verify the status of the port. This only displays if the port is using PPP as its encapsulation type.

KeepAlive Timeout

The number of 100 millisecond increments between generated echo message requests. This only displays if the port is using PPP as its encapsulation type.

DTR Pulse Width

A Data Terminal Ready (DTR) Pulse is sent at the hardware level to determine a port is still synchronized with its far end connection. The Pulse Width is the number of 100 milli-second increments that the pulse lasts. This only displays if the port is using PPP as its encapsulation type.

DTR Pulse Count

A Data Terminal Ready (DTR) Pulse is sent at the hardware level to determine a port is still synchronized with its far end connection. The Pulse Count is the number of pulses generated when a line is down. This only displays if the port is using PPP as its encapsulation type.

Loopback Timeout

Sets the transition time between proprietary messages sent over the link. These messages are analyzed to determine whether the link is in a loopback state. This only displays if the port is using Frame Relay as its encapsulation type.

Connection Function

On serial ports using PPP, it is possible to configure the port to be a DTR dial backup port. In case the primary WAN port fails, a DTR dial backup port will use a synchronous modem to dial out and reestablish the WAN connection. If this feature is disabled, the **Connection Function** field will read **Dedicated**. If it is enabled, the field will read **DTR-Dial**, and the following options will also be displayed:

200) Connect Timeout (seconds) 60
201) Retry Delay (seconds) 10

Connect Timeout

The number of seconds the switch attempts to establish a connection via this port before declaring the attempt a failure. The valid range is 10 to 2147483647. The default is 60. If the value is entered as 0, the attempt will never timeout.

Retry Delay

The number of seconds after a connection failure the switch waits before attempting a new connection. The valid range is 1 to 2147483647. The default is 1.

ISDN-BRI Port Example

The following example displays the configuration view screen for an ISDN-BRI port (port 2). To view 3/2, enter:

```
wpview 3/2
```

or

```
wpv 3/2
```

A screen similar to following displays:

```
Configuration View for Slot 3, Port 2.
1) Admin Status ..... UP
2) Speed in BPS ..... 2048000
3) Clocking ..... Split
4) Protocol Type ..... Frame Relay
8) TOS for Voice Data ..... a0
9) TOS for Voice Signaling Data ..... 60
10) TOS Mask for both TOS Value ..... e0
11) Signaling IP Address ..... 0.0.0.0
12) Signaling IP Mask ..... 255.255.255.255
13) KeepAlive Up Count ..... 0
14) KeepAlive Down Count ..... 0
15) KeepAlive Timeout ..... 10
16) DTR Pulse Width ..... 0
17) DTR Pulse Count ..... 0
```

◆ Note ◆

The ISDN **wpview** menu displays PPP specific line options described in the section *Modifying a Port* on page 28-14. However, they do not apply to an ISDN port, and are not described below.

Admin Status

The options for the Admin Status are **UP** and **DN**. If **UP**, the port has been enabled and can transmit data as long as its Operational Status is also **UP**. If set to **DN**, the port will not pass data even if its physical connection is good.

Speed in BPS

This field displays the access rate for the Frame Relay line to the service provider. This parameter is the speed of the entire connection, not an individual virtual circuit. For example, if you have a 56 kbps line to your service provider, this field should be set to 56000. A full T1 line would have an access rate of 1,544,000 bps, and a full E1 line would have an access rate of 2,048,000 bps. For either T1 or E1, you can also have a fractional service with an access rate that is a multiple of 64 kbps.

Clocking

This field displays either **External**, **Internal**, or **Split**. For a more detailed discussion of clocking, see *Clocking* under *Modifying a Port* on page 28-14.

Protocol Type

The protocol type can be set to either Frame Relay or Point to Point Protocol (PPP). The default setting is Frame Relay.

TOS for Voice Data

Shows the priority for voice data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS for Voice Signaling Data

Shows the priority for voice signaling data streams. The value must be entered in hexadecimal format translated from binary, and can use either IP Precedence or Differentiated Services Code Point (DSCP). See *Type of Service (ToS)* on page 28-2 above for a more detailed explanation of ToS.

TOS Mask for both TOS Value

Shows the mask bits for both voice data and signaling data.

Signaling IP Address

When a data frame cannot be identified by the ToS voice data or ToS signaling data, the Signaling IP Address is checked. Matched frames are loaded on a High Priority Software Queue and a Nominal Hardware Queue.

Signaling IP Mask

The mask associated with the Signaling IP Address described above.

Deleting Ports

The **wpdelete** command allows you to delete configuration information for a WSX port. When you delete a this information, all WAN configuration parameters for the selected port revert back to default settings.

To delete a port configuration, enter the following command:

```
wpdelete slot/port
```

in which **slot** is the slot number for the WSX board and **port** is the port number on the WSX board that you want to delete. For example, to delete port 1 on the WSX board in slot 2, enter:

```
wpdelete 2/1
```

or

```
wpd 2/1
```

This system returns the following prompt to confirm the deletion:

```
This will delete Slot 2, Port 1. Continue? {(Y)es, (N)o} (N)
```

Enter a **Y** to confirm the deletion or press **Enter** to cancel the deletion.

Note

The **wpdelete** command requires that you indicate a slot and port number. For example,

```
wpdelete
```

would be an incorrect usage, whereas,

```
wpdelete 4/2
```

would be correct.

Obtaining Status and Statistical Information

You can obtain general and detailed WAN port statistical information on all WSX boards in the switch, a single WSX board, individual ports, and Frame Relay and PPP protocols. The **wpstatus** command is used to provide this information. This information includes types of physical interface, access rate of the Frame Relay line, and errors. In addition, the **wpstatus** command can display the number of frames received and transmitted.

Obtaining Information on All Boards in a Switch

To obtain status information on all WSX boards in a switch, you enter the **wpstatus** command without any parameters as follows:

```

wpstatus
or
wps

```

This command displays a screen similar to the following (In this example, the port parameters being displayed are for a system that contains a 2-port WSX-BRI module in slot 4, an 8-port WSX module in slot 5, and a 2-port WSX in slot 8.):

Slot/Port	PortType	Intf. Type	Admin/	Protocol	BPS	Speed	Utilization		
			Oper/			Clocking	10s	1m	5m
=====	=====	=====	=====	=====	=====	=====	=====	=====	=====
4/1	Serial	*NONE*	UP/DN	FR	EXT CLK	External	10%	10%	10%
4/2	ISDN	ISDN-ST	UP/DN	PPP	N/A	External	40%	30%	60%
5/1	Serial	V35DCE	UP/UP	PPP	2048000	Split	30%	60%	50%
5/2	Serial	V35DCE	UP/UP	FR	2048000	Split	100%	50%	70%
5/3	Serial	X21DCE	UP/DN	FR	2048000	Split	90%	80%	60%
5/4	Serial	V35DCE	UP/UP	FR	2048000	Split	20%	50%	50%
5/5	Serial	*NONE*	UP/DN	FR	EXT CLK	External	30%	30%	50%
5/6	Serial	*NONE*	UP/DN	FR	EXT CLK	External	100%	50%	80%
5/7	Serial	*NONE*	UP/DN	FR	EXT CLK	External	70%	50%	50%
5/8	Serial	*NONE*	UP/DN	FR	EXT CLK	External	100%	80%	30%
8/1	T1	T1	UP/UP	FR	1544000	External	80%	50%	70%
8/2	Serial	530DCE	UP/UP	FR	2048000	Split	10%	50%	40%

Each row in the table corresponds to a physical port on a WSX board in the switch. The following sections describe the columns shown in this table:

Field Descriptions

The following section explains the fields and their corresponding values.

Slot/Port

The first number in this column is the slot in the switch where this WSX is installed. The second number is the port number on the WSX.

Port Type

This column shows

- Serial
- ISDN
- T1
- E1

Intf Type

This column indicates the physical cable type connected to this port. This cable type is automatically sensed by the WSX hardware. This column indicates the cable type and whether it is DCE or DTE. The following values may appear in this column:

- **V35DTE** (V.35 DTE cable)
- **V35DCE** (V.35 DCE cable)
- **232DTE** (RS-232 DTE cable)
- **232DCE** (RS-232 DCE cable)
- **X21DTE** (X.21 DTE cable)
- **X21DCE** (X.21 DCE cable)
- **530DTE** (RS-530 or RS-449 EIA DTE cable)
- **530DCE** (RS-530 or RS-449 EIA DCE cable)
- **T1**
- **E1**
- **ISDN-ST**
- **ISDN-U**

The WSX sees RS-530 and RS-449 cables the same because they are electrically identical. However, this does not affect the operation of either cable type. Both RS-530 and RS-449 cables are supported.

If no cable is connected to a universal serial port, then this column will display:

NONE

If an error has been detected on the port (e.g., cable type could not be detected), the following value displays:

ERROR!

Admin/Oper State

This column shows the Administrative and Operational State of this WSX port. The value before the slash refers to the Admin Status. If **UP**, the port has been enabled and can transmit data as long as its Operational State is also **UP**. If the Admin Status is **DN**, the port will not pass data even if its physical connection is good.

The value after the slash refers to the Operational State. If **UP**, the port is capable of passing data as long as it has been logically enabled at the Administrative level. If **DN**, the port cannot pass data due to a problem in the physical connection (e.g., cable disconnected, WSX could not detect cable type) or because the port is administratively down. If the Operational State displays **LB**, the port is currently in Loopback (test) mode.

Protocol

The protocol type can be set to either Frame Relay or Point to Point Protocol (PPP).

Speed BPS

This column indicates the speed, or access rate, between the WSX serial port and DSU or other physical DTE device. The speed is expressed in bits per second (bps). This speed is the total bandwidth available on the line connected to this port. Virtual circuits on this port share this bandwidth.

Usually, the WSX port will be a physical DTE device and the speed will be determined by the DSU. In this case, this value will read **EXT CLK**, which means the WSX port gets its clocking from an externally attached DCE device (i.e., DTE cable plugged into WSX port) or no cable is attached. If the WSX port is a physical DCE device (i.e., DCE cable plugged into WSX port), then this value will be the actual clock rate used by the port.

Clocking

Indicates the type of clocking used on this port. The three types of clocking are described in *Clocking* on page 28-17.

Utilization

Indicates the amount of port usage, expressed in bandwidth percentage, over three durations: the previous ten seconds (**10s**), the previous minute (**1m**), and the previous five minutes (**5m**).

Obtaining Information on the Ports for a Single WSX Board

To obtain status information on a single WSX board, enter the **wpstatus** command and the slot number for the WSX board, as follows:

```
wpstatus slot
```

where **slot** is the slot number where the WSX board is installed. For example, if you wanted to obtain status information for the board in slot 4. In the following three examples, the port parameters being displayed are for a system that contains a 2-port WSX ISDN-BRI board in slot 4, an 8-port WSX serial board in slot 5, and a 2-port WSX T1 board in slot 8.)

ISDN-BRI Board Example

In this example, the board in slot 4 is a 2-port ISDN-BRI WSX board. To view the status of slot 4, enter:

```
wpstatus 4
```

or

```
wps 4
```

This command displays a screen similar to the following:

```
WAN Port Status for slot: 4
```

PT	Admin/ Oper Status	Intf Type	Speed BPS	Frames In	Frames Out	Octets In	Octets Out
1	UP/DN	*NONE*	EXT CLK	0	0	0	0
2	UP/DN	ISDN-ST	N/A	0	0	0	0

/Interface/WAN %

Each row in the table corresponds to a port on the WSX you requested information on.

8-Port WSX Board Example

In this example, the board in slot 5 is an 8-port WSX board. To view the status of slot 5, enter:

```
wpstatus 5
```

or

```
wps 5
```

This command displays a screen similar to the following:

WAN Port Status for slot: 5

PT	Admin/ Oper Status	Intf Type	Speed BPS	Frames In	Frames Out	Octets In	Octets Out
1	UP/UP	V35DCE	2048000	3	17	36	276
2	UP/UP	V35DCE	2048000	175	926	2034	25617
3	UP/UP	X21DCE	2048000	123	931	1722	55717
4	UP/UP	V35DCE	2048000	776	189	14430	7531
5	UP/DN	*NONE*	EXT CLK	0	0	0	0
6	UP/DN	*NONE*	EXT CLK	0	0	0	0
7	UP/DN	*NONE*	EXT CLK	0	0	0	0
8	UP/DN	*NONE*	EXT CLK	0	0	0	0

/Interface/WAN %

2-Port Fractional T1 WSX Board Example

In this example, the board in slot 8 is a 2-port Fractional T1 WSX board. To view the status of slot 8, enter:

```
wpstatus 8
```

or

```
wps 8
```

This command displays a screen similar to the following:

```
/Interface/WAN % wps 8
```

WAN Port Status for slot: 8

PT	Admin/ Oper Status	Intf Type	Speed BPS	Frames In	Frames Out	Octets In	Octets Out
1	UP/DN	T1	1544000	0	0	0	0
2	UP/UP	530DCE	2048000	45695	47761	10596229	2560992

/Interface/WAN %

Field Descriptions

The following section explains the fields and their corresponding values.

PT

The port number on the WSX board for which statistics are displayed.

Admin/Oper Status, Int Type, Speed Bps

These columns are described in the section, *Obtaining Information on All Boards in a Switch* on page 28-38. Please refer to this section for detailed information.

Frames In

The total number of frames received on this port since the last time the switch was initialized.

Frames Out

The total number of frames sent on this port since the last time the switch was initialized.

Octets In

The total number of octets, or bytes, received on this port since the last time the switch was initialized. This statistic includes the data and Frame Relay or PPP header fields, but does not include CRC or flag characters.

Octets Out

The total number of octets, or bytes, sent on this port since the last time the switch was initialized. This statistic includes the data and Frame Relay or PPP header fields, but does not include CRC or flag characters.

Viewing Information on a Single Port

To obtain status information on a single WSX port, enter the **wpstatus** command, followed by the slot number for the WSX board and the port number for which you want to receive information, as follows:

```
wpstatus <slot>/<port>
```

or

```
wps <slot>/<port>
```

where **<slot>** is the slot number where the WSX board is installed and **<port>** is the port number on the WSX board.

Frame Relay Example

In the following example, port 1 on slot 4 is configured for Frame Relay. To obtain status information for this port, enter:

```
wpstatus 4/1
```

A screen similar to the following will be displayed:

```

Frame Relay Status for slot 4, port 1:

Applicable to all port types.
Administrative/Operational Status .....Up/Up
Port Type.....Universal Serial Port
Protocol.....Frame Relay

Physical Level Information.
Speed      Intf.      Receive      Receive      Receive      Transmit      Signal
BPS        Type      CRC Errors   Aborts       Overruns     Overruns     Errors
=====
2048000    V35DCE    0            0            0            0            0

Displays for serial ports only
Control    DTR      RTS      DSR      CTS      DCD
Signal     ON       ON       ON       ON       OFF

Logical (Frame Relay) Information
Frame Relay Information:
UniCast   Discarded   Error
Octets    Frames     Frames     Count
=====
IN         941079     0          0
Out        21334     0          0
IN+OUT    962413     0          0

Administrative/Operational Phase .... Up/Up

Last Error Type .....No Error Since Reset
Last Error Time .....0 days, 00:00:00
Interface failures .....0
Last interface failure time .....0 days, 00:00:00

Virtual Circuit Level Information
DLCI Information:
Admin/
DLCI Oper  DLCI      Frames     Frames     Octets     Octets
Num  Status  Type      In         Out        In         Out
====
0    UP/UP   Configured  1021      1021      16044     1494
31   UP/UP   Learned    17716     136       2746651   12663
32   UP/DN   Learned     0         0         0         0
    
```

This command displays three (3) layers of information. The top section provides information on the physical interface. The middle section provides information on the logical, or Frame Relay, interface. The bottom section provides information on the virtual circuits associated with this physical port.

For detailed descriptions of the fields, refer to Chapter 29 “Managing Frame Relay.”

PPP Example

In the following example, port 1 on slot 4 is configured for Point-To-Point Protocol (PPP). To obtain status information for this port, enter:

wpstatus 5/1

A screen similar to the following will display:

```

/Interface/WAN % wps 5/1
WAN Port Status for slot 5, port 1:
Administrative/Operation Status: ..... UP/UP
Port Type ..... Universal Serial Port
Protocol ..... PPP

Speed      Intf.      Receive      Receive      Receive      Transmit      Signal
BPS        Type      CRC Errors   Aborts       Overruns     Underruns     Errors
=====
2048000    V35DCE           0           0           0           0           0

Control    DTR  RTS  DSR  CTS  DCD
Signals    ON  ON  ON  ON  ON

PPP Management Statistics:

Admin      IP      IPX      BCP      CCP
Status     Mode   Oper    Oper    Oper    Oper
=====
UP         Normal Open    Close   Open    Open

LCP Pkts   IPCP Pkts  IPX Pkts  BCP Pkts  CCP Pkt
IN/OUT     IN/OUT     IN/OUT    IN/OUT    IN/OUT
=====
3/4        2/2        4/0       2/2       3/3

          Packets  Packets  Packets  Octets  Octets  %In  %Out
          In      Out      In+Out  In      Out      %In  %Out
          =====
Total          284    5809    6093    100333  344187
Ethernet       0     1337    1337      0    157846      0    45
8025           0         0         0         0         0      0     0
FDDI           0         0         0         0         0      0     0
IP            281     282     563    100216  22931     99     6
IPX           0         0         0         0         0      0     0
BPDU           3    4190    4193     117   163410      0    47

STAC-LZS      Compressed  Compressed  Uncompressed  Compression
Compression:  Frames      Octets      Octets         Ratio
=====
In              284         8635         100333        11.6:1
Out             5809        96794        449230         4.6:1
In+Out          6093       105429        549563         5.2:1

/Interface/WAN %

```

◆ Note ◆

The section devoted to compressed data traffic statistics will be displayed only if the port has been configured for STAC-LZS compression.

For detailed descriptions of the fields, refer to Chapter 30, “Point-to-Point Protocol.”

Configuring 31 Timeslots on a WAN E1 Port

On WSX E1 ports, the unframed format is not supported since WSXs only support standard E1 framing for PPP or Frame Relay (the “unframed” format is only supported for unstructured Circuit Emulation T1 or E1 ports). WSX E1 ports *must* be set to one of the standard E1 Framing types (E1, E1-CRC, E1-MF, E1-MF-CRC) with the **temod** command. (See Chapter 33, “Managing T1 and E1 Ports,” for more information on the **temod** command.)

Most E1 services only allow a maximum of 30 usable timeslots since timeslot 0 is always used for Frame Synchronization (which is why you cannot use unframed for Frame Relay or PPP ports since you *must* specify how timeslot 0 is used) and timeslot 16 is usually used for multiframe sequencing.

The WSX can support 31 timeslots for cases where timeslot 16 is not used for multiframe control. When you configure the timeslots for a WSX E1 port, you specify a starting timeslot followed by a number of timeslots by using the **wpmodify** command. (See *Modifying a Port* on page 28-14 for more information on the **wpmodify** command.)

Normally, the WSX will use a default configuration that skips timeslot 16 automatically. In this way, it will select the E1 frame to generate E1 timeslot 0 (the “synchronization” timeslot), but leave timeslot 16 (the “multiframe control” timeslot) free. The WAN port configuration software when configured for 31 timeslots will then use all timeslots from 1 to 31 to give you a full E1 where timeslot 16 is also used for data. Again, this should only be done for facilities that do not require E1 Multi-Frame. For those types of E1 lines, they can support a maximum of 30 timeslots. Only those E1 lines that do not require E1 multiframe can be configured in the method described below.

To configure a WAN E1 port for 31 timeslots, follow the steps below:

1. Enter **temod <slot>/<port>** at the system prompt, where **<slot>** is the slot number of the module with the E1 port and **<port>** is the port number of the E1 port. For example, to configure WSX E1 port 4/2, enter **temod 4/2**.
2. Enter **2=4** at the prompt to set the frame type to E1 or enter **2=5** at the prompt to set the frame type to E1-CRC.
3. Enter **save** at the prompt to save your settings.
4. Enter **wpmodify <slot>/<port>** or **wpm <slot>/<port>** at the system prompt, where **<slot>** is the slot number of the module with the E1 port and **<port>** is the port number of the E1 port. For example, to configure WSX E1 port 4/2, enter **wpm 4/2**. (Note: **wpm** is the abbreviated form of **wpmodify**.)
5. Enter **3=1** to set the starting timeslot to 1.
6. Enter **4=31** to set the number of timeslots to 31.
7. Enter **save** at the prompt to save your settings.

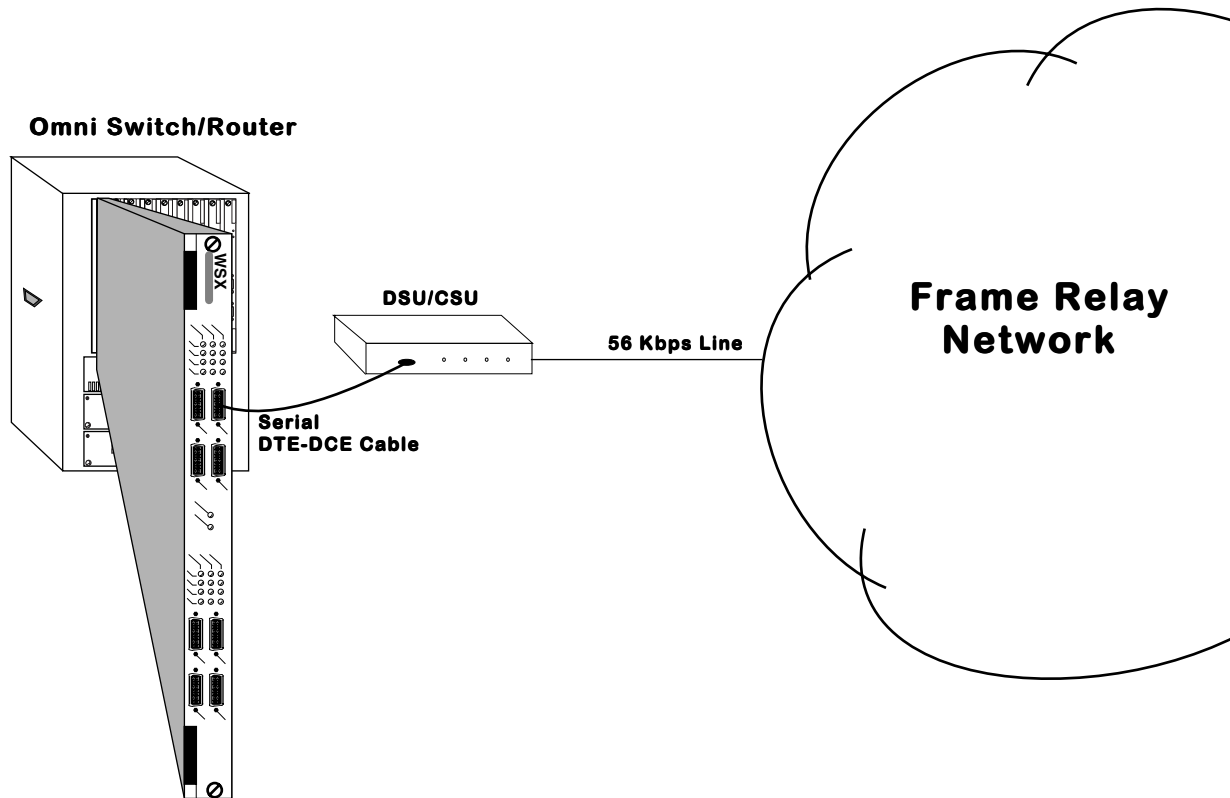
29 Managing Frame Relay

The WAN Switching Module (WSX) family supports Frame Relay on universal serial, T1 or E1 ports. Management, data handling, compression, and multi-protocol encapsulation are compatible with current Frame Relay standards, such as RFC 1490 and FRF.9. The WSX supports all three major DLCMI management protocols.

WSX frame relay extends the power and flexibility of LAN switching over large geographic distances using a Frame Relay network or a leased line, such as a T1. In a Frame Relay network configuration, the WSX provides a cost effective link supporting multiple virtual circuits. In a leased line configuration, the WSX provides dedicated bandwidth to a single remote site.

VLAN architectures are preserved and consistent on both sides of a WAN link. The WSX supports frame relay trunking, so VLAN Groups on one side of a Frame Relay link are compatible with those on the other side. In addition, the WSX is capable of Frame Relay IP and IPX routing and complies with Inverse Address Resolution Protocol (InARP) RFC 1293.

In a typical configuration, the WSX occupies one slot in an Omni Switch/Router. Since it is compatible with Omni Switch/Router any-to-any switching and VLAN architecture, you can switch other topologies in the LAN to Frame Relay. The following diagram shows a typical WSX setup using a 56 Kbps Frame Relay line (up to 2 Mbps access rates are supported).



Typical WSX Frame Relay Setup Using Serial Ports

The WSX supports automatic detection of cable types attached to universal serial ports. It also supports three types of DLCMI management: LMI Rev. 1.0, ANSI T1.617 Annex D, and CCITT/ITU-T Q.933 Annex A.

Software in the switch allows you to configure access rate, clocking, DLCMI type, compression, and congestions controls, such as the Committed Information Rate (CIR). Additional software commands allow you to view the status of the Frame Relay connection at the WSX board, port, or virtual circuit level. Extensive statistics are provided at each level, including a breakdown of traffic by frame type (Ethernet, IP, IPX, or BPDU) at the virtual circuit level.


The WSX is designed to require as little configuration as possible. It senses the cable type installed and automatically maps virtual circuits to virtual ports as soon as you plug in the cable. The WSX supports 256 Permanent Virtual Circuits (PVCs), which is equivalent to the number of virtual ports allowed in an Omni Switch/Router.

In addition, you can set up a default bridging and a default routing Group. Virtual circuits are automatically assigned to these Groups as soon as they are configured or learned, which means Frame Relay frames can be bridged or routed without user-configuration.

Back-to-Back Frame Relay Configurations

Frame Relay switching modules may be connected “back-to-back” without an intervening Frame Relay network or switch. Such connections are made by using private leased lines, such as T1 lines, instead of public Frame Relay networks usually over large geographic distances.

No special user configuration is necessary for back-to-back connections. The WSX software automatically detects that a Frame Relay Logical DCE (i.e., Frame Relay switch) is not present and that there is another Frame Relay Logical DTE (i.e., another WSX, FRAD, bridge/router) on the other end of the WAN connection. The WSX then automatically brings up a Permanent Virtual Circuit identified with a DLCI of 32, which is the same value IBM uses in this scenario. The WSX does not bring up PVC DLCI 32 until it knows that it has established communication with another DTE device rather than a Frame Relay switch.



Back-to-Back Frame Relay Configuration Using Serial Ports

Universal Serial Port Cable Interfaces

The WSX automatically senses the cable type that you plug into one of its universal serial ports. It can sense whether the cable type is DCE or DTE and whether it is one of the following interfaces:

- RS-232
- RS-449
- RS-530
- V.35
- X.21 (European)

All cable types, except RS-232, are capable of access rates from 9.6 Kbps to 2 Mbps. The RS-232 cable is not compatible with speeds greater than 64 Kbps. Each cable type is illustrated and described in Appendix D, “Custom Cables.”

The WSX serial port is normally considered a physical DTE device. It is possible to turn it into a physical DCE device simply by plugging in a DCE cable. The WSX board internally senses whether a DCE or DTE cable is connected.

DTE/DCE Type and Transmit/Receive Pins

The RS-232 protocol, which is employed at the physical level for all cable types, always defines Transmit and Receive pins in relation to the DTE. So, the type of cable you attach (DCE or DTE) determines the direction of data flow on your connector's Transmit and Receive pins.

If the WSX serial port is a physical DTE, which is probably the most common configuration, then data is received on Receive pins and transmitted on Transmit pins. If you are using a WSX port as a physical DCE, then data is *transmitted* on the Receive pins and *received* on the Transmit pins.

“Physical” and “Logical” Devices

This chapter refers to “physical” and “logical” DTE (Data Terminal Equipment) and DCE (Data Communication Equipment) devices. A physical device operates on the network layer, and is normally an actual piece of hardware, such as a WSX or CSU/DSU. Physical devices may further be differentiated as DTE and DCE devices. A physical DTE device would be a piece of hardware, such as a WSX, that does not control the access rate for virtual circuits. The physical DTE device is a conduit for data traffic but not a controller of data traffic. A physical DCE device is hardware, such as a CSU/DSU, that does control access rates of Frame Relay traffic. Normally physical DTE and DCE devices are directly connected to one another.

Logical devices operate on the Frame Relay protocol layer, and are sometimes referred to as “Frame Relay logical” devices. Logical devices can also be broken down into DTE and DCE devices. Logical DTE devices, again like the WSX, do not have direct control over the Frame Relay network and the various congestion and control parameters that govern it. Logical DTE devices do not control such actions as bringing up and tearing down virtual circuits; they act upon updates and commands generated by the Frame Relay network. Logical DCE devices, such as a Frame Relay switch, have a large span of control over Frame Relay network traffic. They bring up and tear down virtual circuits, set congestion control bits in packets, and communicate status to logical DTE devices.

Compression

Data compression allows you to get more data through the Frame Relay pipeline, further enhancing cost benefits. A typical data compression ratio on the WSX board at the hardware level is 4:1. In addition, the compression processor (STAC 9705) has its own DRAM that can store up to 100 virtual circuits (on a 4-port WSX) without performance degradation. An 8-port WSX can store up to 200 virtual circuits without performance degradation. Support for more than 100 compressed VCs (or 200 VCs on an 8-port WSX) is possible through swapping within memory, but compression performance may decrease at these levels.

The WSX will only compress data if you enable Compression Negotiation through software and the Bridge/Router on the other end of the Frame Relay virtual circuit supports standard FRF.9 compression. (An Omni Switch/Router-to-Omni Switch/Router connection would support compression.) Negotiation is necessary because if compressed data is sent to a Bridge/Router that does not support compression, then this Bridge/Router will not recognize the data and will automatically drop the unrecognizable frames.

If you enable Compression Negotiation, the WSX will query the Frame Relay device on the other end of the circuit (according to FRF.9 specifications) to see if it supports compression. If it does, then the WSX compresses all data except DLCMI (management) data. If it doesn't, then data on that virtual circuit is sent uncompressed. See *Setting Configuration Parameters* on page 29-22 for information on enabling compression.

Note

Compression is not supported on the 2 universal serial port Omni Switch/Router WSX and Omni Switch/Router WSX modules.

Virtual Circuits and DLCIs

The WSX supports Permanent Virtual Circuits (PVCs), but not Switched Virtual Circuits (SVCs). Most service carriers do not currently offer SVCs. PVCs are either static (configured) or dynamic (learned). Static PVCs are user-configured and consist of Management, or Control, PVCs and any configured Data PVCs. Management VCs are used by the WSX to communicate with the Frame Relay network. Dynamic PVCs are usually data circuits, which are controlled by the Frame Relay network and not configured in advance. A logical Frame Relay DTE device like the WSX does not create or control dynamic data VCs. It is only informed of their status through periodic Status updates from the Frame Relay network.

Each virtual circuit is locally defined by a Data Link Connection Identifier (DLCI). The Frame Relay network assigns the DLCIs and informs the WSX about them.

DLCI numbers from 0 to 15 and 992 to 1023 are reserved for Control VCs. If you are using Annex A or Annex D as your DLCMI, the management control VC will be assigned DLCI 0. If you are using the LMI Revision 1.0 DLCMI, then the management control VC will be assigned DLCI 1023.

DLCI numbers from 16 to 991 are reserved for Data VCs.


You may have up to 256 virtual circuits and up to 128 virtual ports on a WSX.

A VC may or may not have the same DLCI on each side of a WAN link. For example, if a WSX physical port contains three Frame Relay VCs on its local network with DLCIs 16, 17, and 18, these same VCs on the other side of the Frame Relay network might be 30, 31, and 32. The two sets of DLCIs are technically part of the same virtual circuits, but their values may or may not be different. DLCIs are only significant locally.

At any one time, a virtual circuit will be active, inactive or deleted. If a virtual circuit is Active it can transmit and receive data. If it is Inactive, the Frame Relay network still sees the virtual circuit, but there is a problem with it and it is discarding data. If the virtual circuit is Deleted, then the virtual circuit is not transmitting or receiving data and no DLCI exists for it.

WSX Self-Configuration and Virtual Circuits

The following diagram summarizes the self-configuration features of the WSX. This example assumes no configuration parameters are entered for the WSX. Default bridging is set up on Group 1, and no Routing or Trunking are configured.



WSX Initial Port and Virtual Circuit Configuration

After mapping virtual circuits to virtual ports, the WSX is ready to send data. STATUS ENQUIRIES are repeated periodically by the WSX. The intervals between STATUS ENQUIRES can be configured through software. See *Setting Configuration Parameters* on page 29-22 for information on setting these parameters.

Congestion Control

Use of Frame Relay lines tends to be “bursty,” with heavy use at times and light use at others. During heavy periods of congestion, data may be discarded. However, Frame Relay uses several software-configurable parameters and techniques to control congestion and to avoid data loss on the network during these heavy periods. These software parameters are set on a VC-by-VC basis. This section describes these parameters.

Note

The parameters in this section describe how the Frame Relay network handles congestion. The WSX supports these parameters, but they must match those used by your Frame Relay service provider.

Regulation Parameters

The **Committed Information Rate (CIR)**, which is also referred to as “VC Throughput,” is the minimum bandwidth a virtual circuit will provide under normal circumstances. Frames transmitted within the CIR are not tagged by the Frame Relay network as being eligible for discard. Frames transmitted above the CIR are tagged for discard, but they will normally only be discarded if the virtual circuit or network becomes congested. For example, if the CIR is 16 Kbps and you have a 56 Kbps line, then this virtual circuit will always get at least 16 of the available 56 Kbps. The extra 40 Kbps ($56-16=40$) is normally available to this virtual circuit as long as it is not being used by other virtual circuits and depending on how you have configured the **Committed Burst Size (Bc)** and **Excess Burst Size (Be)**, which are described below.

The CIR is normally a rate given by your service provider. Your service provider may not allow a CIR, in which case your CIR would be 0 (no committed data rate for the virtual circuit).

The **Committed Burst Size (Bc)** is the amount of data that the network will guarantee to transfer under normal conditions. The data may or may not be contiguous and is expressed in kilobits. This number is related to your CIR. In fact, the CIR is Bc divided by Tc where Tc is the time interval used to express the CIR. If Tc is equal to 1 second (a typical value for Tc) and your Bc is 16 kilobits, then your CIR is equal to 16 Kbps. So in many cases the Committed Burst rate will be the same number as the CIR expressed as a *quantity* of data (kilobits) rather than a data *rate* (kilobits per second).

The **Excess Burst Size (Be)** is the amount of data over-and-above the Committed Burst Size (Bc) that the network will transmit as long as excess bandwidth is available on the virtual circuit. The number is also expressed in kilobits. Data at this level is not guaranteed transfer. Any data exceeding the Committed Burst Size may be part of the Excess Burst Size. If there is no bandwidth available on the virtual circuit or if the network is congested, the first data to be dropped is part of this Excess Burst data.

The Excess Burst Size is related to the Committed Burst Size and the access rate of the Frame Relay line. The Excess Burst Size plus the Committed Burst Size should be less than or equal to the access rate of the Frame Relay line. So, if you have a 56 Kbps line and the Committed Burst size is 16 kilobits, then the Excess Burst Size could range from 0 to 40 kilobits.


By default all of these congestion control parameters are set to zero (0), meaning that congestion control is disabled and data flows at the access rate for learned virtual circuits. Congestion control is not enabled until you set one or more of these parameters to a non-zero number.

Discard Eligibility (DE) Flag

The Frame Relay network keeps track of data that is eligible for discard by using a single bit within each frame. When the data rate exceeds the CIR, frames are tagged (i.e., the DE bit is set to 1). If congestion in the network nears saturation, those frames tagged with the DE bit will be dropped before untagged frames. Unless totally congested, data below the CIR level on all virtual circuits is usually guaranteed delivery. Normally, frames are not dropped on an entire Frame Relay connection, but only those frames that exceed the pre-defined CIR level.

Interaction Among Congestion Parameters

The following example helps illustrate the interaction among congestion regulation parameters. A Frame Relay line has an access rate of 64 Kbps. The guaranteed Committed Information Rate (CIR) is 16 Kbps. The Committed Burst Size is 16 Kilobits and the Excess Burst Size is also 16 Kilobits. These parameters mean that any data exceeding 16 Kilobits (within a Tc sample period) normally will be tagged with a Discard Eligibility flag and could be discarded if congestion occurs on the virtual circuit. In addition, since the Excess Burst Size is 16 kilobits, any frames sent exceeding 32 Kbps will have a higher probability of being discarded. The following graph illustrates this example.



Effect of Congestion Control Parameters on Data


Congestion Control

Frames are shown as broken lines below the Access Rate line. The space between frames indicates the delay between the transmission of each frame. For each second, frames sent within the white zone below the diagonal Access Rate line get through. The shaded area just above the white area contains frames that are stamped for Discard Eligibility that will get through as long as the VC is not congested. The darkest shaded area shows frames that may not get through because they exceed the Excess Burst Size allowed in one second.

Notification By BECN

Each data link header contains a congestion control flag called BECN (Backwards Explicit Congestion Notification), which is usually pronounced “beckon.” Normally this flag is turned off. As with other WAN packet-based networks, frames in Frame Relay may build up in queues at certain points. When a queue is full, due to congestion, frames will be dropped. The senders of this data (Bridge/Router or WSX) may not be aware of the congestion. Frame Relay uses a congestion notification technique to notify the Bridge/Router that traffic is jammed further down the circuit.

When a frame on one side of the bi-directional virtual circuit sees data congested on the other side, the Frame Relay network sets the frame’s BECN flag On. Any subsequent frames that see the congestion also have their BECN flag set On. These BECN frames continue down the virtual circuit until they reach the Bridge/Router or WSX on the other end. The receiving WSX sees the BECN flags and adjusts data flow in the opposite direction. Normally the WSX will slow the speed of data down to the CIR. If the BECNs persist, then data flow is stepped down even further. Data flow will gradually increase back up to the normal rate as soon as BECNs or FECNs (see below) are no longer received.




Congestion Notification Using a BECN

BECN notification only works if traffic flows in both directions. If traffic in the uncongested direction did not exist then there would be no frames for the Frame Relay network to set BECN flags on.

Notification By FECN

Frame Relay headers also contain a congestion control bit called FECN (Forwards Explicit Congestion Notification), which is usually pronounced “Feckon.” Like BECN, the FECN bit also notifies a WSX or Bridge/Router of congestions problems. However, it is set by the Frame Relay network in frames that are actually experiencing congestion. When the WSX receives frames with their FECN bit set, it knows that congestion is already occurring on the virtual circuit in the direction that these FECN frames are travelling. The WSX reacts by reducing the data flow down to the CIR for data in the opposite direction. If the FECNs persist, then data flow is stepped down even further. Data flow will gradually increase back up to the normal rate as soon as FECNs or BECNs are no longer received.



Congestion Notification Using a FECN

Frame Formats Supported

Frames coming in from the Frame Relay network are not translated, but they are manipulated to be compatible for transport over the switch's VBUS. Incoming frames must contain RFC 1490 headers. The following standard 1490 frame types are supported:

- BPDU
- Ethernet 802.3
- Token Ring 802.5 (see Note below)
- FDDI (see Note below)
- IP Routed
- ARP/InARP Routed
- IPX Routed
- Compressed (which decompresses to one of the above supported formats)

Note

Source Routing is not supported on Token Ring and FDDI frames.

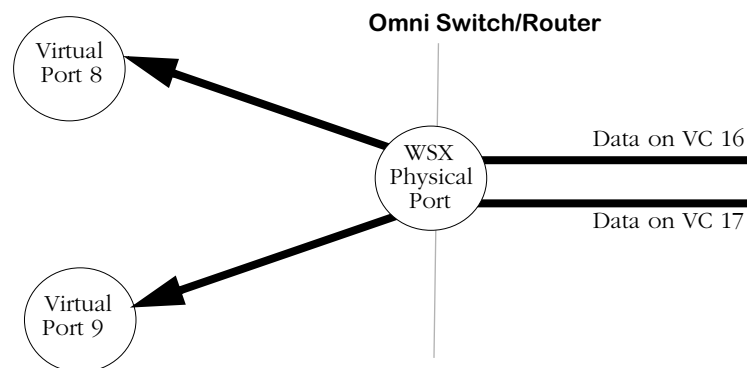
All other frames types from the network are discarded at the physical port level.

Frames coming from the switch to the Frame Relay network are optionally translated if they are a non-Ethernet frame (e.g., FDDI and Token Ring) for a Bridged VLAN. In this case, the frame is translated to an Ethernet frame before it is sent to the Frame Relay interface. Frames from non-Ethernet interfaces can also be sent as is without translation. This translation, which is called Default Bridging Mode, can be configured at the service or port level. In addition, BPDU and Routed frames (IP, ARP, InARP, IPX) are accepted.

Bridging Services

All Frame Relay Virtual Circuits (VCs) belong to a service, whether it be a Bridge, Router, or Trunk service. By default, a virtual circuit belongs to a bridge service. No configuration is necessary for a VC to support bridging on Group 1. However, configuration is necessary for a VC to support Frame Relay Routing, Trunking, or Bridging on a Group other than Group 1.

For bridging there is a one-to-one map between Frame Relay virtual circuits and switch virtual ports. When data is received from a virtual circuit at the physical port level it automatically maps to the corresponding virtual port. For example, if Frame Relay virtual circuit 16 maps to virtual port 8, then all incoming data on this circuit would be incoming data on switch virtual port 8. And if virtual circuit 17 maps to virtual port 9, then all incoming data would be on virtual port 9.



One-to-One Mapping Between Virtual Ports and Virtual Circuits

Frame Relay bridging uses standard Spanning Tree as defined in 802.1d. Typically, one bridge port within the WAN will act as the designated root bridge (and may be the actual root bridge) and maintain a single path through the Frame Relay network. To avoid duplication and loops, some paths will not be allowed.

As far as Spanning Tree is concerned, the virtual ports that map off a Frame Relay physical port are LAN ports. Each port will come up as default bridging on VLAN 1.

A unique aspect of Frame Relay bridging is that MAC addresses must be learned for each DLCI and for each virtual port. So, although the virtual circuits map directly to virtual ports, the bridge must still learn their MAC addresses separately. Also, Frame Relay BPDUs do not have MAC addresses.

One of the disadvantages of bridging in Frame Relay is that broadcasts must be sent across all virtual circuits that are associated with a given physical port for a given group. This requirement can create duplication across the Frame Relay network. At the extreme, on a full T1 line with 96 virtual circuits defined, 96 copies of each broadcast would have to be sent for the same Group. When using access rates at the higher end of the Frame Relay spectrum, you could separate virtual circuits into separate Groups to decrease the size of each broadcast domain. Or, you could use a Routing (IP or IPX) or Trunking configuration to more efficiently manage the data flow.


The configuration of bridging services is described in *Configuring a Bridging Service* on page 29-57.

Frame Relay IP Routing

Frame Relay routing is different than standard LAN IP Routing. In normal LAN IP Routing MAC addresses are used as source and destination addresses. In Frame Relay IP Routing, no MAC addresses are included in a routed frame. In fact, the only address in a routed Frame Relay frame is the DLCI, or virtual circuit identifier. The DLCI is the main identifier for source and destination addresses.

Because Frame Relay uses 10-bit DLCIs as the main addressing units, routed Frame Relay frames require less overhead than LAN IP frames, which use LAN standard 48-bit addresses. However, due to the nature of DLCIs on a WAN, Frame Relay routing requires a special version of the IP protocol. The DLCI for a single VC may or may not be different on both sides of a Frame Relay connection. That's why Frame Relay uses the Inverse Address Resolution Protocol (InARP) to resolve DLCI issues and to automatically learn the IP addresses of remote routers.

The InARP protocol ensures that before any data passes between two Frame Relay routers, those routers notify each other of their IP addresses and associated DLCIs. So, the first communication over a routed Frame Relay network is normally initiated by InARP.




Frame Relay IP Routing

An InARP message is sent between the two routers indicating their IP addresses and associated VC. Once they know each other's IP address and the DLCI of the VC on each end of the link (the same VC may have a different DLCIs on each end), then they can begin normal routing of RIP frames, etc.

The Frame Relay Subnet and "Split Horizon"

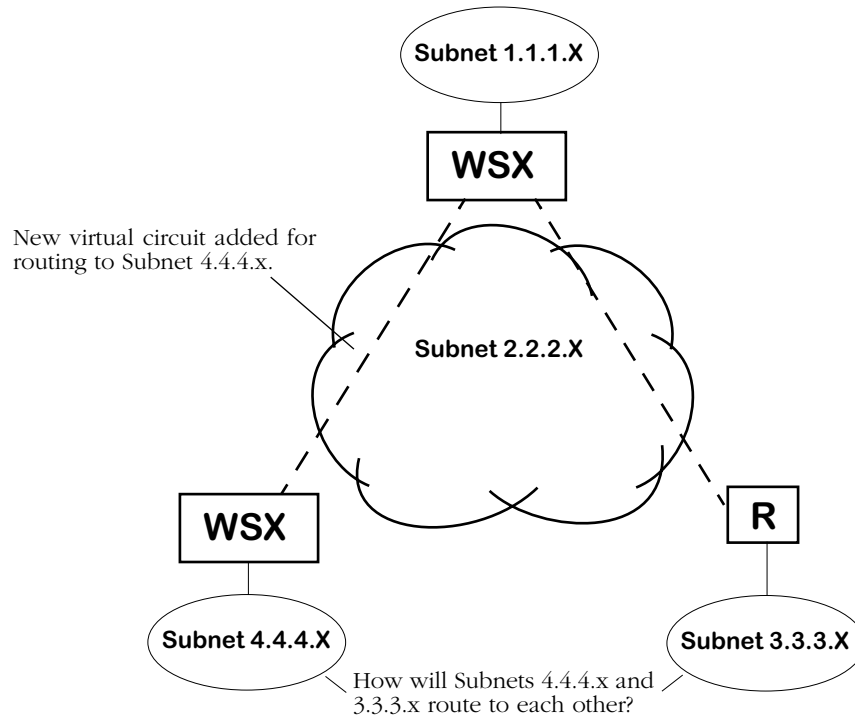
The IP protocol must account for the Frame Relay network in making routing decisions. After all, the WAN network is more than just a single cable, or even several cables, attaching two routers. The solution is to assign the Frame Relay network a unique IP subnet.



Frame Relay Network Is an IP SubNet

In the configuration shown above, one virtual circuit connects the WSX router on IP Subnet 1.1.1.x and the other router on IP Subnet 3.3.3.x. The Frame Relay network, for routing purposes, is considered to be IP Subnet 2.2.2.x. Routing decisions are straightforward in this setup. But if another Router and another IP Subnet were added, a special routing technique must be devised.

If an additional Router and Subnet were added to the network and a new VC was added to connect the new location, then much of the WAN routing load would fall on the WSX attached to Subnet 1.1.1.x.



Adding A New Router Raises New Questions

The new WSX attached to Subnet 4.4.4.x connects to the WAN through the addition of a new virtual circuit connecting directly to the WSX attached to Subnet 1.1.1.x. However, for the new WSX to route to Subnet 3.3.3.x it must go through the WSX router attached to Subnet 1.1.1.x. This is okay for the initial routed path decision. But IP will try to find the most efficient route between Subnet 4.4.4.x and 3.3.3.x. Unfortunately the most efficient route—which would be a direct path between the two routers—is not possible because no WAN link exists between the two.

Frame Relay routing allows the new Subnet, 4.4.4.x, and Subnet 3.3.3.x to route through the WSX router attached to Subnet 1.1.1.x. Normal IP would have a problem with this solution because it does not allow “backtracking” through IP Subnets, which is exactly what must be done in this case. Routed frames actually pass through the Frame Relay Network Subnet 2.2.2.x twice—once to get the WSX Router attached to Subnet 1.1.1.x and another time to get to the Router attached to either Subnet 4.4.4.x or 3.3.3.x.

Standard routing uses a technique called “split horizon” that prevent loops through the same Subnet from occurring. *Frame Relay enhances split horizon to account for the nature of virtual circuits.* Loops through a LAN Subnet are inefficient, but Frame Relay routing makes allowances to compensate for the fact that a WAN does not enjoy the same flexibility with router connections as a LAN.

Note

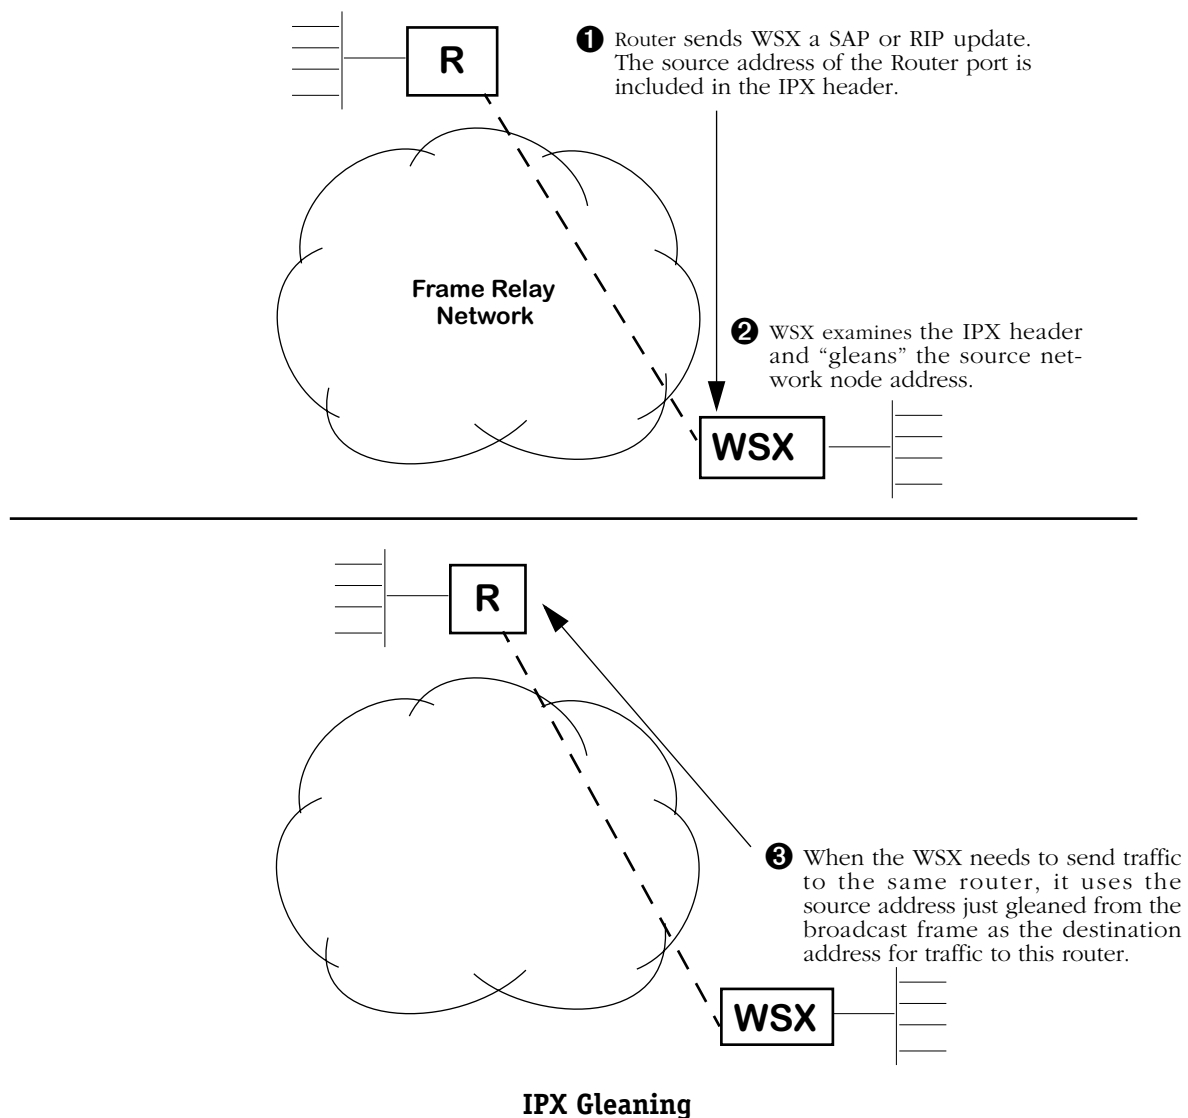
Backtracking in InARP is allowed only through the IP Subnet defined for the Frame Relay network.

The configuration of WSX routing services is described in *Configuring a WAN Routing Service* on page 29-59.

Frame Relay IPX Routing

Frame Relay IPX and IP routing differ in the way they determine the address of a router at each end of a virtual circuit. Instead of using Inverse ARP, IPX uses a process called “gleaning” to determine routing information. In gleaning, the IPX routing protocol on one end of a virtual circuit obtains the network node number for the router at other end of the virtual circuit.

A WSX or router continuously receives RIP and SAP updates on a given virtual circuit. When it receives the first such broadcast, the IPX process looks at, or gleans, the source address from the frame’s IPX header. When the router needs to send traffic on that router later, it uses the source address it just obtained as the destination address for that router. The following diagram illustrates IPX Gleaning.



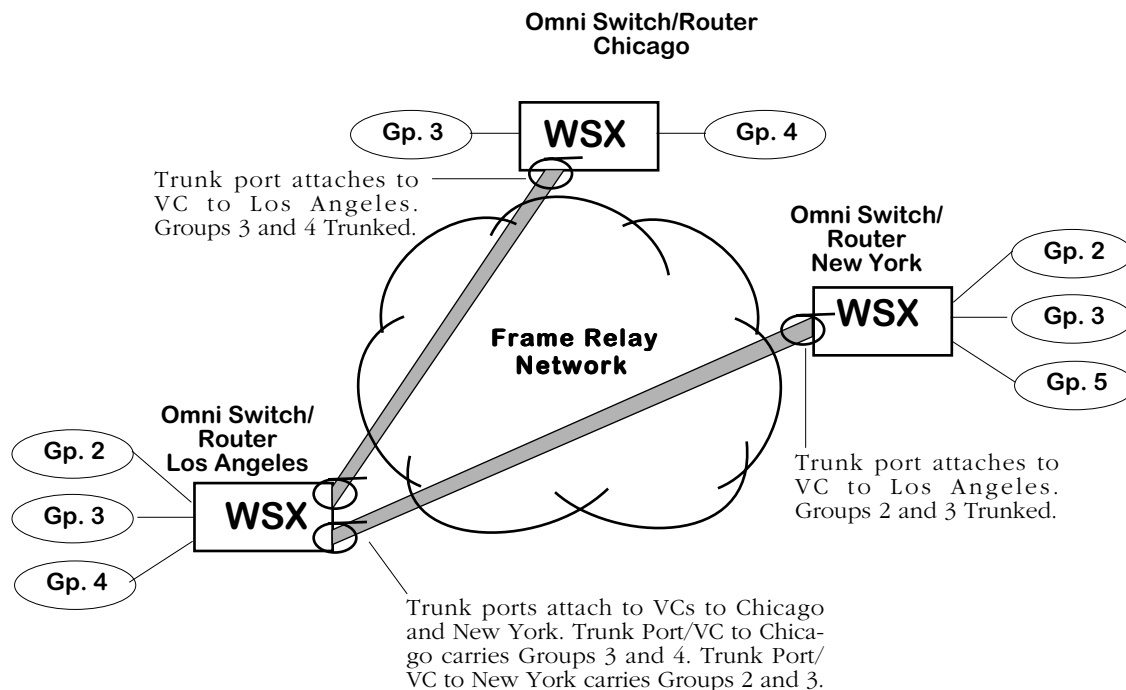
Not all Routers support IPX gleaning. If you need to interoperate with a Router that does not support gleaning, then you may need to statically map addresses on that Router.

The configuration of WSX routing services is described in *Configuring a WAN Routing Service* on page 29-59.

Trunking

A trunking service must be set up for each virtual circuit that will support trunking. When trunking is set up, you specify the slot, port, DLCI, and Groups that are going to be trunked over the virtual circuit.

The illustration below shows a sample trunking configuration. The WSX in Los Angeles has two trunk ports, one to Chicago and one to New York.



Trunk Ports and Virtual Circuits Over Frame Relay Network

Frame Relay virtual ports are mapped one-to-one to virtual circuits, so each of these trunk ports is connected to a virtual circuit. When setting up Trunking you need to be aware of your virtual circuit configurations, their DLCIs, and their termination points. Configuring a Trunking Service is described in *Configuring a Trunking Service* on page 29-62.

Note

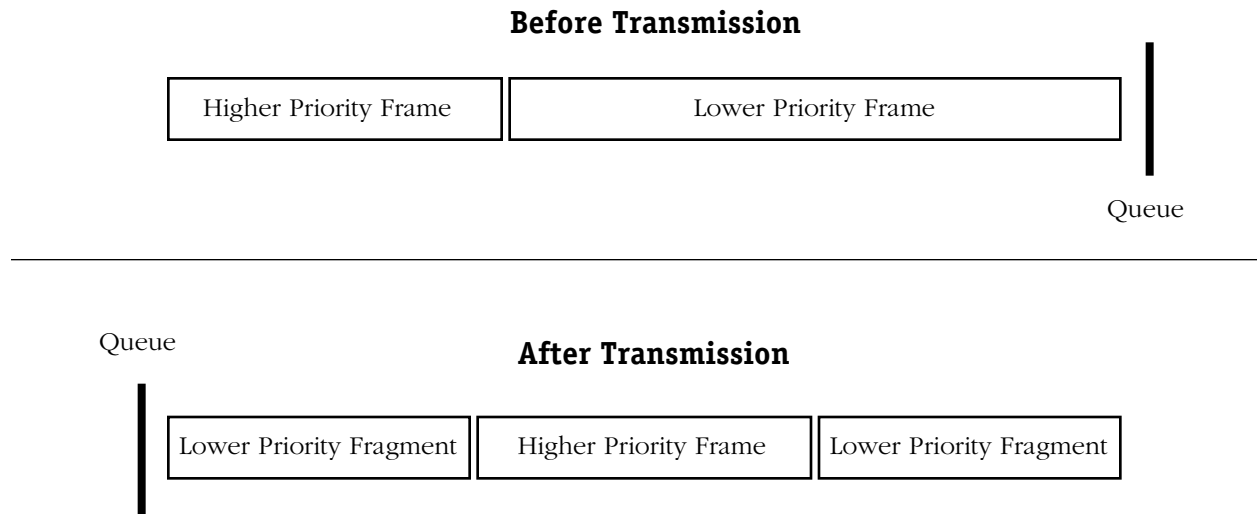
No standard exists for trunking Groups or VLANs over Frame Relay. Therefore, you must configure Trunking using Alcatel's method.

Frame Relay Fragmentation Interleaving

The fragmentation interleaving feature allows for the transmission of higher priority traffic to be expedited by setting a maximum frame size. If lower priority traffic exceeds the set frame size value, it is fragmented into smaller pieces less than or equal to the set maximum frame size.

When the fragmentation feature is enabled, frame traffic is examined for priority. High priority traffic is not fragmented. Fragmented frames are reassembled on the far side (destination) of the link.

The following diagram illustrates this concept:



Frame Relay Fragmentation Process

Fragmentation must be enabled on both ends of the data link for successful fragmentation and reassembly to occur.

The maximum frame size should not be set so low that average frame traffic is broken into more than 64 fragments.

For information on enabling Frame Relay fragmentation, see *Modifying a Virtual Circuit* on page 29-29.

The Frame Relay Software Menu

User interface commands for Frame Relay are on a separate menu that you can access through the **fr** command. The Frame Relay menu is a sub-menu of the **Interface/WAN** menu. Typing **fr** at any system prompt displays the following menu:

Command	Frame-Relay Menu
frstatus	Status of entire chassis, slot, port, and DLCI (e.g., 4/1/32).
frview	View a given slot, port, or DLCI (e.g., 4/1/32).
frmodify	Modify a given slot, port, or DLCI (e.g., 4/1/32).
frdelete	Delete a given port or DLCI (e.g., 4/1/32).
fradd	Add a DLCI with slot, port, DLCI (e.g., 4/1/32)
Main	File Summary VLAN Networking
Interface	Security System Services Help

You can start any of the commands by typing just the first three (3) letters of the command name. For example, to use the **frview** command you could type only **frv**.

The following sections describe the use of commands on the Frame Relay menu.

Setting Configuration Parameters

When you plug in a WSX board it is automatically configured with default settings. The WSX board will default the WAN port protocol to frame relay for WSX serial ports, T1 and E1 ports. Commands generic to the WSX module can be found in Chapter 49.

By default the WSX frame relay software uses ANSI T1.617 Annex D for the Data Link Control Management Interface (DLCMI) and uses a Committed Information Rate (CIR) of 0. In addition, the access rate defaults to 64 Kbps for RS-232 cables and to 2 Mbps for all other cable types. You can change these settings as well as several other settings with the **frmodify** command.

You have a choice of modifying parameters at the port or DLCI (virtual circuit) level. You receive different configuration choices depending upon which level you choose. The two sections below describe both ways to use the **frmodify** command.

Modifying a Port

To modify a port, enter the following command

```
frmodify <slot>/<port>
```

where **<slot>** is the slot number where the WSX board is located, and **<port>** is the port number on the WSX board that you want to modify. For example, if you wanted to modify port number 1 on the WSX board in switch slot 3, you would enter

```
frmodify 3/1
```

or

```
frm 3/1
```

A screen similar to the following displays:

Modifying Frame Relay port for Slot 2, Port 1.

- 1) Description..... =
 {Enter Up to 30 Characters}
- 2) Administrative Status = Up
 {(U)p, (D)own}
- 3) DLCMI Type = ANSI T1.617 Annex D
 {(L)MI Rev. 1.0, T1.617 Annex (D), Q.933 Annex (A), (N)one }
- 31) LMI Procedure Type = Bidirectional
 { (B)idirectional, (U)ser, (N)etwork }
- 4) Polling Interval T391/nT1 = 10
 {1 through 255 seconds}
- 41) Poll Verification Interval T392 (seconds). = 15
 {1 through 255 seconds}
- 5) Full Status Interval N391/nN1 = 6
 {1 through 10}
- 6) Error Threshold N392/nN2 = 3
 {1 through 10}
- 61) Network Error Threshold N392 = 3
 {1 through 10}
- 7) Monitored Events Counter N393/nN3 = 4
 {1 through 10}
- 71) Network Monitored Events Counter N393 = 4
 {1 through 10}
- 8) Default Bridging Group..... = 1
 {1-65535}
- 9) Default Frame Relay Bridging Mode..... = Bridge All
 {1=Bridge All, 2=Ethernet only,
 (AN) Bridge All No FCS, (EN) Ethernet Only No FCS}
- 10) Default Routing Group..... = 0
 {1-65535}
- 11) Default Compression Admin Status = Enabled
 {(E)nable, (D)isable}
- 12) Default Compression PRetry Time = 3
 {1-10}
- 13) Default Compression PRetry Count = 10
 {3-255}

To change a value, enter the corresponding number, an '=', and the new value. For example to set a new description, use

: 2=My new Description

To clear an entry specify the value as '.' as in

: 2=.

When complete enter "save" to save all changes, or cancel or Ctrl-C to cancel all changes. Enter ? to view the new configuration.

(save/quit/cancel)

:

You make changes by entering the line number for the option you want to change, an equal sign (=), and then the value for the new parameter. When you are done entering all new values, type **save** at the colon prompt (:) and all new parameters will be saved. The following sections describe the options you can alter through this menu.

◆ Caution ◆

Several of the parameters in this menu (**Polling Interval, Full Status Interval, Error Threshold, and Monitored Events Counter**) are set to Frame Relay defaults and do not need to be changed except in rare cases. These options should only be modified by experienced Frame Relay network administrators. Changes to these options will probably also require coordination with the service provider.

In addition, the **DLCMI Type** option must be entered correctly or the WSX will not be able to communicate with the Frame Relay network. The WSX board is self-configuring in many ways, but it cannot compensate for an incorrect DLCMI Type.

1) Description

Enter a description for this port. The description can be up to 30 characters long.

2) Administrative Status

This option enables or disables the port. If set to **UP**, then the port has been enabled and can transmit data as long as its Operational Status is also UP. If set to **DN**, then the port will not pass data even if its physical connection is good.

3) DLCMI Type

This field specifies the Data Link Control Management Interface (DLCMI) that you want to use for Frame Relay and virtual circuit management. You have three choices for this protocol, each of which corresponds to an existing widely-used protocol. The letters used in the **frmodify** screen correspond to the following DLCMIs:

- L** LMI rev. 1.0 (LMI)
- D** ANSI T1.617 Annex D
- A** CCITT-ITU-T Q.933 Annex A
- N** None

Enter your choice by specifying the letter corresponding to your choice.

◆ Important Note ◆

The DLCMI protocol that you enter must match that used by your service provider. Entering an incorrect DLCMI protocol may cause the port to not operate. The WSX needs to know the protocol you are using to establish communication with the Frame Relay network.

31) LMI Procedure Type

This field specifies the Local Management Interface (LMI) procedure type for this Frame Relay port. You have three choices for the LMI procedure type. The letters used in the **frmodify** screen correspond to the following:

- B** Bidirectional
- U** User (the default)
- N** Network

Enter your choice by specifying the letter corresponding to your choice.

4) Polling Interval T391/nT1

This interval is the time in seconds between WSX port polls of the Frame Relay network. The WSX port polls the network by sending STATUS ENQUIRY messages, which check the link integrity of the Frame Relay connection. By default this interval is set to 10 seconds, but you can increase or decrease it. The default is the standard Frame Relay value. Increasing the polling interval lightens the data load on the port, as it does not have to poll as often. The interval may range from 1 second to 4 minutes and 15 seconds (255 seconds).

◆ Important Note ◆

This option should only be modified by experienced Frame Relay network administrators.

5) Full Status Interval N391/nN1

This interval is the time in seconds between FULL STATUS ENQUIRIES initiated by the WSX to the Frame Relay network. The Frame Relay network returns a list of all virtual circuits and whether they are active or inactive. You can set this interval from 1 to 10 seconds. By default, this interval is set to 6 seconds, which is the standard Frame Relay default value.

◆ Important Note ◆

This option should only be modified by experienced Frame Relay network administrators.

6) Error Threshold N392/nN2

The number of DLCMI protocol errors that will be tolerated before determining the Frame Relay line is down and all associated virtual circuits are inactive. These errors may include timeouts from STATUS ENQUIRY polls and invalid STATUS messages returned from the Frame Relay network. By default, this threshold is set to 3, which is the standard Frame Relay default value.

◆ Important Note ◆

This option should only be modified by experienced Frame Relay network administrators.

7) Monitored Events Counter N393/nN3

The number of status polling intervals over which the **Error Threshold** is counted. This value should be greater than or equal to the **Error Threshold**. If the station received the number of errors specified in **Error Threshold** within the number of polling intervals specified for the **Monitored Events Counter**, then the Frame Relay line is considered down and all associated virtual circuits are considered inactive. By default, this counter is set to 4, which is the standard Frame Relay default value.

◆ Important Note ◆

This option should only be modified by experienced Frame Relay network administrators.

8) Default Bridging Group

The default Group for bridging any virtual circuits (user-configured or learned from the Frame Relay network) that are not specifically assigned to a Bridging service. If you set this value to 0, then virtual circuits will not perform bridging unless assigned to a bridging service. By default, the Default Bridging Group is set to 1. By entering a value here you can change the default for this port.

◆ Important Note ◆

The **Default Bridging Group** only applies to user-side (i.e., the LMI Procedure Type has been set to **User**) Frame Relay ports.

9) Default Frame-Relay Bridging Mode

This field sets the default translation option for this port. When set to **All**, no translation is performed on frames before they are sent out to the Frame Relay network; frames are sent as is. When set to **Eth-only**, non-Ethernet frames are first translated to the default Ethernet frame format for this port before they are sent out to the Frame Relay network. Any MAC translations configured through the Switch menu are valid.

◆ Important Note ◆

The **Default Frame-Relay Bridging Mode** only applies to user-side (i.e., the LMI Procedure Type has been set to **User**) Frame Relay ports.

10) Default Routing Group

The default Group for bridging any virtual circuits (user-configured or learned from the Frame Relay network) that are not specifically assigned to a Routing service. If you set this value to 0 (the default value), then virtual circuits will not perform Routing unless specifically assigned to a Routing service.

This option is intended to simplify Routing configuration if you do not need to route many Groups over a Frame Relay physical port. The WSX learns about Data virtual circuits from the Frame Relay network. To enable routing on each of these learned virtual circuits, you would have to set up each circuit individually. However, if you already know the Routing Group for your VCs, then you can specify it here and all VCs will be placed in that Group with an extra configuration on your part. Note that you still need to set up a Frame Relay Routing Group through the **crgp** command. See *Configuring a WAN Routing Service* on page 29-59 for more information.

◆ Important Note ◆

The **Default Routing Group** only applies to user-side (i.e., the LMI Procedure Type has been set to **User**) Frame Relay ports.

11) Default Compression Admin Status

This option indicates whether compression negotiation is enabled or disabled for virtual circuits that are learned from the Frame Relay network. Configured virtual circuits are enabled for compression through the **fradd** or **frmodify** (virtual circuit level) commands. The compression negotiation status that you set up for a specific virtual circuit overrides the status you enter here for the physical port.

12) Default Compression PRetry Time

This option sets the number of seconds between compression negotiation messages. If compression negotiation is enabled, the WSX will send compression negotiation messages as many times as you indicate in the Default Compression PRetry Count. The time between these tries is indicated in this field. The number of seconds between retries may range between 1 and 10 seconds. The default is 3 seconds. This default can be by using the **frmodify** command on an individual virtual circuit.

◆ Important Note ◆

The **Default Compression PRetry Time** should only be modified by experienced Frame Relay network administrators. In addition, it should match the setting for the remote Omni Switch/Router or Bridge/Router.

13) Default Compression PRetry Count

This option sets the total number of compression negotiation messages that will be sent before giving up and not running compression. You enter the time between these retries in the Default Compression PRetry Time field. The number of retries can range from 3 to 255. The default is 10. This default can be by using the **frmodify** command on an individual virtual circuit.

◆ Important Note ◆

The **Default Compression PRetry Time** should only be modified by experienced Frame Relay network administrators. In addition, it should match the setting for the remote switch or Bridge/Router.

Modifying a Virtual Circuit

To modify a virtual circuit, enter the following command:

```
frmodify <slot>/<port>/<DLCI>
```

where **<slot>** is the slot number where the WSX board is located, **<port>** is the port number on the WSX board, and **<DLCI>** is the number used to identify the virtual circuit that you want to modify. For example, if you wanted to modify DLCI 17 on Port number 1 of the WSX board in slot 3, you would enter

```
frmodify 3/1/17
```

or

```
frm 3/1/17
```

A screen similar to the following displays:

Modifying Frame Relay DLCI for Slot 3, Port 1, DLCI 17.

- 1) Administrative State = U
 {(U)p, (D)own}
- 2) Committed Information Rate (CIR) in BPS = 0
 {0 through line speed in BPS}
- 3) Committed Burst Rate(Bc) = 0
 {0 through positive number in bits}
- 4) Excess Burst Rate(Be) = 0
 {0 through positive number in bits}
- 5) Compression Administrative Status = Enabled
 {(E)nabled, (D)isabled}
- 6) Compression PRetry Time = 3
 {1..10}
- 7) Compression PRetry Count = 10
 {3..255}
- 8) Fragmentation Interleaving = Disable
 {(E)nabled, (D)isabled}

To change a value, enter the corresponding number, an '=', and the new value. For example to set a new DLCI Active/Inactive Traps, use
: 5=d

When complete enter "save" to save all changes, or cancel or Ctrl-C to cancel all changes. Enter ? to view the new configuration.

You make changes by entering the line number for the option you want to change, an equal sign (=), and then the value for the new parameter. When you are done entering all new values, type **save** at the colon prompt (:) and all new parameters will be saved. The following sections describe the options you can alter through this menu.

Administrative State

This option enables and disables the virtual circuit you are modifying. Setting this option to **Up** enables the circuit and allows data to be sent or received on it as long as the Operational Status is also Up. Setting this option to **Down** disables the circuit; no data can be sent on the circuit. This may be a good option to use when preconfiguring a virtual circuit in advance of live network operation.

Committed Information Rate (CIR)

This field sets the Committed Information Rate (CIR) for this virtual circuit. See *Congestion Control* on page 29-8 for further information on the CIR.

◆ Important Note ◆

The **CIR** that you enter must match that used by your service provider. This option should only be modified by experienced Frame Relay network administrators.

Committed Burst Size (Bc)

The Committed Burst Size (BC) is the amount of data that the network will guarantee to transfer under normal conditions. See *Congestion Control* on page 29-8 for further information.

◆ Important Note ◆

The **Committed Burst Rate** that you enter must match that used by your service provider. This option should only be modified by experienced Frame Relay network administrators.

Excess Burst Size (Be)

The Excess Burst Size (Be) is the amount of data over-and-above the Committed Burst Size (BC) that the network will transmit as long as excess bandwidth is available. See *Congestion Control* on page 29-8 for further information.

◆ Important Note ◆

The **Excess Burst Rate** that you enter must match that used by your service provider. This option should only be modified by experienced Frame Relay network administrators.

Compression Administrative State

This field enables and disables compression negotiation for this virtual circuit. If set to enable, then the WSX will query the Bridge/Router on the other end of the Frame Relay link as to whether it supports compression. Compressed data will be sent only when the other Bridge/Router also supports compression. If the Bridge/Router on the other end is an Omni Switch/Router, then data would be sent compressed as long as you set the Compression Administrative State to Enabled.

Disabling Compression Administrative State means that data will not be sent compressed even if the other Bridge/Router supports compression. Data compression is always negotiated before it is activated.

Compression PRetry Time

This option sets the number of seconds between compression negotiation messages on this virtual circuit. If compression negotiation is enabled, the WSX will send compression negotiation messages as many times as you indicate in the Compression PRetry Count. The time between these tries is indicated in this field. The number of seconds between retries may range between 1 and 10 seconds. The default is 3 seconds. The value you enter for this field overrides the **Default Compression PRetry Time** set up for the physical port with which this virtual circuit is associated.

◆ Important Note ◆

The **Compression PRetry Time** that should only be modified by experienced Frame Relay network administrators. In addition, it should match the setting for the remote Omni Switch/Router or Bridge/Router.

Compression PRetry Count

This option sets the total number of compression negotiation messages that will be sent before giving up and not running compression on this virtual circuit. You enter the time between these retries in the Compression PRetry Time field. The number of retries can range from 3 to 255. The default is 10. The value you enter for this field overrides the **Default Compression PRetry Count** set up for the physical port with which this virtual circuit is associated.

◆ Important Note ◆

The **Compression PRetry Count** that should only be modified by experienced Frame Relay network administrators. In addition, it should match the setting for the remote switch or Bridge/Router.

Fragmentation Interleaving

The fragmentation feature for Frame Relay, when activated, allows for lower priority frames exceeding a preset size to be fragmented into multiple frames so that higher priority traffic can be sent more quickly.

Fragmented frames are reassembled at destination.

Selected whether to **enable** or **disable** Frame Relay fragmentation. When fragmentation is **enabled**, the fragmentation menu is expanded to include a parameter for setting the maximum frame size, as shown.

81) Fragment Size = 53
 Valid size between 24 - 1600

Frame size is in kilobytes.

Adding a Virtual Circuit

Data virtual circuits and their DLCIs are normally learned through status messages with the Frame Relay network. However, it may be convenient to pre-configure these virtual circuits before connecting to a live network. In such a case you will need to use the **fradd** command to set parameters for the virtual circuit. The information for the virtual circuit will be stored in the WSX database. This method of configuration is different than using the **frmodify** command, which changes virtual circuit parameters after the circuit has been learned from the network or configured through **fradd**.

To set up a data virtual circuit, enter the following command

```
fradd <slot>/<port>/<DLCI>
```

where **<slot>** is the slot number where the WSX board is located, **<port>** is the port number on the WSX board, and **<DLCI>** is the number used to identify the virtual circuit that you want to add. For example, if you wanted to add DLCI 32 on Port number 1 of the WSX board in slot 2, you would enter

```
fradd 2/1/32
```

or

```
fra 2/1/32
```

A screen similar to the following displays:

Adding Frame Relay port for Slot: 2, Port: 1 Dlci: 32.

- 1) Administrative State = UP
 {(U)p, (D)own}
- 2) Committed Information Rate (CIR) in BPS = 0
 {0 through line speed in BPS}
- 3) Committed Burst Rate (Bc) in bits = 0
 {0 through positive number in bits}
- 4) Excess Burst Rate (Be) in bits = 0
 {0 through positive number in bits}
- 5) Compression Administrative Status = Enabled
 {(E)nabled, (D)isabled}
- 6) Compression PRetry Time = 3
 {1..10}
- 7) Compression PRetry Count = 10
 {3..255}
- 8) Fragmentation Interleaving = Disable
 {(E)nabled, (D)isabled}

Enter the value for each parameter after the colon prompt (:). An additional field, **DLCI Number**, is displayed if you do not specify a DLCI number in the **fradd** command. The remaining parameters are the same ones used for the **frmodify** command. See *Modifying a Virtual Circuit* on page 29-29 for information on each of these parameters.

When you have entered values in all fields, the following prompt displays

```
Do you want to configure additional DLCIs? {(Y)es, (N)o}
```

Enter a **Y** to set up additional virtual circuits or enter **N** to exit the **fradd** command. If you enter **Y**, then you are prompted for all virtual circuit parameters again.

Viewing Configuration Parameters for the WSX

You can view all current parameters for a WSX port or an individual virtual circuit using the **frview** command. These parameters will be either the default parameters or parameters you modified using the **frmodify** command or network management software.

You have a choice of viewing parameters at the chassis, port or DLCI (virtual circuit) level. You receive different configuration choices depending upon which level you choose. The sections below describe both ways to use the **frview** command.

Viewing Parameters for all WSXs in the Chassis

To view port parameters for all WSX boards in a chassis, enter the following command

```
frview
```

```
or
```

```
frv
```

A screen similar to following displays:

Frame Relay Configuration for Chassis:

Slot/Port	Intf Type	Speed BPS	Clocking	Default Bridging Grp	Default Routing Grp
3/1	V35DTE	0	External	1	0
3/2	V35DCE	0	External	1	0
3/3	*NONE*	0	External	1	0
3/4	*NONE*	0	External	1	0

Only ports configured as frame relay (see the **wpm** command in Chapter 49) will be displayed in this screen. This screen lists all the current values for the listed parameters. These parameters are the same ones set through the **frmodify** command. For detailed information on these values, see *Modifying a Port* on page 29-22. For detailed information on the **Intf Type** column, see *Intf Type* on page 29-39.

Viewing Port Parameters

To view port parameters, enter the following command

```
frview <slot>/<port>
```

where **<slot>** is the slot number where the WSX board is located, and **<port>** is the port number on the WSX board on which you want to view information. For example, if you wanted to view configuration parameters for Port number 1 on the WSX board in slot 2, you would enter

```
frview 2/1
or
frv 2/1
```

A screen similar to following displays:

Frame Relay port for Slot 2, Port 1.

```
1) Description..... = Port1
2) Administrative Status ..... = UP
3) DLCMI Type..... = ANSI T1.617 Annex D
   31) DLCMI Type ..... = User
4) Poll Verification Interval T392 (seconds ) ..... = 15
5) Full Status Interval N391/nN1 ..... = 6
6) Error Threshold N392/nN2 ..... = 3
7) Monitored Events Counter N393 ..... = 4

8) Default Bridging Group ..... = 1
9) Default Frame-Relay Bridging Mode ..... = Bridge All
10) Default Routing Group ..... = 0
11) Default Compression Admin Status..... = Enabled
12) Default Compression PRetry Time ..... = 3
13) Default Compression PRetry Count ..... = 10
```

This screen lists all the current values for the listed parameters. These parameters are the same ones set through the **frmodify** command. For detailed information on these values, see *Modifying a Port* on page 29-22.

Viewing Virtual Circuit Parameters

To view virtual circuit parameters, enter the following command

```
frview <slot>/<port>/<DLCI>
```

where **<slot>** is the slot number where the WSX board is located, **<port>** is the port number on the WSX board, and **<DLCI>** is the number used to identify the virtual circuit that you want to view. For example, if you wanted to view configuration parameters for DLCI 17 on Port number 1 of the WSX board in switch slot 3, you would enter

```
frview 3/1/17
```

or

```
frv 3/1/17
```

A screen similar to the following displays:

Frame Relay DLCI for Slot 3, Port 1, DLCI 17.

```
1) Administrative State ..... = UP
2) Committed Information Rate (CIR) in BPS ..... = 16000
3) Committed Burst Rate(Bc) in bits ..... = 16000
4) Excess Burst Rate(Be) in bits..... = 40000
5) Compression Administrative Status..... = Enabled
6) Compression PRetry Time ..... = 3
7) Compression PRetry Count ..... = 10
8) Fragmentation Interleaving ..... = Enabled
81) Fragmentation Size ..... = 53
```

This screen lists all the current values for the listed parameters. These parameters are the same ones set through the **frmodify** command. For detailed information on these values, see *Modifying a Virtual Circuit* on page 29-29.

Deleting Ports and Virtual Circuits

You can delete a WSX port or virtual circuit. When you delete a port or virtual circuit all configuration parameters revert back to default settings. You can use the **frdelete** command to delete:

- a single virtual circuit, or
- a port and all of its associated virtual circuits

The **frdelete** command always requires you to indicate at least a slot and port number. You cannot, for example, enter **frdelete** along with no slot and port parameters.

Deleting a Virtual Circuit

You can delete a single virtual circuit as long as you know its DLCI number and the WSX port where it exists. Deleting a virtual circuit resets the configuration parameters on that circuit to configuration and bridging defaults. By default, a virtual circuit is assigned to Group 1.

Virtual circuits are also not actually “deleted” when you use **frdelete**. The Frame Relay network stills sees them as active or inactive. If the virtual circuit was configured (management circuit or a circuit configured through **frmodify**), then the database record for the circuit is deleted; the VC is still present as long as it was present before you deleted it. If the virtual circuit is learned (through status updates from the Frame Relay network), then the database record for the circuit is deleted, but the circuit is still present.

To delete a virtual circuit, enter the following command

```
frdelete <slot>/<port>/<DLCI>
```

where **<slot>** is the Omni Switch/Router slot number for the WSX board, **<port>** is the port to which the virtual circuit maps, and **<DLCI>** is the identification number for the virtual circuit. For example, if you wanted to delete virtual circuit 32 on Port 1 of the WSX board in slot 2, then would enter:

```
frdelete 2/1/32
```

or

```
frd 2/1/32
```

This system returns the following prompt to confirm the deletion:

```
This will delete Slot 2, Port 1, DLCI 32. Continue? {(Y)es, (N)o} (N)
```

Enter a **Y** to confirm the deletion or press **<Enter>** to cancel the deletion.

Deleting a Port and Its Virtual Circuits

You can delete a port as well as all of its associated virtual circuits. Deleting a port means that all configuration parameters on the port and all learned virtual circuits will revert back to default settings. The port is not logically deleted, and can still be reconfigured after the delete. To truly “delete” a port you must disconnect its cable or set its Administrative Status to Disable.

To delete a virtual circuit, enter the following command:

```
frdelete <slot>/<port>
```

where **<slot>** is the Omni Switch/Router slot number for the WSX board, **<port>** is the port number on the WSX board that you want to delete. For example, if you wanted to delete Port 1 of the WSX board in slot 2, then would enter:

```
frdelete 2/1
```

or

```
frd 2/1
```

This system returns the following prompt to confirm the deletion:

```
This will delete Slot 2, Port 1 and its DLCIs. Continue? {(Y)es, (N)o} (N)
```

Enter a **Y** to confirm the deletion or press **<Enter>** to cancel the deletion.

Obtaining Status and Statistical Information

You can obtain general and detailed Frame Relay statistical information on all WSX boards in the switch, a single WSX board, individual ports, and individual virtual circuits. The **frstatus** command is used to provide this information. This information includes types of physical interface, access rate of the Frame Relay line, and errors. In addition, the **frstatus** command can display the number of frames received and transmitted categorized by frame type (i.e., compressed/uncompressed, Ethernet, IP, IPX, BPDU).

Information on All Boards in a Switch

To obtain status information on all WSX boards in a switch, you enter the **frstatus** command without any parameters as follows:

```
frstatus
```

This command displays a screen similar to the following:

```

Frame Relay Status for the Chassis:
      Admin/      Intf      Speed      VC's
      Oper      Status      Type      BPS      Clocking      Active/
      Status      Type      BPS      Clocking      Inactive
=====
 4/1  UP/UP  V35DCE  2048000  Split        2/0
 4/2  DN/DN  *NONE*  EXT CLK  External     0/0
 4/3  UP/DN  *NONE*  EXT CLK  External     0/0
 4/4  UP/UP  232DCE   56000   Internal    19/1
    
```

Only ports configured as frame relay (see the **wpm** command in Chapter 49) will be displayed in this screen. Each row in the table corresponds to a physical port on a WSX board in the switch. The following sections describe the columns shown in this table:

Slot/Port

The first number in this column is the slot in the switch where this WSX is installed. The second number is the port number on the WSX.

Admin/Oper Status

This column shows the Administrative and Operational Status of this WSX port. The status indicator before the slash refers to the Administrative Status. If **UP**, then the port has been enabled and can transmit data as long as its Operational Status is also **UP**. If the Administrative Status is **DN**, then the port will not pass data even if its physical connection is good.

The status indicator after the slash refers to the Operational Status. If **UP**, then the port is capable of passing data as long as it has been logically enabled at the Administrative level. If **DN**, then the port cannot pass data because of a problem in the physical connection (e.g., cable disconnected, WSX could not detect cable type) or because the port is Administratively Down.

Intf Type

This column indicates the physical cable type connected to this port. This cable type is automatically sensed by the WSX hardware. This column indicates the cable type and whether it is DCE or DTE. The following values may display in this column

- **V35DTE** (V.35 DTE cable)
- **V35DCE** (V.35 DCE cable)
- **232DTE** (RS-232 DTE cable)
- **232DCE** (RS-232 DCE cable)
- **X21DTE** (X.21 DTE cable)
- **X21DCE** (X.21 DCE cable)
- **530DTE** (RS-530 or RS-449 EIA DTE cable)
- **530DCE** (RS-530 or RS-449 EIA DCE cable)
- **T1** (T1 port)
- **E1** (E1 port)

The WSX sees RS-530 and RS-449 cables the same because they are electrically identical. However, this does not affect the operation of either cable type. Both RS-530 and RS-449 cables are supported. If no cable is connected to a port, then this column will display

NONE

If an error has been detected on the port (e.g., cable type could not be detected), then the following value displays:

ERROR!

Speed BPS

This column indicates the speed, or access rate, between the WSX serial port and DSU or other “physical” DTE device. The speed is expressed in bits per second (bps). This speed is the total available bandwidth on the line connected to this port. Virtual circuits on this port share this bandwidth.

Usually, the WSX port will be a physical DTE device and the speed will be determined by the DSU. In this case, this value will read **EXT CLK**, which means the WSX port gets its clocking from an externally attached DCE device (i.e., DTE cable plugged into WSX port) or no cable is attached. If the WSX port is a physical DCE device (i.e., DCE cable plugged into WSX port), then this value will be the actual clock rate used by the port. The speed on a T1 port will always be 1544000; the speed for an E1 port will always be 2048000.

Clocking

This field indicates the type of clocking used to clock transmit and receive data in and out of the serial port. When the clock is out-of-phase, you receive errors. If this value is set to External, then clocking is controlled by the external DCE (a DSU or other DCE device on the other end of the cable from the WSX port). External clocking is the default option when the WSX is a physical DTE device (i.e., controlled by an external DCE device).

Note

See Chapter 49, “Managing WAN Modules,” for documentation on setting the clocking mode for serial ports, see Chapter 54, “Managing T1 and E1 Ports,” for documentation on setting the clocking mode for T1 and E1 ports, and see Chapter 55, “Managing DS3/E3 Modules,” for documentation on setting the clocking mode for DS3 and E3 ports.

If this value is set to Internal, then clocking is controlled by the internal DCE (the WSX). Internal clocking should only be selected if the WSX is a physical DCE device and you are using an RS-232 cable. Internal clocking is the default when the WSX is a physical DCE device and an RS-232 DCE cable is connected to this port. For T1 and E1 ports, internal clocking is equivalent to local timing.

Note

The Clocking value only makes a difference if the WSX port is a physical DCE port (i.e., DCE cable plugged into the WSX port). If the WSX port is a physical DTE port, then Clocking will default to External.

Split clocking, which is also known as “loop timing,” uses additional control signals (TXCE and RXCE) to keep the WSX and DSU clocking in sync. Split clocking takes the incoming clock signals (TX clock and RX clock) and loops them back out to the DSU. The WSX and DSU uses these additional signals to communicate the current status of their clocks. Split clocking should only be used if the WSX is a physical DCE device and you are using a non-RS-232 cable, such as V.35.

◆ Important Note ◆

Split clocking is required if the access rate of the Frame Relay line is greater than 256 Kbps. If Split clocking is not used at these data rates, then data out-of-phase errors, aborts, or CRC errors may occur.

Split clocking is the default when the WSX is a physical DCE device and a non-RS-232 DCE cable is connected to the port. For T1 and e1 ports, external or split clocking is the same as loop timing.

VCs Active/Inactive

Each port will have one or more associated virtual circuits. This column tells you the current status of *Data* virtual circuits. These counts do not apply to management virtual circuits. The first number is the number of active VCs and the second is the number of inactive VCs. An **Active** virtual circuit is one that is operationally Up and capable of transmitting data; it may not necessarily be transmitting at this time. An **Inactive** virtual circuit is present, but for some reason is operationally Down. It is not capable of passing data because either its administrative status was set to Down or the Frame Relay network indicated it was present but Down.

Information on the Ports for One WSX Board

To obtain status information on a single WSX board, you enter the **frstatus** command along with the slot number for the WSX board, as follows:

```
frstatus <slot>
```

where **<slot>** is the slot number where the WSX board is installed. For example, if you wanted to obtain status information for the board in slot 4, you would enter:

```
frstatus 4
```

This command displays a screen similar to the following:

Frame Relay Status for slot: 4

	Admin/ Oper PTStatus	Intf Type	Speed BPS	VCs Active/ Inactive	Frames In	Frames Out	Octets In	Octets Out
	=====	=====	=====	=====	=====	=====	=====	=====
1	UP/UP	V35DTE	2048000	2/0	364	128	8962	2650
2	DN/DN	*NONE*	9600	0/0	0	0	0	0
3	UP/DN	232DTE	56000	0/0	89	90	890	895
4	UP/UP	V35DTE	256000	19/1	9	21	124	245

Each row in the table corresponds to a port on the WSX you requested information on.

PT

The Port number on the WSX board for which statistics are displayed.

Admin/Oper Status, Int Type, Speed Bps, DLCI Active/Inactive

These columns are described in the section, *Information on All Boards in a Switch* on page 29-38. Please refer to this section for detailed information.

Frames In

The total number of frames received on this port since the last time the switch was initialized.

Frames Out

The total number of frames sent on this port since the last time the switch was initialized.

Octets In

The total number of Octets, or bytes, received on this port since the last time the switch was initialized. This statistic includes the data and Frame Relay header fields, but does not include CRC or flag characters.

Octets Out

The total number of Octets, or bytes, sent on this port since the last time the switch was initialized. This statistic includes the data and Frame Relay header fields, but does not include CRC or flag characters.

Information on One Port

To obtain status information on a single WSX port, you enter the **frstatus** command along with the slot number for the WSX board and the port number for which you want to receive information, as follows:

```
frstatus <slot>/<port>
```

where **<slot>** is the slot number where the WSX board is installed and **<port>** is the port number on the WSX board. For example, if you wanted to obtain status information for Port 1 on the WSX module in Slot 4, you would enter:

```
frstatus 4/1
```

This command displays a screen similar to the one shown on the following page:

Obtaining Status and Statistical Information

Frame Relay Status for slot 3, port 1:

Physical Level Information	<pre> Administrative/Operational Status Up/Up Speed Intf. Receive Receive Receive Transmit Signal BPS Type CRC Errors Aborts Overruns Overruns Errors ===== 2048000 V35DTE 18 12 0 0 2 Control DTR RTS DSR CTS DCD Signal ON ON ON ON OFF </pre>
Logical (Frame Relay) Information	<pre> Frame Relay Information: Octets UniCast Discarded Error ===== Frames Frames Count IN 8962 120 2 0 Out 2650 24 5 0 IN+OUT 11612 144 7 0 Administrative/Operational Phase Up/Up Last Error TypeNo Error Since Reset Last Error Time 0 Seconds Interface failures 0 Last interface failure time 0 Seconds </pre>
Virtual Circuit Level Information	<pre> DLCI Information: Admin/ DLCI Oper DLCI Frames Frames Octets Octets Num Status Type In Out In Out ==== ===== 0 UP/UP Configured 10 10 160 140 31 UP/UP Learned 31 20 4196 1250 32 UP/DN Learned 145 110 4813 1450 </pre>
Fragmentation Information	<pre> Frame Relay Fragmentation Information: DLCI Frag Size Frag Status In Frag Out frag Dropped Frag ===== 32 0 Disabled 0 0 0 </pre>

This command displays three (4) layers of information. The top section provides information on the physical interface. The middle section provides information on the logical, or Frame Relay, interface. The third section provides information on the virtual circuits associated with this physical port. The fourth section shows Frame Relay fragmentation information.

Physical Layer Information

The statistics shown in this section are taken at the physical, or serial, interface level.

Administrative/Operational Status

This field shows the Administrative and Operational Status of this WSX port. The status indicator before the slash refers to the Administrative Status. If **UP**, then the port has been enabled and can transmit data as long as its Operational Status is also UP. If the Administrative Status is **DN**, then the port will not pass data even if its physical connection is good.

The status indicator after the slash refers to the Operational Status. If UP, then the port is capable of passing data as long as it has been logically enabled at the Administrative level. If **DN**, then the port cannot pass data because of a problem in the physical connection (e.g., cable disconnected, WSX could not detect cable type) or because the port is Administratively Down.

Speed BPS

The configured speed of the port. For a physical DTE port, the actual rate is determined by the DCE device to which the WSX is attached (i.e., a modem or DSU). For a physical DCE port, the actual rate is the rate configured through the **fmodify** command.

Intf Type

The type of cable that is plugged into the WSX port. The cable may be DCE or DTE and one of 5 different serial types. See *Intf Type* on page 29-39 for further information.

Receive CRC Errors

The total number of frames with an invalid frame check sequence received on the port since the last time the switch was initialized.

Receive Aborts

The total number of frames received that were terminated with an HDLC abort sequence since the last time the switch was initialized. An abort sequence consists of 7 contiguous bits of ones (1111111).

Receive Overruns

The total number of frames that were not received on the port because the system could not keep up with the data flow. Receive overrun errors include buffer errors and errors reported by the RISC processor.

Transmit Overruns

The total number of frames that were not transmitted on the port because the system could not keep up with the data flow. Transmit overrun errors include buffer errors and errors reported by the RISC processor.

Signal Errors

The total number of frames that failed to be received or transmitted due to a loss of modem signals since the last time the switch was initialized. If the WSX port is a physical DTE, then this count is the number of frames dropped due to a loss of the Data Set Ready (DSR) signal. If the WSX port is a physical DCE, then this count is the number of frames dropped due to a loss of the Data Terminal Ready (DTR) signal.

Control Signal

This table (which displays only for serial ports, not T1 or E1 ports) lists two or more control signals along with their current state. If a V.35, RS-232, RS-530, or RS-449 cable is attached then this table lists the following signals:

- **DTR** (Data Terminal Ready.)
- **RTS** (Request To Send.)
- **DSR** (Data Set Ready.)
- **CTS** (Clear To Send.)
- **DCD** (Data Carrier Detect.)

The ON/OFF indicator below the signal name tells you the current status of the signal. Under normal operating conditions (physical connection is good and VC is administratively enabled), all signals should be On.

Whether the signal is an input or an output depends on whether the WSX is a physical DTE or DCE. The following table shows the Input/Output status of each signal type.

Signal	Signal Direction When Port Is...	
	DCE	DTE
DTR	In	Out
RTS	In	Out
DSR	Out	In
CTS	Out	In
DCD	Out	In

If using an X.21 cable, then the table shown in the sample display is replaced by the following table:

Control Signal	C(Control) ON	I(Indicator) ON
----------------	---------------	-----------------

This X.21 table shows 2 rather than 5 signal statuses. The **C** signal is similar to the RTS (Request To Send) signal. The **I** signal is similar to the DCD (Data Carrier Detect) signal. Under normal operating conditions, both the **C** and **I** signals should be On.

Whether the signal is an input or an output depends on whether the WSX is a physical DTE or DCE. The following table shows the Input/Output status of each signal type.

Signal	Signal Direction When Port Is...	
	DCE	DTE
C	In	Out
I	Out	In

Frame Relay Information

The statistics shown in the section are gathered at the Frame Relay protocol level.

Octets

The total octets, or bytes, received (first row) and sent (second row) on this port. The third row shows the cumulative number of octets that have passed through the port (sent and received). This statistic includes the data and Frame Relay header fields, but does not include CRC or flag characters.

UniCast Frames

The total number of Unicast frames received (first row) and sent (second row) on this port. The third row shows the cumulative number of Unicast frames that have passed through this port (sent and received).

Unicast frames are destined for a specific virtual circuit, and are normally sent from one local DLCI to the corresponding DLCI on the other side of the Frame Relay link. In Frame Relay terms, these unicast frames are sent from a logical DTE, such as a WSX port, to a Remote logical DTE, such as a WSX port on the other side of the Frame Relay link.

Discarded Frames

The number of frames discarded due to an error.

Error Count

Frames that contained Frame Relay type errors, such as DLCMI protocol errors and invalid frame format. This count does not include standard physical errors, such as CRC and abort errors.

Administrative/Operational Status

This field shows the Administrative and Operational Status of this WSX port. The status indicator before the slash refers to the Administrative Status. If **UP**, then the port has been enabled and can transmit data as long as its Operational Status is also UP. If the Administrative Status is **DN**, then the port will not pass data even if its physical connection is good.

The status indicator after the slash refers to the Operational Status. If **UP**, then the port is capable of passing data as long as it has been logically enabled at the Administrative level. If **DN**, then the port cannot pass data due to a problem in the physical connection (e.g., cable disconnected, WSX could not detect cable type) or because the port is Administratively Down.

Last Error Type

The last type of Frame Relay DLCMI protocol error received on this port. The following list describes the error types displayed:

Unknown Error	An error occurred but it can not be classified into one of the standard Frame Relay error types.
Receive Short	The receive frame was not long enough to allow demultiplexing. The address field was incomplete, or the protocol identifier was missing or incomplete.
Receive Long	The receive frame exceeded the maximum length for this port.
Illegal DlcI	The DLCI address field in a frame did not match the configured format.
Unknown DlcI	A frame was received on a virtual circuit that was not active or was administratively disabled.
Dlcmi Protocol Error	An Unspecified error occurred while trying to interpret the Link Maintenance frame.
Dlcmi Unknown IE	DLCMI Unknown Information Element. The Link Maintenance frame contained an Information Element type that is not valid for the configured DLCMI protocol.
Dlcmi Sequence Error	The Link Maintenance frame contained a sequence flag that was different than the expected flag.
Dlcmi Unknown RPT	DLCMI Unknown Report Type. The Link Maintenance frame contained a Report Type Information Element with a value that is not valid for the configured DLCMI protocol.
No Error Since Reset	No error has occurred since the last time this port was initialized.

Last Error Time

The time since the last Frame Relay protocol error was received. A value of zero (0) indicates no Frame Relay protocol errors have been received. The type of error that was last received is indicated in the **Last Error Type** field.

Interface Failures

The number of times this Frame Relay port has gone down since it was initialized.

Last Interface Failure Time

The time since the interface was taken down due to excessive errors. Excessive errors are defined as the time when a DLCMI error exceeds the **Error Threshold** or the errors within the **Monitored Events Counter**. A value of zero (0) indicates the interface has not been taken down due to excessive errors. These error parameters are configured through **frmodify** and in most cases should be set to defaults. See *Setting Configuration Parameters* on page 29-22 for more information.

DLCI Layer Information

The information in this section of the display provides statistics on virtual circuits. Each row in this table corresponds to one virtual circuit.

DLCI Num

The DLCI number assigned to this virtual circuit. This value is only valid locally; the same virtual circuit on the other end of the Frame Relay line may or may not use the same DLCI for this VC.

Admin/Oper Status

This field shows the Administrative and Operational Status of this virtual circuit. The status indicator before the slash refers to the Administrative Status. If **UP**, then the virtual circuit has been enabled and can transmit data as long as its Operational Status is also UP. If the Administrative Status is **DN**, then the VC will not pass data even if its physical connection is good.

The status indicator after the slash refers to the Operational Status. If UP, then the virtual circuit is capable of passing data. If **DN**, then the VC cannot pass data because the network has declared the virtual circuit inactive, the network does not respond to STATUS ENQUIRY messages, or the VC is Administratively Down.

DLCI Type

The type of virtual circuit will be either **Configured** or **Learned**. Configured means this VC is a management, or control, circuit that is used by Frame Relay protocols, such as the DLCMI protocols, to pass various status messages. The Frame Relay network does not self-configure management virtual circuits. Data VCs can become “configured” if you use **fmodify** to change any of the default settings for the Data VC. Learned means this is a Data VC that the Frame Relay network informed the WSX module about through status messages (using a Control VC).

Note

The **VC Type** of the management DLCI (0 or 1023) is always **configured** since the Frame Relay network does not dynamically configure management virtual circuits.

Frames In

The number of frames received on this VC since it was created.

Frames Out

The number of frames transmitted on this VC since it was created.

Octets In

The number of octets, or bytes, received on this VC since it was created.

Octets Out

The number of octets, or bytes, transmitted on this VC since it was created.

Fragmentation Information

The information in this section of the display provides statistics on fragmentation. Each row in this table corresponds to one virtual circuit.

DCLI

The virtual circuit that the fragmentation statistics apply to.

Frag Size

The maximum size of a frame if fragmentation is enabled.

Frag Status

Whether fragmentation is **enabled** or **disable**.

In Frag

The number of frame fragments received on this virtual circuit.

Out Frag

The number of frame fragments sent on this virtual circuit.

Dropped Frag

The number of frame fragments dropped from this virtual circuit.

Information on One Virtual Circuit

To obtain status information on a single virtual circuit, you enter the **frstatus** command along with the slot number for the WSX board, the port number, and DLCI number for the virtual circuit on which you want information, as follows:

```
frstatus <slot>/<port>/<DLCI>
```

where **<slot>** is the slot number where the WSX board is installed, **<port>** is the port number on the WSX board, and **<DLCI>** is the virtual circuit identifier. For example, if you wanted to obtain status information for the board in slot 4, port 1, DLCI 32, you would enter:

```
frstatus 4/1/32
```

This command displays a screen similar to the following:

```

Frame Relay Status for slot 4, port 1, DLCI 32
Admin/Oper Status: UP:UP for 0 days, 00:34:40.59
Compression Administrative Status/Operational Phase: Enabled/Operation

```

	Frames In	Frames Out	Frames In+Out	Octets In	Octets Out	%In	%Out
Total	200	250	450	20000	17000		
Ethernet	100	150	250	10000	11000	50	65
802.5	0	0	0	0	0	0	0
FDDI	0	0	0	0	0	0	0
IP	0	4	4	0	2000	0	12
IPX	90	99	185	9560	3960	48	23
BPDU	10	1	11	440	40	2	<1
DE Bit	10	0	10				
FECN Bit	5						
BE CN Bit	7						
Discarded	0						

Frame Relay Fragmentation Information:

DLCI	Frag Size	Frag Status	In Frag	Out frag	Dropped Frag
32	0	Disabled	0	0	0

FRF.9 Compression:	Compressed Frames	Compressed Octets	Uncompressed Octets	Compression Ratio
In	200	10000	20000	2.0:1
Out	250	15000	17000	1.2:1
In+Out	450	25000	37000	1.5:1

The top of the display provides information on the status of this virtual circuit. The **Admin/Oper Status** field indicates the current Administrative and Operation Status for this virtual circuit. The next informational field, **Compression Administrative Status/Operational Phase**, indicates the current Administrative and Operational status for Compression Negotiation on this VC. The Administrative Status will be either **Enabled** or **Disabled**. The Operational Phase will be **Disabled** (compression negotiation not enabled), **Initialization** (compression negotiation in progress), or **Operation** (negotiation successful, data being compressed).

The table below the status information breaks down traffic on the virtual circuit by protocol type. Each row corresponds to a frame type, such as Ethernet or IPX. For each frame type, the number of frames received, frames transmitted, octets received, and octets transmitted is given. The final two columns of the table (**%In** and **%Out**) represent the total percentage of traffic (octets, not frames) for that protocol type.

The Frame Relay Fragmentation Information gives a break down of the fragmented traffic received and sent by this virtual circuit, with indications if fragmentation is currently enabled and the maximum frame size.

The final table provides information on compressed data on this virtual circuit. The following sections describe information in the table.

Total (Protocol)

Statistics in this row indicate traffic for all protocol (Ethernet, IP, IPX, and BPDU) frames and octets on this VC. Statistics for octets, or bytes, include the data and Frame Relay header fields, but they do not include CRC or flag characters.

Ethernet

Statistics in this row indicate traffic for Ethernet (bridged 802.3 or trunked format) frames and octets on this virtual circuit. Statistics for octets, or bytes, include the data and Frame Relay header fields, but they do not include CRC or flag characters.

802.5

Statistics in this row indicate traffic for Token Ring (802.5 format) frames and octets on this virtual circuit. Statistics for octets, or bytes, include the data and Frame Relay header fields, but they do not include CRC or flag characters.

FDDI

Statistics in this row indicate traffic for FDDI frames and octets on this virtual circuit. Statistics for octets, or bytes, include the data and Frame Relay header fields, but they do not include CRC or flag characters.

IP

Statistics in this row indicate traffic for routed IP, ARP, and Inverse ARP format frames and octets on this virtual circuit. Statistics for octets, or bytes, include the data and Frame Relay header fields, but they do not include CRC or flag characters.

IPX

Statistics in this row indicate traffic for routed IPX format frames and octets on this virtual circuit. Statistics for octets, or bytes, include the data and Frame Relay header fields, but they do not include CRC or flag characters.

BPDU

Statistics in this row indicate traffic for BPDU frames and octets on this virtual circuit. Statistics for octets, or bytes, include the data and Frame Relay header fields, but they do not include CRC or flag characters.

DE Bit

Statistics in this row indicate the number of frames sent and received that have been marked for Discard Eligibility (the DE bit in the frame is set to 1). No statistics are given for Octets in this row. See *Discard Eligibility (DE) Flag* on page 29-9 for more information on the DE bit.

FECN Bit

This value indicates the total number of frames received from the network indicating forward congestion. This occurs when the Frame Relay network sets the frame's Forward Discard Eligibility (FECN) flag. These frames experienced congestion coming over the virtual circuit. Statistics are given only for Frames In for FECN Bit since the Frame Relay network sets it. See *Notification By FECN* on page 29-12 for more information on the FECN bit.

BECN Bit

This value indicates the number of frames received from the network indicating backward congestion. This occurs when the Frame Relay network sets a frame's Backward Discard Eligibility (BECN) flag. These frames observed congestion occurring in the opposite direction during their path over the virtual circuit. Statistics are given only for Frames In since the Frame Relay network sets the BECN bit. See *Notification By BECN* on page 29-11 for more information on the BECN bit.

Discarded

The number of inbound frames that were dropped due to format errors or because the VC was inactive.

Compressed Frames

Statistics in this column indicate traffic for compressed frames on this virtual circuit. Compressed frames are only sent if both sides of a Frame Relay link successfully negotiate for compression (i.e., both must support compression).

Compressed Octets

Statistics in this column indicate traffic for compressed octets on this virtual circuit. Compressed frames are only sent and received if both sides of a Frame Relay link successfully negotiate for compression (i.e., both must support compression). Statistics for octets include the data, Frame Relay header, and Data Compression header fields, but they do not include CRC or flag characters.

Uncompressed Octets

Statistics in this column indicate traffic for uncompressed octets on this virtual circuit. These values apply to the compressed data before compression or just after decompression. Statistics for octets include the uncompressed data and Frame Relay header fields, but they do not include CRC or flag characters.

Compression Ratio

Statistics in this column indicate the compression that was achieved for this type of traffic. For example, in the sample table Outgoing traffic had compression ration of

1.2:1

meaning that each compressed octet is 1.2 uncompressed octets.

Resetting Statistics Counters

You can reset the statistics counters for a single WSX board, a WSX port, or a specific DLCI. The statistics that are cleared on those that are displayed through the **frstatus** commands. The **frclear** command is used to reset statistics.

Resetting Statistics for a WSX Board

To reset statistics on a single WSX board, enter the **frclear** command along with the slot number for the WSX board, as follows:

```
frclear <slot>
```

where **<slot>** is the slot number where the WSX board is installed. For example, if you wanted to clear statistics for the board in slot 4, you would enter:

```
frclear 4
```

or

```
frc 4
```

Resetting Statistics for a WSX Port

To reset statistics on a single WSX port, enter the **frclear** command along with the slot number for the WSX board and the port number as follows:

```
frclear <slot>/<port>
```

where **<slot>** is the slot number where the WSX board is installed and **<port>** is the port number on the WSX board. For example, if you wanted to reset statistics for Port 1 on the WSX module in Slot 4, you would enter:

```
frclear4/1
```

or

```
frc 4/1
```

Resetting Statistics for a Virtual Circuit (DLCI)

To reset statistics on a single virtual circuit, you enter the **frclear** command along with the slot number for the WSX board, the port number, and DLCI number for the virtual circuit on which you want to reset statistics, as follows:

```
frclear <slot>/<port>/<DLCI>
```

where **<slot>** is the slot number where the WSX board is installed, **<port>** is the port number on the WSX board, and **<DLCI>** is the virtual circuit identifier. For example, if you wanted to reset statistics for the board in slot 4, port 1, DLCI 32, you would enter:

```
frclear 4/1/32
```

or

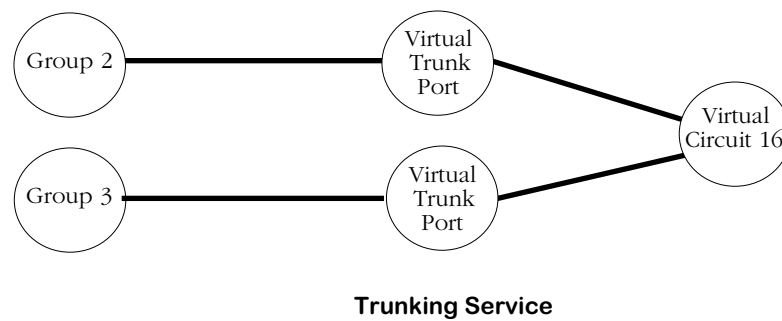


```
frc 4/1/32
```

Managing Frame Relay Services

By default, all virtual circuits on a WSX port have a Bridging service and are assigned to Group 1. The **frmodify** command allows you to change this default bridging service to another Group and to set up a default routing service for the port. See *Setting Configuration Parameters* on page 29-22 for information on the **frmodify** command.

To extend your control over a Frame Relay service, you can use Service menu commands. These command allow you to create and modify bridging, routing, and trunking services by assigning specific virtual circuits and Groups to the services.

Setting up a bridging service requires you to map a virtual circuit to a Group. Setting up a routing service requires you to map one or more virtual circuits to a Group. And setting up a Trunking service requires you to map a single virtual circuit to one or more Groups. The diagrams below illustrate the relationship between Groups, virtual ports and virtual circuits for each Frame Relay service type:



An overview of each type of service and how each operates in a Frame Relay environment can be found earlier in this chapter in the following sections:

- Bridging See *Bridging Services* on page 29-14.
- Routing See *Frame Relay IP Routing* on page 29-15 and *Frame Relay IPX Routing* on page 29-18.
- Trunking See *Trunking* on page 29-19.

The decision to set up one service over another is determined by your network configuration and amount of traffic. In general, you can follow these guidelines:

1. If all your Frame Relay connections are through Omni Switch/Routers, then Trunking is probably the best choice. Trunking is normally set up exclusively for a virtual circuit. No bridging or Routing service needs to be configured on the same virtual circuit where a Trunking service has already been set up.
2. If interoperability is important, then Bridging or Routing is a good choice. In an environment where broadcast traffic is low and high CIRs are deployed, Bridging is a simpler and better choice. In environments with higher broadcast traffic and lower CIRs, Routing is a good solution. However, if you choose to set up a Routing service in an environment with different types of routers, all must support RFC 1490 encapsulation.
3. Bridging and routing services may share a virtual circuit.

The following sections describe how to configure each service type and then how to modify, view, and delete your Frame Relay services.

Configuring a Bridging Service

Frame Relay traffic is automatically bridged for Group 1 in a switch. You can alter this default through two different commands: **frmodify** and **cas**.

The **frmodify** command allows you to change the default Bridging Group from Group 1 to another Group or to turn off bridging completely. This command configures bridging on a port-by-port basis, but does not configure bridging on a virtual circuit basis—all virtual circuits may also be assigned to the Group specified in **frmodify**. See *Modifying a Port* on page 29-22 for more information on the **frmodify** command.

The **cas** command provides more control over bridging service configuration. In addition to naming, enabling and disabling bridging services through **cas**, you can assign specific virtual circuits to a bridging service. Follow the steps below to set up a bridging service through the **cas** command.

1. Enter the **cas** command followed the slot number, a slash (/), the port number, and then the service number for the bridging service:

```
cas 2/3 3
```

A screen similar to the following displays:

```
Slot 1 Port 2 Service 3 Configuration
1) Description ..... = Frame-Relay
   {Enter up to 30 characters}
2) Service Type ..... = Bridging
   {(T)runking, (R)outing, (B)ridging}
3) Administrative Status ..... = Enabled
   {(E)nable, (D)isable}
4) VC(s) ..... = 0
5) VLAN Group(s) ..... = 0
6) Frame-Relay Bridging Mode (Applies to Bridging Only).. = Bridge All
   {Bridge (a)ll, (E)thernet only}
```

```
(save/quit/cancel)
```

```
:
```

You make changes to the options in this screen at the colon prompt (:). You make changes by entering the line number for the option you want to change, an equal sign (=), and then the value for the new parameter.

2. Enter a description of this bridging service by entering 1, an equal sign (=), and then a description for this service. Your description can be up to 30 characters long.

```
1=<bridge service name>
```

When you are done entering a description, press **<Enter>**.

3. Specify that this is a bridging service by entering a 2, an equal sign, and a **B** as follows:

```
2=B
```

This specifies that you want to set up a bridging service, as opposed to a Trunking or Routing service. Press **<Enter>**.

4. By default, the bridging service is Enabled. This means that as soon as you are done configuring the service, it will begin bridging Frame Relay traffic. If you would like to disable this bridging service now and enable it later, enter **3=D** and press **<Enter>**.

5. You need to specify the DLCI for the virtual circuit to include in this bridging service. Only one virtual circuit may be specified for each bridging service. There is a one-to-one mapping between the Group and the virtual circuit. Enter a 4, an equal sign (=), and the DLCI number for the virtual circuit. The example below includes the virtual circuit with DLCI 16 in the bridging service:

4=16

Press **<Enter>**.

6. Specify the Group number that you want to be part of this bridging service. Enter a 5, an equal sign (=), and the Group number. Remember, by default a virtual circuit already bridges on Group 1. The example below includes Group 3 in the bridging service:

5=3

Press **<Enter>**.

7. Indicate whether or not you want frames to be translated on this virtual bridge port. When the **Frame-Relay Bridging Mode** field is set to **Bridge all**, no translation is performed on frames before they are sent out to the Frame Relay network; enter an **A** at this field to select this option.

When the **Frame-Relay Bridging Mode** field is set to **Ethernet only**, non-Ethernet frames are first translated to the default Ethernet frame format for this port before they are sent out to the Frame Relay network. Any MAC translations configured through the Switch menu are valid. Enter an **E** at this field to select this option.

8. Type **save** at the colon prompt (:) and press **<Enter>**. All parameters for this bridging service are saved.

Configuring a WAN Routing Service

There are two main steps to configuring WAN routing for frame relay:

1. Enable and configure routing for a specific WAN Routing group with the **crgrp** command. (Frame Relay Groups are different from other Groups as far as router configurations are concerned.)
2. Set up a WAN routing service through the **cas** command.

Both of these steps are described in the next two sections.

Step 1. Set Up a Frame Relay Routing Group

You enable WAN routing for a Group when you create the Group through the **crgrp** command. The steps for setting up a Group are described in Chapter 24, “Managing Groups and Ports.” Please see that chapter for the generic steps used to create a Group. Also, understand the following points where WAN Groups differ from other Groups.

- During the process of configuring the Group, the **crgrp** command will prompt you with the following prompt:

Enable WAN Routing? (n):

If you want to configure WAN routing on this Group, then you must answer Yes to this prompt. Otherwise, the Group will not be tagged correctly and will not be able to route Frame Relay traffic.

- When configuring IP and IPX Routing, you do not specify a default framing type since Frame Relay routing always uses 1490 encapsulation.
- You do not set up physical interfaces (virtual ports) through the **crgrp** command. All physical mappings for Frame Relay are done through services, as described in Step 2 of this section.

You can configure all virtual circuits to automatically be assigned to the WAN Routing Group you set up in this step. The **frmodify** command contains a parameter, **Default Routing Group**, that you can set to a WAN routing Group. All dynamically learned virtual circuits will automatically be assigned to this Group without any configuration required. See *Modifying a Port* on page 29-22 for more information on the **frmodify** command.

You can also configure a Frame Relay service using the **cas** command as described in *Step 2. Set Up a Frame Relay Routing Service* on page 29-60.

Step 2. Set Up a Frame Relay Routing Service

You create a Frame Relay routing service using the **cas** command. Follow the steps below to set up a routing service.

1. Enter the **cas** command followed the slot number, a slash (/), the port number, and then the service number for the routing service:

```
cas 2/3 1
```

A screen similar to the following displays:

```
Slot 1 Port 2 Service 3 Configuration
1) Description ..... = Frame-Relay
   {Enter up to 30 characters}
2) Service Type ..... = Bridging
   {(T)runking, (R)outing, (B)ridging}
3) Administrative Status ..... = Enabled
   {(E)nable, (D)isable}
4) VC(s)..... = 0
5) VLAN Group(s)..... = 0
6) Frame-Relay Bridging Mode (Applies to Bridging Only).. = Bridge All
   {Bridge (a)ll, (E)thernet only}
```

```
(save/quit/cancel)
```

```
:
```

You make changes to the options in this screen at the colon prompt (:). You make changes by entering the line number for the option you want to change, an equal sign (=), and then the value for the new parameter.

2. Enter a description of this routing service by entering 1, an equal sign (=), and then a description for this service. Your description can be up to 30 characters long.

```
1=<router service name>
```

When you are done entering a description, press **<Enter>**.

3. Specify that this is a routing service by entering a 2, an equal sign, and an **R** as follows:

```
2=5
```

This specifies that you want to set up a routing service, as opposed to a Trunking or Bridging service. Press **<Enter>**.

4. By default, the routing service is Enabled. This means that as soon as you are done configuring the service, it will begin routing Frame Relay traffic. If you would like to disable this routing service now and enable it later, enter **3=D** and press **<Enter>**.
5. You need to specify the DLCIs of the virtual circuits to include in this routing service. Multiple VCs may be configured for a single routing service and all configured VCs will map to a single virtual router port. Enter a 4, an equal sign (=), and then the DLCI numbers for each virtual circuit. Separate DLCIs with spaces, as shown in the example below.

```
4=16 17
```

Press **<Enter>** after you enter all virtual circuit DLCIs.

6. Specify the Group number to which this router port belongs. Enter a 5, an equal sign (=), and the Group number. The example below includes Group 4 in the routing service:

5=4

Press **<Enter>**.

You must have previously configured this Group as a Frame Relay Routing Group through the **crgp** command. If you have not configured the Group for Frame Relay routing, then the following message displays:

Given Vlan Group is not a Frame-Relay Router Group

See the section, *Step 1. Set Up a Frame Relay Routing Group* on page 29-59 for further information on setting up a Frame Relay Group.

7. Disregard the **Frame-Relay Bridging Mode** field. It does not apply to virtual router ports.
8. Type **save** at the colon prompt (:) and press **<Enter>**. All parameters for this bridging service are saved.

Configuring a Trunking Service

To configure a Frame Relay Trunking service, you must use the **cas** command. Perform the following steps:

1. Enter the **cas** command followed the slot number, a slash (/), the port number, and then the service number for the Trunking service:

```
cas 2/3 1
```

A screen similar to the following displays:

```
Slot 1 Port 2 Service 3 Configuration
1) Description ..... = Frame-Relay
   {Enter up to 30 characters}
2) Service Type ..... = Bridging
   {(T)runking, (R)outing, (B)ridging}
3) Administrative Status ..... = Enabled
   {(E)nable, (D)isable}
4) VC(s) ..... = 0
5) VLAN Group(s) ..... = 0
6) Frame-Relay Bridging Mode (Applies to Bridging Only).. = Bridge All
   {Bridge (a)ll, (E)thernet only}
```

```
(save/quit/cancel)
```

```
:
```

You make changes to the options in this screen at the colon prompt (:). You make changes by entering the line number for the option you want to change, an equal sign (=), and then the value for the new parameter.

2. Enter a description of this Trunking service by entering 1, an equal sign (=), and then a description for this service. Your description can be up to 30 characters long.

```
1=<trunk service name>
```

When you are done entering a description, press **<Enter>**.

3. Specify that this is a Trunking service by entering a 2, an equal sign, and a **T** as follows:

```
2=T
```

This specifies that you want to set up a Trunking service, as opposed to a bridging or Routing service. Press **<Enter>**.

4. By default, the Trunking service is Enabled. This means that as soon as you are done configuring the service, it will begin Trunking Frame Relay traffic as you configure it through this menu. If you would like to disable this Trunking service now and enable it later, enter **3=D** and press **<Enter>**.

5. You need to specify the DLCI for virtual circuit that will be used to trunk traffic over the Frame Relay network. Only one virtual circuit may be specified for each Trunking service. Enter a 4, an equal sign (=), and the DLCI number for the virtual circuit similar to the example below:

4=16

Press **<Enter>**.

6. Specify the Group number or numbers that you want to be Trunked over the specified virtual circuit. A separate virtual Trunk port is created for each Group you specify here. Each Group and Trunk port maps down to a single virtual circuit. Enter a 5, an equal sign (=), and the Group number(s). The example below includes Groups 5 and 6 in the trunking service:

5=5 6

Press **<Enter>**.

7. Disregard the **Frame-Relay Bridging Mode** field. It does not apply to virtual trunk ports.
8. Type **save** at the colon prompt (:) and press **<Enter>**. All parameters for this bridging service are saved.

Viewing Frame Relay Services

You can view all Frame Relay services for an entire switch, a single WSX board, or a single WSX port. Use the **vas** command with the following parameters:

```
vas <slot>/<port> <service number>
```

The <slot>, <port> and <service number> parameters are not required but may be specified to narrow the range of the information displayed.

The following is an example of the Frame Relay portion of the **vas** command display:

Frame-Relay Services							
Slot	Port	VCs	Groups	Service Number	Vport	Description	Service Type
====	====	====	====	====	====	====	====
3	2	16	1	1	10	Virtual port (#10)	Bridging
3	3	16	1	1	11	Virtual port (#11)	Bridging
3	3	17	1	2	13	Virtual port (#13)	Bridging
3	2	17	1	2	14	Virtual port (#14)	Bridging
3	3	17	3	3	17	Virtual port (#17)	Routing
3	4	18	2	1	18	Virtual port (#18)	Trunking

The following sections describe the columns in this table.

Slot

The slot number where this WSX module is installed.

Port

The port number to which this service maps. A port may be listed more than once if multiple virtual circuits or multiple services are configured for it. The port is listed for each virtual circuit and for each service. For example, in the sample screen above Port 3 is listed three times—twice as a bridging service for virtual circuits 16 and 17 and again as a routing service for virtual circuit 17.

VCs

The DLCI of the virtual circuit supported by this service. A virtual circuit can be attached to more than one port and be supported by more than one service type.

Groups

The Group or Groups associated with this service. Only one Group is supported by a bridging or routing service. Trunking services may support multiple Groups.

Service Number

Each service for a port is assigned a number. This column lists the number for this service on this particular port. Note that in the sample screen, Port 2 has two services associated with it (Bridging for VC 16 and 17) and Port 3 has three services associated with it (Bridging for VC 16 and 17 and Routing for VC 17).

Vport

The virtual port associated with this service. For bridging services, there is a one-to-one mapping between a virtual port and a virtual circuit. For routing services, multiple virtual circuits may map to a single virtual port. For trunking services, multiple virtual ports can map to a single virtual circuit.

Description

The textual description given to this service when you set it up through the **cas** or **mas** command.

Service Type

A Frame Relay service may be **Bridging**, **Routing** or **Trunking**. All three service types are set up through the **cas** command. Bridging and Routing services may coexist on the same virtual circuit. Trunking cannot coexist with either Bridging or Routing on the same virtual circuit.

Modifying a Frame Relay Service

You can modify previously created Frame Relay services using the **mas** command. The **mas** command uses the same screen as the **cas** command. Simply enter **mas**, the slot, slash (/), port and service number. For example:

```
mas 2/3 1
```

would modify the first service on Port 3 for the WSX board in Slot 2. This command displays the same screen as the **cas** command. See the appropriate section for modifying the service type:

- Bridging See *Configuring a Bridging Service* on page 29-57.
- Routing See *Configuring a WAN Routing Service* on page 29-59.
- Trunking See *Configuring a Trunking Service* on page 29-62.

Deleting a Frame Relay Service

You can delete a Frame Relay service using the **das** command as follows:

1. Enter **das** followed by the slot, port and service number for the Frame Relay service that you want to delete. You can obtain the service number by using the **vas** command. See *Viewing Frame Relay Services* on page 29-64. For example, if you wanted to delete service number 2 for Port 2 on the WSX board in Slot 3, you would enter

das 3/2 2

and the following screen would display:

Frame-Relay Services							
Slot	Port	VCs	Groups	Service Number	Vport	Description	Service Type
====	====	====	====	====	====	====	====
3	2	16	1	1	10	Virtual port (#10)	Bridging
3	2	17	1	2	14	Virtual port (#14)	Bridging

Remove Frame Relay Slot 3 Port 2 Service 2 (n)? :

2. Enter **1** and press **<Enter>** to confirm the deletion of this service. The following messages display confirming the deletion of the service:

Removing Frame Relay Slot 3 Port 2 Service 2, please wait...

Frame Relay Slot 3 Port 2 Service 2 removed

30 Point-to-Point Protocol

The Point-to-Point Protocol (PPP) provides a standard method for transporting multi-protocol datagrams over point-to-point links. The base protocol is specified in RFC 1661. Many other RFCs define additional capabilities for network protocol negotiation, management information databases (MIBs), and PPP operation over different kinds of serial channels.

PPP is comprised of three main components. The first component is a method of encapsulating multi-protocol datagrams so that the underlying protocol can be identified; the second component is the Link Control Protocol (LCP) that is used for establishing, configuring, and testing the datalink connection; the third component is a family of Network Control Protocols (NCPs) that are used for establishing and configuring different network-layer protocols such as IP and IPX.

The implementation of PPP for the Omni Switch/Router WAN Switching Modules supports bridging, IP routing and IPX routing. Data compression of the PPP packets is also supported when the WSX module contains a STAC 9705 Data Compression Coprocessor.

PPP Connection Phases

There are five phases to a PPP connection: Dead, Establish, Authenticate, Network, and Terminate:

Dead. The first phase is called the “Dead” phase because the physical channel has not yet been activated.

Establish. After the physical channel has been activated, the PPP connection enters the second phase, called “Establish,” wherein it attempts to negotiate link-level parameters and options using the Link Control Protocol (LCP). This phase ends when the LCP enters its own “open” state.

Authenticate. After LCP has reached its “open” state, the PPP connection enters the phase called “Authenticate” wherein it tries to identify the peer with which it is attempting to establish a connection. If the authentication option is enabled, either the Password Authentication Protocol (PAP) or the Challenge Handshake Authentication Protocol (CHAP) is used to perform the authentication. If authentication is not enabled, the PPP connection proceeds to the next phase, “Network.”

Network. After the “Authenticate” phase is successful (or when it is not enabled), the PPP connection proceeds to the next phase, called “Network,” wherein the network protocols are negotiated using the appropriate Network Control Protocol (NCP). For example, to negotiate the use of IP over the PPP connection, the Internet Protocol Control Protocol (IPCP) is used. The details of the negotiation are specific to each network protocol, but may include such tasks as assigning network layer addresses. A network layer protocol must be negotiated successfully before the exchange of protocol packets can proceed; but, once negotiated, the protocol can begin to freely exchange packets. The PPP connection spends most of its time in the “Network” phase, because this is where the active transmission of data occurs.

Terminate. The final phase of a PPP connection is called the “Terminate” phase. This phase begins when authentication is unsuccessful or the channel becomes inoperative. Very often, this phase is simply bypassed, and PPP will return to the idle (Dead) phase when a channel is disconnected.

Data Compression

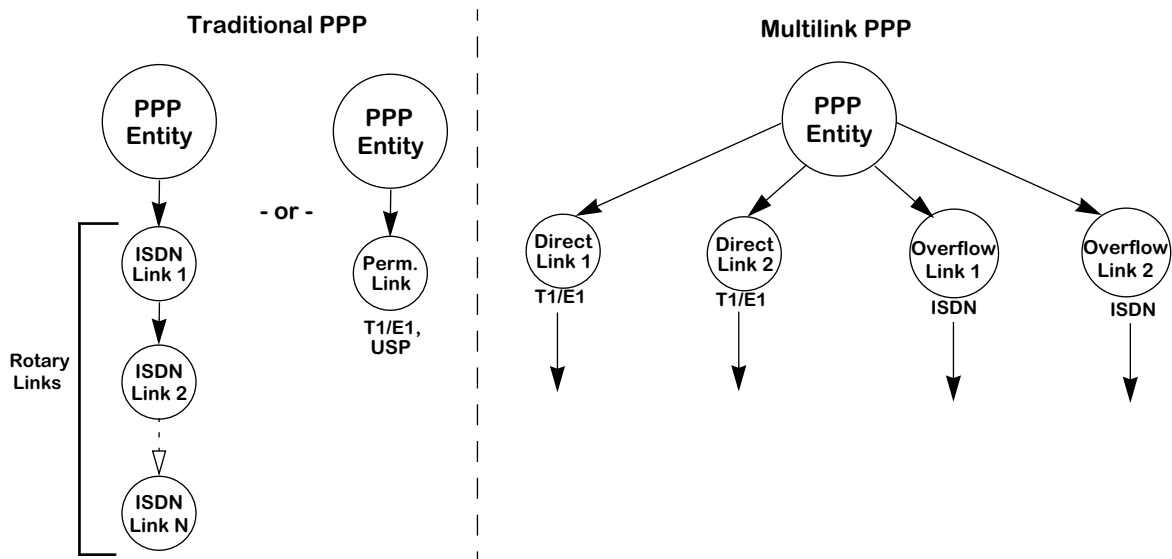
RFC 1974 specifies the use of STAC-LZS compression with PPP. Data compression allows the payload of a PPP packet, including the protocol ID, to be compressed, saving valuable bandwidth. Compression is negotiated during the Network phase using Compression Control Protocol (CCP), which includes the negotiation of a data compression algorithm and any parameters specific to the algorithm. Once negotiated, all data packets (i.e., non-control protocol packets) from all successfully negotiated protocols are compressed before transmission. The compression algorithm negotiated includes any mechanism for synchronizing the compressor and decompressor.

STAC-LZS's maximum data compression ratio is 30:1. The LZS algorithm is optimized to compress all file types as efficiently as possible. Even string matches as short as two octets are effectively compressed. The STAC-LZS compression algorithm supports both single compression history communication and multiple compression history communication.

Often, many streams of information are interleaved over the same link. Each virtual link will transmit data that is independent of other virtual links. Using multiple compression histories can improve the compression ratio of a communication link.

Multi-Link PPP

The main limitation of PPP is implicit in its name: Point-to-Point Protocol, meaning that it is limited to connecting two points over a single physical connection. Multi-Link PPP (MLPPP) extends the functionality of PPP by combining multiple PPP links into a single logical data pipeline, called a "bundle." Unlike standard PPP, MLPPP is not limited to individual links; both physical and virtual connections can be bundled.



Traditional vs. Multilink PPP

Multilink Modes of Operation

Multilink PPP supports combinations of both permanent and switched connections. This results in two possible modes of operation:

- permanent connection only
- switched connection only

Note

One important thing to remember when setting up multilinks is that all links to be bundled must exist on the same slot.

Permanent Connection Only

This mode allows multiple links to be joined into a single bundle. Permanent connections can be universal serial ports or fractional T1/E1 ports.

Switched Connection Only

This mode supports only switched connections. The only switched connections currently supported are ISDN calls. This allows multiple switched connections to be joined into a single bundle. In this mode, the first call is initiated as a demand connection, if a frame is available for the peer, or a backup connection, if the primary link becomes inactive, according to the configuration of the ISDN link.

◆ Note ◆

ISDN MLPPP bundles are limited to 2 B-channels

PPP Fragmentation Interleaving

The PPP Fragmentation/Interleaving functionality creates two prioritized virtual streams within a single PPP connection. The lower priority stream uses an MLPPP header to sequence the frames while the higher priority stream uses a standard PPP header without MLPPP sequence numbers. The lower priority stream is fragmented according to maximum delay parameters so that a higher priority frame can be injected in the middle of the low priority frame and not have to wait for the entire low priority frame to be transmitted.

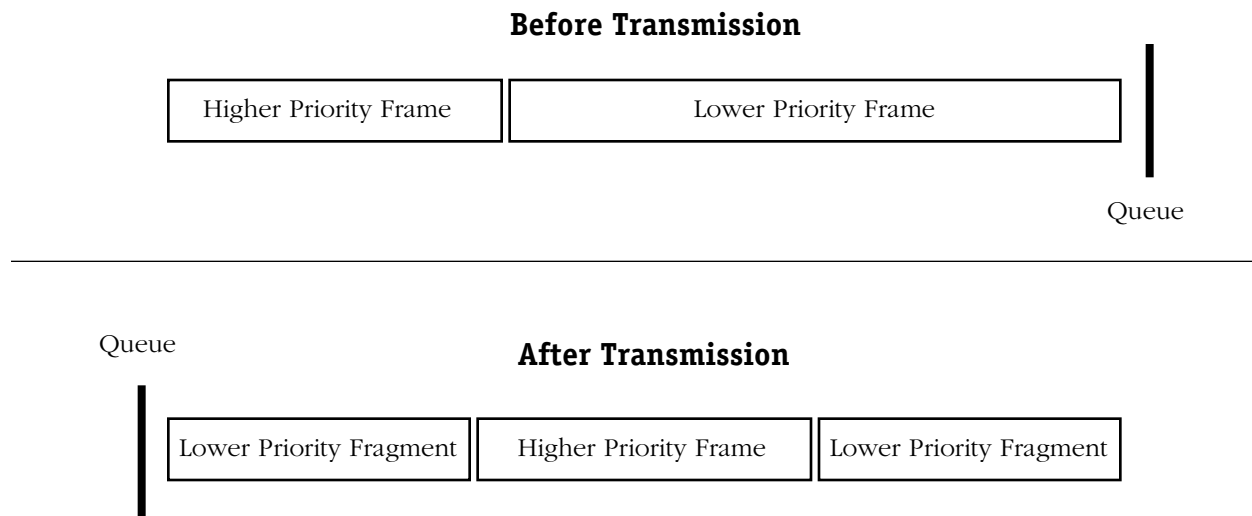
On the transmitting side, when low priority frames are being transmitted they are divided into multiple fragments. The size of each fragment is determined by the configured maximum delay parameter and the speed of the physical interface. The fragments are encapsulated with a standard MLPPP header, which contains a sequence number to identify lost fragments and beginning and ending flags to identify frame boundaries. When high priority frames are transmitted they are sent whole with standard PPP headers.

The delay of the high priority frame is the time it takes to finish transmitting the current frame or fragment plus the time it would take to transmit any other high priority frames in queue. On the OA-512, which has a hardware based high and low priority transmit queues, the high priority frame would be sent as soon as the current fragment/frame is finished. On the WSX, which has software based high and low priority transmit queues, it depends on how many frame/fragments have been committed to transmit buffer descriptors ahead of it. As part of the WSX transmit data flow improvements, the number of frames/fragments committed to buffer descriptors will be kept to a minimum, but because the queues are software based, will not be able to match the delays of the OA-512.

On the receiving side, as low priority frames are being received they will be put into the MLPPP reassembly queue, as supported by the existing software. As complete frames are received they will be forwarded to the normal PPP processing. When high priority frames are received, since they will always be sent complete, they will immediately be forwarded to the normal PPP processing.

The only configurable parameter that has been added is the maximum delay. The feature is enabled when this parameter is set to a non-zero value. A flag has been utilized to force a 16 fragment maximum for a fragmented frame to make this process compatible with Cisco products.

The following diagram illustrates this concept:



PPP Fragmentation Process

Overview of PPP Configuration Procedures

The configuration of a PPP connection on your switch is divided into three separate tasks. This three-phase strategy was chosen to allow PPP connections to be configured over *any* serial channel interface without requiring the use of multiple PPP configuration displays for each separate type of interface.

Step 1. Configure the Physical Interface to be Used for PPP

The information configured at the physical interface level includes the specification of the type of WSX interface and of any information that is specific to the given type of interface. The interfaces that can support PPP are ISDN, T1/E1, and the Universal Serial Port on all WSX boards.

An ISDN interface (WSX-BRI) requires the specification of the switch type, the local telephone number, and the Service Profile Identifiers (SPIDs) if appropriate for the switch type. The UI commands used to configure ISDN interfaces allow for modifying and viewing ISDN port's configuration and the display of its operational status. See Chapter 32 titled "Managing ISDN Ports" for detailed information on configuring an ISDN interface for PPP.

The configuration of a T1/E1 interface is described in Chapter 33 titled “Managing T1 and E1 Ports.”

The configuration of a universal serial port (USP) on a WSX-S board is described in Chapter 28 titled “Managing WAN Switching Modules.”

Step 2. Configure the Operation of PPP Itself

The information configured at the PPP level includes the remote and local user IDs and passwords, network protocol information, the use of data compression, and retry and delay information to be used during PPP connection establishment with LCP. The UI commands used to configure PPP connections (called “PPP Entities”) allow for the adding, modifying, and viewing of PPP connections and their operational status. This chapter describes the configuration of PPP Entities (connection configurations) using the **pppadd**, **pppmodify**, **pppdelete**, **pppview**, and **pppstatus** commands.

Step 3. Configure a Link Between the Physical Interface and PPP

As mentioned above, three kinds of physical interfaces can support PPP connections: Universal Serial Ports (on all WSX boards,), T1/E1 channels (on the WSX-FT1/E1 board), and ISDN lines (on the WSX-BRI board).

The “WAN Links” used to support PPP connections vary somewhat, depending upon which type of physical interface is being used for PPP. When the physical interface is a Universal Serial Port (USP) or a fractional T1/E1 channel (which are permanent channels), the port is dedicated to the PPP connection and the “WAN Link” simply identifies the physical interface in terms of the slot and port. When the physical interface being used is an ISDN interface (which provides dynamic, switched connections), the “WAN Link” identifies the numbering information that is to be used to establish the serial connection and the slot/port if necessary. The UI commands used to configure WAN Links allow for the adding, modifying, and viewing of the links, and the display of their operational status. See Chapter 31 titled “WAN Links” for detailed information on the commands used to configure WAN Links.

Multiple links can be configured when employing Multilink PPP, one for each link in the bundle. For Multilink PPP over ISDN, each link configured for a PPP entity is called every time the connection is attempted and Multilink PPP is successfully negotiated. For normal PPP over ISDN, when a connection with a PPP entity is attempted, each link is called until one is successful.

The PPP Submenu

The WAN menu contains a submenu, named **PPP**, containing commands specific to the Point-to-Point-Protocol (PPP).

To display the **PPP** menu, enter the following commands:

```
PPP
?
```

A screen similar to the following displays:

Command	PPP Menu
pppglobal	Add PPP Global configuration record
pppadd	Add PPP configuration record
pppmodify	Modify PPP configuration record
pppdelete	Delete PPP configuration record
pppview	View PPP configuration record(s)
pppstatus	Get Status of PPP configuration records and associated links

Main **File** **Summary** **VLAN** **Networking**
Interface **Security** **System** **Services** **Help**

PPP Configuration Overview

Your first configuration step is to create a global PPP configuration record using the **pppglobal** command. This global record is used to provide default settings to be used for incoming calls. Then, you can add individual PPP configuration records (called “PPP Entities”) for each peer (i.e., for each remote site) with which you wish to be able to establish a point-to-point connection. You will need to know specific information about the remote peers with which you wish to connect in order to successfully configure the PPP Entity.

After you have configured at least one PPP Entity, you can use the other commands on the PPP Menu to modify, delete, view, and display its operational status. You can then add PPP Entities as you need them to support additional PPP connection requirements.

When a port is configured for PPP via the **wpm** command, a PPP entity and a WAN link entry are created automatically. For more information, see Chapter 28 titled “Managing WAN Switching Modules.”

Setting Global PPP Parameters

The **pppglobal** command is used to set global configuration parameters that are used by the PPP protocol. These parameters are termed “global” because they are the default settings used by the switch to establish connections with incoming calls. These global settings are not tied to a specific peer (i.e., a PPP Entity; see *Adding a PPP Entity* on page 30-9).

To set the global PPP parameters, enter the following command:

```
pppglobal
```

A screen similar to the following displays:

```
PPP Global Configuration:
```

- ```

1) Default Authentication Type PAP
 {(N)one, (P)AP, (C)HAP}
2) Global User ID sent to remote for Authentication. =
 {8 characters userid}
3) Global Password sent to remote for Authentication =
 {8 characters password}
4) Default Compression Type = STAC-LZS
 {(N)one, STAC-(L)ZS}
5) Default Bridge Config Admin Status = Disabled
 {(E)nable, (D)isable}
6) Default IP Config Admin Status..... = Enabled
 {(E)nable, (D)isable}
7) Default IPX Config Admin Status = Disabled
 {(E)nable, (D)isable}

```

```
(save/quit/cancel)
```

```
:
```

The fields on this screen have the following meanings:

### Default Authentication Type

Specifies the type of authentication that is to be expected on incoming calls. The options are **None**, **PAP**, and **CHAP**. Set this parameter to the type of authentication that you expect your callers to be using. If you enable either PAP or CHAP authentication, the next two parameters must also be set (user ID and password) or the caller’s connection requests will be refused. If you set this parameter to **None**, you must also set the Default Bridge, IP and IPX Configuration Administration Status parameters or the caller’s connection requests will be refused.

### Global User ID sent to remote for Authentication

Specifies the user ID that will be sent to a peer on incoming calls. Enter the text you will transmit on incoming calls. This parameter must contain a value if either PAP or CHAP authentication is being used. The User ID and password received from the peer will be checked against the list of peers (PPP Entities) to attempt to identify the remote peer.

### Global Password sent to remote for Authentication

Specifies the password that will be sent to a peer on incoming calls. Enter the text you will transmit on incoming calls. This parameter must contain a value if either PAP or CHAP authentication is being used.

### **Default Compression Type**

Specifies the type of compression that is to be expected on incoming calls. The options are **None** and **STAC-LZS**. If you set this parameter to **None** and your callers are using compression, the caller's connection request may be refused. See *Data Compression* on page 30-2 for a description of STAC-LZS data compression.

### **Default Bridge Config Admin Status**

Specifies whether the bridging function is to be negotiated for incoming calls. More information on the bridging function can be found in *Adding a PPP Entity* on page 30-9. If this parameter is disabled here, but disabled on the caller, the caller's connection request may be refused.

### **Default IP Config Admin Status**

Specifies whether the IP routing function is to be negotiated for incoming calls. More information on the IP routing function can be found in *Adding a PPP Entity* on page 30-9. If this parameter is disabled here, but disabled on the caller, the caller's connection request may be refused.

### **Default IPX Config Admin Status**

Specifies whether the IPX routing function is to be negotiated for incoming calls. More information on the IPX routing function can be found in *Adding a PPP Entity* on page 30-9. If this parameter is disabled here, but disabled on the caller, the caller's connection request may be refused.

## Adding a PPP Entity

The **pppadd** command is used to add a PPP Entity configuration record. The PPP Entities you create are identified by numbers called Peer IDs. When you enter the **pppadd** command, you may enter a Peer ID number with the command like this:

```
pppadd <ID number>
```

Alternatively, you can enter the command alone and you will be prompted for a Peer ID. The prompt will identify the next available, unique ID number.

After you enter the **pppadd** command as described above, a screen will be displayed that contains the configuration parameters that make up the PPP Entity. The steps that begin below will take you through the process of adding a PPP Entity.

After you have set the PPP Entity's configuration parameters, you must save them to actually create the PPP Entity. After saving, you will be prompted to add one or more links to be used with the PPP Entity. In other words, the software will automatically issue a **linkadd** command for you. This was designed to help you to quickly create working PPP Entities as they must be associated with at least one link in order to operate. The **linkadd** command, as well as the other commands on the Link menu, are described in Chapter 31 titled "WAN Links."

1. To add a PPP Entity, enter the following command:

```
pppadd
```

A screen similar to the following will display:

```
Add PPP configuration record. Please specify a unique
ID number to identify this record and the remote Peer to communicate with.
```

```
Peer ID (1):
```

This prompt is asking you to enter a Peer ID as well as indicating that the next available number is 1. If other Peers have already been configured, the number indicated will be different than is shown above.

2. To answer the prompt, for example, for Peer ID 1, you would enter the following command:

```
1
```

If you have enabled the verbose mode, you will see the following text immediately before the prompts:

```
To change a value, enter the corresponding number, an '=', and the new
value. For example to set a new description, use
: 2=My new Description
To clear an entry specify the value as '.' as in
2=.
When complete enter "save" to save all changes, or cancel or Ctrl-C to
cancel all changes. Enter ? to view the new configuration.
```

This text provides brief help on entering commands at the following screens. In the steps that follow below, this help text will *not* be shown.

A screen similar to the following will display:

```
Adding PPP configuration record for Peer ID: 1
Enter PPP parameters:

1) Description :
 {Enter text up to 30 characters}
2) Administrative Status Enabled
 {(E)nabled, (D)isable}
3) PPP Mode Normal
 {(N)ormal, (M)ultilink}
4) Compression Type None
 {(N)one, STAC-(L)ZS}
5) Bridging Group 1
 {1-65535 or 0 if no Bridging}
50) Bridge Config Admin Status Enabled
 {(E)nabled, (D)isable}
51) PPP Bridging Mode Ethernet Only
 {Bridge (A)ll, (E)thernet Only}
6) Routing Group 0
 {1-65535 or 0 if no Routing}
7) Authentication Type None
 {(N)one, (P)AP, (C)HAP}
8) Max Failure Count 3
 {1..65535}
9) Max Configure Count 3
 {1..65535}
10) Max Terminate Count 3
 {1..65535}
11) Retry Timeout Value 5
 {Retry Timeout in Second(s) 1..65535}
12) Fragmentation Interleaving No
 {Fragmentation Interleaving Yes or No}

(save/quit/cancel)
:
```

The prompts for Bridging, Routing, Authentication, and Fragmentation Interleaving (numbered 5, 6, 7, and 12 above), contain suboptions that are displayed only if you have enabled those features. These expanded menus are shown below in the relevant sections describing the UI fields.

3. When you have made the changes you need to the prompts on this screen, enter the following command to save the PPP Entity:

```
save
```

The following prompt will display:

```
Normal (non-multilink) PPP configuration record created.
Do you wish to define the link at this time y/n (y):
```

If you answer yes to this prompt, a **linkadd** command will be automatically executed for this PPP Entity. For complete details on using the **linkadd** command, see the relevant section in Chapter 31, entitled “WAN Links.”

If you answer No to this prompt, a message will appear indicating that the link was not added, but the PPP Entity itself was added.

**Note**

You can add the link needed for a PPP Entity later if you decide not to do so now. The automatic execution of the **linkadd** command is done here only as a convenience to you.

The fields on the **pppadd** configuration screen have the following meanings:

**Description**

A textual description for this PPP Entity. You can enter any text you like (up to 30 characters).

**Administrative Status**

Indicates the Administrative Status of this PPP Entity. **Enabled** will allow the PPP Entity to operate. **Disabled** will disable the PPP Entity without deleting it.

**PPP Mode**

Can be set to either **Multilink** or **Normal** (single PPP connection).

**Compression Type**

Controls whether this PPP Entity will perform compression. The one type of compression currently available is STAC-LZS. See *Data Compression* on page 30-2 for details on STAC-LZS compression.

**Bridging Group**

Indicates the VLAN Group to be used for PPP Bridging. A value of zero (0) indicates that this PPP Entity will not perform a bridging service and will discard all bridged format packets received or transmitted. The suboptions under this heading are:

***Bridge Config Admin Status***

Used to enable or disable the bridging function for this PPP Entity.

***PPP Bridging Mode***

Used to select the operational mode for bridging. The options are **Ethernet**, which will enable bridging on Ethernet interfaces only, or **All**, which enables it for all interfaces.

**Routing Group**

Indicates the VLAN Group to be used for PPP Routing of the IP and IPX protocols. A value of zero (0) indicates that this PPP Entity will not perform a routing service and will discard all routed format packets received or transmitted.

Enabling Routing expands the menu with the following suboptions:

- 6) **Routing Group** ..... 1  
{1-65535 or 0 if no Routing}
- 60) **IP Config Admin Status** ..... Enabled  
{(E)nabled, (D)isable}
- 61) **Remote IP Address (Only valid if IP is enabled)** ..... 0.0.0.0  
{Valid IP address notation e.g., x.x.x.x}
- 62) **IPX Config Admin Status** ..... Enabled  
{(E)nabled, (D)isable}

The suboptions under this heading are:

### ***IP Config Admin Status***

Used to enable or disable the routing of IP packets over PPP. The options are **Enabled** and **Disabled**.

### ***Remote IP Address (Only valid if IP is enabled)***

Used to specify the Remote IP address of the PPP connection when IP routing is enabled. Valid IP address notation must be used. If this parameter is set to 0.0.0.0 and IP routing is enabled, the Remote IP address will be learned during Internet Protocol Control Protocol (IPCP) negotiation.

### ***IPX Config Admin Status***

Used to enable or disable routing of IPX packets over PPP. The options are **Enabled** and **Disabled**.

## **Authentication Type**

Indicates the type of authentication to be used by this PPP Entity. The options are **None**, **PAP**, and **CHAP**.

Enabling Authentication expands the menu with the following suboptions:

- 7) **Authentication Type** ..... PAP  
{(N)one, (P)AP, (C)HAP}
- 70) **User ID received from remote for Authentication** ...  
{8 characters userid}
- 71) **Password received from remote for Authentication** .  
{8 characters password}
- 72) **User ID sent to remote for Authentication** .....  
{8 characters userid}
- 73) **Password sent to remote for Authentication** .....  
{8 characters password}

The suboptions under this heading are:

### ***User ID received from remote for Authentication***

Used to specify the User ID to be expected from the remote end during PAP or CHAP authentication.

### ***Password received from remote for Authentication***

Used to specify the password to be expected from the remote end during PAP or CHAP authentication.

**User ID sent to remote for Authentication**

Used to specify the User ID to be sent to the remote end during PAP or CHAP authentication. This parameter is used only for outgoing calls. Incoming calls use the global defaults (see *Setting Global PPP Parameters* on page 30-7 for details).

**Password sent to remote for Authentication**

Used to specify the password to be sent to the remote end during PAP or CHAP authentication. This parameter is used only for outgoing calls. Incoming calls use the global defaults (see *Setting Global PPP Parameters* on page 30-7 for details).

**Max Failure Counter**

The maximum number of times a CONFIGURATION\_REQUEST packet will be sent when the previous attempts received responses, but did not receive a CONFIGURATION\_ACK. This counter applies to all LCP and NCP negotiations.

**Max Configure Counter**

The maximum number of times a CONFIGURATION\_REQUEST packet will be sent when the previous attempts did not receive any responses. This counter applies to all LCP and NCP negotiations.

**Max Terminate Counter**

The maximum number of TERMINATE\_REQUEST packets that will be sent without receiving a TERMINATE\_ACK packet. This counter applies to all LCP and NCP negotiations.

**Retry Timeout Value**

Indicates the number of seconds to wait between CONFIGURATION\_REQUEST retries that do not receive a response. This timeout value applies to all LCP and NCP negotiations.

**Fragmentation Interleaving**

Fragmentation Interleaving allows you to break up lower priority packets into smaller pieces and insert higher priority packets inbetween. This is useful when sending time-critical information streams such as voice or video data.

Enabling Fragmentation Interleaving expands the menu with the following suboptions:

- 12) Fragmentation Interleaving ..... No  
     {Fragmentation Interleaving Yes or No
- 121) Fragmentation Delay ..... 0  
         {Fragmentation Delay in milliseconds(ms)}
- 122) Limit Maximum number of fragmentation to 16 ..... No  
         {(Y) meant cisco compatible. (N) meant native}

**Fragmentation Delay**

This field specifies a millisecond count for determining when to fragment a PPP packet. If higher priority data will remain in the queue for over the set amount of time, then the packet is fragmented.

### ***Limit Maximum number of fragments to 16***

This flag is set to make the interface compatible with Cisco products. When set to **Yes**, a PPP packet is never fragmented into more than 16 smaller packets.



## Modifying a PPP Entity

The **pppmodify** command is used to modify the parameters of an existing PPP Entity. To modify a specific PPP Entity, for example Peer ID 1, enter the following command:

```
pppmodify p1
```

A screen similar to the following displays:

```

Modify PPP for communication to Peer ID: 1
Enter PPP parameters:
1) Description :
 {Enter text up to 30 characters}
2) Administrative Status Enabled
 {(E)nable, (D)isable}
3) PPP Mode Normal
 {(N)ormal, (M)ultilink}
4) Compression Type None
 {(N)one, STAC-(L)ZS}
5) Bridging Group 1
 {1-65535 or 0 if no Bridging}
50) Bridge Config Admin Status Enabled
 {(E)nabled, (D)isable}
51) PPP Bridging Mode Ethernet Only
 {Bridge (A)ll, (E)thernet Only}
6) Routing Group 2
 {1-65535 or 0 if no Routing}
60) IP Config Admin Status Enabled
 {(E)nabled, (D)isable}
61) Remote IP Address 0.0.0.0
 {IP address or 0.0.0.0 = learn, if IP enabled}
62) IPX Config Admin Status Disabled
 {(E)nable, (D)isable}
7) Authentication Type PAP
 {(N)one, (P)AP, (C)HAP}
70) User ID received from remote for Authentication ...
 {0 (No ID) to 8 ASCII characters}
71) Password received from remote for Authentication .
 {0 (No Password) to 8 ASCII characters}
72) User ID sent to remote for Authentication
 {0 (No ID) to 8 ASCII characters}
73) Password sent to remote for Authentication
 {0 (No Password) to 8 ASCII characters}
8) Max Failure Count 3
 {1..65535}
9) Max Configure Count 3
 {1..65535}
10) Max Terminate Count 5
 {1..65535}
11) Retry Timeout Value 5
 {Retry Timeout in Second(s) 1..65535}
12) Fragmentation Interleaving No
 {Fragmentation Interleaving Yes or No}
121) Fragmentation Delay 0
 {Fragmentation Delay in milliseconds(ms)}
122) Limit Maximum number of fragmentation to 16 No
 {(Y) meant cisco compatible. (N) meant native}
:

```

The fields on this screen are the same as those produced by the **pppadd** command. See *Adding a PPP Entity* on page 30-9 for descriptions of each of these fields.

Make the desired changes to any of the parameters, then enter the **save** command to implement the changes. You will then be returned to the system prompt.

## Viewing PPP Entity Configurations

The `pppview` command is used to view the configuration parameters of existing PPP Entities.

### Displaying the Configuration of All PPP Entities

To view configuration information on all PPP Entities, enter the following command:

```
pppview
```

A screen similar to the following displays:

PPP Configuration for Chassis:

| Peer ID | Admin Status | Mode      | Authentication | Compression | Bridging Group | Routing Group |
|---------|--------------|-----------|----------------|-------------|----------------|---------------|
| 1       | UP           | Normal    | None           | None        | 1              | 0             |
| 2       | DN           | Multilink | PAP            | STAC-LZS    | 1              | 2             |
| 3       | UP           | Normal    | CHAP           | None        | 0              | 2             |

The fields on this screen have the following meanings:

#### Peer ID

The number assigned to this PPP Entity when it was added. Used to identify a specific PPP Entity that you want to examine with the `pppview` or `pppstatus` commands.

#### Admin Status

Indicates the Administrative Status of this PPP Entity. **UP** means that this entity is enabled, or operative. **DN** means that this entity is disabled, or inoperative.

#### Mode

Indicates whether **Normal** or **Multilink** operation is used by this PPP Entity. Multilink operation is described under the heading *Multi-Link PPP* on page 30-2.

#### Authentication

Indicates the type of authentication used by this PPP Entity. The options are **None**, **PAP** and **CHAP**. These are two well-established standards currently used for PPP authentication.

#### Compression

Indicates the type of data compression configured to operate with this PPP Entity. The options are **None** or **STAC-LZS**. See *Data Compression* on page 30-2 for information on STAC-LZS compression.

#### Bridging Group

Indicates the VLAN Group to be used for PPP Bridging. A value of zero (0) indicates that this PPP Entity will not perform a bridging service and will discard all bridged format packets received or transmitted.

## Routing Group

Indicates the VLAN Group to be used for PPP Routing of the IP and IPX protocols. A value of zero (0) indicates that this PPP Entity will not perform a routing service and will discard all routed format packets received or transmitted.

## Displaying the Configuration of a Specific PPP Entity

To view configuration information on a *specific* PPP Entity, you must enter a Peer ID number with the **pppview** command. For example, to examine Peer ID 1, you would enter the following command:

```
pppview p1
```

A screen similar to the following displays:

```
View PPP configuration record for communication to Peer ID: 1
1) Description : Entry Peer ID 1
2) Administrative Status Enabled
3) PPP Mode Normal
4) Compression Type Disabled
5) Bridging Group 1
 50) Bridge Config Admin Status Enabled
 51) PPP Bridging Mode Ethernet Only
6) Routing Group 1
 60) IP Config Admin Status Enabled
 61) Remote IP Address 0.0.0.0
 62) IPX Config Admin Status Disabled
7) Authentication Type PAP
 70) User ID received from remote for Authentication ...
 71) Password received from remote for Authentication .
 72) User ID sent to remote for Authentication
 73) Password sent to remote for Authentication
8) Max Failure Count 3
9) Max Configure Count 3
10) Max Terminate Count..... 3
11) Retry Timeout Value..... 5
12) Fragmentation Interleaving No
 {Fragmentation Interleaving Yes or No
 121) Fragmentation Delay 0
 {Fragmentation Delay in milliseconds(ms)}
 122) Limit Maximum number of fragmentation to 16 No
 {(Y) meant cisco compatible. (N) meant native}
```

The fields on this screen are similar to those produced by the **pppadd** command. A few differences are noted in the descriptions that are given below. Note that you cannot make changes to the parameters on this screen. To do so, you must use the **pppmodify** command instead (see *Modifying a PPP Entity* on page 30-15 for complete information).

## Displaying PPP Entity Status

The `pppstatus` command is used to view the operational status of one or more PPP Entities.

### Displaying the Status of All PPP Entities

To view the operational status of *all* PPP Entities, enter the following command:

```
pppstatus
```

A screen similar to the following displays:

| Peer ID | Admin State | Mode      | IP Oper State | IPX Oper State | BCP Oper State | CCP Oper State |
|---------|-------------|-----------|---------------|----------------|----------------|----------------|
| 1       | UP/UP       | Normal    | Open          | Close          | Open           | Open           |
| 2       | UP/UP       | Multilink | Open          | Open           | Open           | Open           |

The fields on this screen have the following meanings:

#### Peer ID

The number assigned to this PPP peer.

#### Admin State

Indicates the Administrative Status of this PPP Entity. **UP** means that this entity is enabled, or operative. **DN** means that this entity is disabled, or inoperative.

#### Mode

Indicates whether **Normal** or **Multilink** operation is used by this PPP Entity. Multilink operation is described under the heading *Multi-Link PPP* on page 30-2.

#### IP Oper State

Indicates the operational state of the IP Routing option. **Open** means that IP has successfully negotiated a connection and is able to pass IP packets. **Closed** means that IP has not yet reached the open state, and is therefore unable to pass IP packets. The reasons why the state may be closed are: 1) the call has been disconnected, 2) the protocol is in the process of making a connection, or 3) the IP Routing option was not configured.

#### IPX Oper State

Indicates the operational state of the IPX Routing option. **Open** means that IPX has successfully negotiated a connection and is able to pass IPX packets. **Closed** means that IPX has not yet reached the open state, and is therefore unable to pass IPX packets. The reasons why the state may be closed are: 1) the call has been disconnected, 2) the protocol is in the process of making a connection, or 3) the IPX Routing option was not configured.

### BCP Oper State

Indicates the operational state of the Bridging Control Protocol option. **Open** means that the bridging operation is active. **Closed** means that the bridging operation has not yet reached the open state. The reasons why the state may be closed are: 1) the call has been disconnected, 2) the protocol is in the process of making a connection, or 3) the Bridging option was not configured.

### CCP Oper State

The operational state of the compression control protocol option. **Open** means that compression is active. **Closed** means that compression has not reached the open state. The reasons why the state may be closed are: 1) the call has been disconnected, 2) the protocol is in the process of making a connection, or 3) the compression option was not configured.

## Displaying the Status of a Specific PPP Entity

To view both the operational status and the relevant statistics of a specific PPP Entity, for example, Peer ID 1, enter the following command:

```
pppstatus p1
```

A screen similar to the following displays:

```

PPP statistics for Peer ID: 2
Admin IP IPX BCP CCP
State Mode Oper Oper Oper Oper
===== =====
UP Normal Open Close Open Close

LCP Pkts IPCP Pkts IPCP Pkts BCP Pkts CCP Pkts
IN/OUT IN/OUT IN/OUT IN/OUT IN/OUT
===== =====
3/4 2/2 0/0 4/4 0/0

 Packets Packets Packets Octets Octets
 In Out In+Out In Out %In %Out
 =====
Total 2232 1475 3707 91751 66034
Ethernet 0 146 146 0 13413 0 20
8025 0 0 0 0 0 0 0
FDDI 0 0 0 0 0 0 0
IP 79 158 237 7784 6952 8 10
IPX 0 0 0 0 0 0 0
BPDU 2153 1171 3324 83967 45669 91 69

STAC-LZS Compressed Compressed Uncompressed Compression
Compression Frames Octets Octets Ratio
 =====
In 0 0 0 0.0:1
Out 0 0 0 0.0:1
IN+Out 0 0 0 0.0:1

```

The additional fields produced by the **pppstatus** command when a specific Peer ID is entered with the command are as follows:

### LCP Pkts IN/OUT

The total number of Link Control Protocol (LCP) packets received (**In**) and transmitted (**Out**) on this PPP connection.

### **IPCP Pkts IN/OUT**

The total number of IP Control Protocol (IPCP) packets received (**In**) and transmitted (**Out**) on this PPP connection.

### **IPCP Pkts IN/OUT**

The total number of IP Control Protocol (IPCP) packets received (**In**) and transmitted (**Out**) on this PPP connection.

### **BCP Pkts IN/OUT**

The total number of BCP packets received (**In**) and transmitted (**Out**) for this PPP connection.

### **CCP Pkts IN/OUT**

The total number of CCP packets received (**In**) and transmitted (**Out**) for this PPP connection.

Also shown on this screen are two tables of statistics. The first table shows various data transmission statistics shown both as a total and sorted by the type of frame encapsulation being used (**Total**, **Ethernet**, **8025**, **FDDI**, **IP**, **IPX**, and **BPDU**). The columns in the first table show the following information for each type of frame encapsulation: the number of packets received (**Packets In**), the number of packets transmitted (**Packets Out**), the sum of received and transmitted packets (**Packets In+Out**), the number of octets received (**Octets In**), the number of octets transmitted (**Octets Out**), and the percentages received (**%In**) and transmitted (**%Out**) for each type of frame encapsulation.

The second table shows statistics related to the performance of STAC-LZS compression sorted by **In**, **Out**, and **In+Out** categories. The column headings show the number of compressed frames and octets, the number of uncompressed frames and octets, and the overall compression ratio represented by the previous figures.

## Deleting a PPP Entity

The **pppdelete** command is used to delete an existing PPP Entity.

1. Before you can delete a PPP Entity, you must first delete all the links associated with it. You do so using the **linkdelete** command (see Chapter 31 titled “WAN Links”). If you try to delete a PPP Entity that still has links associated with it, the following message will be displayed:

```
Delete PPP Peer ID: 1 aborted because the following link(s) attach to it.
Link Index: 1, Description: Link Entry: 1, Peer ID: 1
```

2. To delete a specific PPP Entity (after deleting all links associated with it), enter the Peer ID number along with the **pppdelete** command. For example, to delete Peer ID 2, enter the following command:

```
pppdelete p2
```

A screen similar to the following displays:

```
This will delete the configuration for PPP Peer ID: 2
Continue ? {(Y)es, (N)o} : N
```

3. To delete this entity, enter **y** and press **Enter**. If you decide to cancel out of the deletion, press **Enter** to accept the default answer of No. The system prompt will reappear.





# 31 WAN Links

## Introduction

This chapter describes the procedures for configuring a “WAN link” between an already created PPP Entity (see Chapter 30, *Point-to-Point Protocol*) and the physical interface that will be used to carry PPP traffic. The procedures described in this chapter comprise the third and final step in the three-step process for configuring the operation of PPP on your Omni Switch/Router (the complete three-step process was also described in Chapter 30).

Here is a brief review of the PPP configuration process: the first step is to configure the physical interfaces that will carry PPP traffic. The second step is to configure the operation of PPP itself by creating “PPP Entities.” The third step is to configure the “link” between an existing PPP Entity and the physical interface that will be used to carry PPP traffic (hence the name “WAN Links”).

## Configuring WAN Interfaces

Three kinds of physical WAN interfaces can support PPP connections: serial ports (WSX-S), T1/E1 channels (WSX-FT1/E1), and ISDN lines (WSX-BRI). The “WAN Links” you create to support PPP connections vary somewhat, depending upon the type of physical interface being used. When the physical interface being used is a Universal Serial Port (USP) or a fractional T1/E1 channel (which are permanent channels), the port is dedicated to the PPP connection and the “WAN Link” simply identifies the physical interface in terms of the slot and port. When the physical interface being used is an ISDN interface (which provides dynamic, switched connections), the “WAN Link” identifies the numbering information that is to be used to establish the serial connection and the slot/port if necessary.

### Related Hardware Chapters

The configuration of an ISDN interface is described in Chapter 32, *Managing ISDN Ports*. The configuration of a T1/E1 interface is described in Chapter 33, *Managing T1 and E1 Ports*. The configuration of a universal serial port (USP) on a WSX-S board is described in Chapter 28, *Managing WAN Switching Modules*. The ISDN WSX board (WSX-BRI) also contains a USP; this port on the WSX-BRI board may be configured in a similar manner to the USP ports on the WSX-S board.

## The Link Submenu

The WAN menu contains a submenu named **link** which contains commands for creating the WAN Links needed to support the Point-to-Point Protocol (PPP) over various hardware interfaces. WAN links can either be “fixed” (i.e., configured for a serial port or T1/E1 port), or dial-based (i.e., configured for an ISDN port). The link UI commands also provide a means of modifying and viewing existing WAN Links and displaying their operational status.

To switch to, and to display, the **link** menu, enter the following commands:

```
link
?
```

A screen similar to the following displays:

| <b>Command</b>    | <b>Link Menu</b>                                   |
|-------------------|----------------------------------------------------|
| <b>linkadd</b>    | <b>Add a Link configuration entry</b>              |
| <b>linkmodify</b> | <b>Modify an existing Link configuration entry</b> |
| <b>linkdelete</b> | <b>Delete an existing Link configuration entry</b> |
| <b>linkview</b>   | <b>View configuration of WAN Link(s)</b>           |
| <b>linkstatus</b> | <b>Status of WAN Link(s)</b>                       |

|                  |                 |                |                 |                   |
|------------------|-----------------|----------------|-----------------|-------------------|
| <b>Main</b>      | <b>File</b>     | <b>Summary</b> | <b>VLAN</b>     | <b>Networking</b> |
| <b>Interface</b> | <b>Security</b> | <b>System</b>  | <b>Services</b> | <b>Help</b>       |

Each of the commands on this menu is described in the following sections.

## Adding a WAN Link

The **linkadd** command is used to add link configuration records, or “WAN Links” to the switch. This command defaults to a WSX physical port (serial or Fractional T1/E1). When the **linkadd** command is used to create links over WSX ports, all of the parameters needed to create the link are contained on one screen. However, when you select to create a link over an ISDN port, a second screen will be displayed after you enter and save the initial parameters on the first screen.

The first subheading (*Adding WSX Port Links*) below shows the sequence of screens when creating a link over a WSX port. The second subheading below (*Adding ISDN Call Links* on page 31-4) shows the sequence of screens when creating a link over an ISDN port.

### Adding WSX Port Links

1. To add a link over a WSX port, you must enter a Peer ID number (associated with a “PPP Entity”) with the command. See Chapter 30, *Point-to-Point Protocol*, for details on creating Peer IDs.

For example, to create a link for Peer ID 1, enter the following command (where **p1** is the Peer ID number):

```
linkadd p1
```

A screen similar to the following displays:

```
Adding Link for Peer ID 1, Link Index 1:
```

- ```

1) Description : Link Entry: 2, Peer ID: 1
   {Enter text up to 31 characters}
2) Administrative Status ..... = Enabled
   {(E)nabled, (D)isabled}
3) Link Type..... = WSM Port
   {(W)SM Port, (I)SDN call}
4) Link Slot..... = 0
   {Slot number or 0 if not tied to a slot}
5) Link Port ..... = 0
   {Port number or 0 if not tied to a port}

```

```
(save/quit/cancel)
```

```
:
```

To alter a parameter, enter the line number for the parameter, followed by an equal sign (=), then the new value. For example, to change the **Link Type** (line 3) from WSX Port to ISDN call, you would enter:

```
3=l
```

When you have completed configuring parameters, enter **save**. Your new values will be saved and you will exit this menu. If you want to exit this menu without saving changes, simply enter **quit** or **cancel**.

The fields on this screen have the following meanings:

Description

A textual description used to identify this WAN Link. The default text indicates the link entry number and the Peer ID number.

Administrative Status

Sets the Administrative Status of this WAN Link. The options are “**Enabled**,” which will enable this link and “**Disabled**,” which will disable the link but not delete it.

Link Type

Specifies the type of physical connection that will carry the link. The options are “**WSM Port**,” which means a serial or Fractional T1/E1 connection and “**ISDN**,” which means an ISDN call will be used to make the connection.

Link Slot

Specifies the switch slot number to be used by this WAN Link.

Link Port

Specifies the switch port number to be used by this WAN Link.

2. To make a change to the values for any of the fields on this screen, enter the field's line number followed by the desired value.
3. To add the link for a WSX port, you must specify which switch port and slot is to be used. To do so, you must make changes to the values for items 4 and 5. For example, if your WSX port is in slot 5, port 2, you would enter the following three commands:

```
: 4=5
: 5=2
: save
```

After entering the **save** command, you will be returned to the system prompt.

Adding ISDN Call Links

1. To create a link over ISDN, you must enter a Peer ID number (associated with a “PPP Entity”) with the command. See Chapter 30, *Point-to-Point Protocol*, for details on creating Peer IDs.

For example, to create a link for Peer ID 1, you would enter the following command (where **p1** is the Peer ID number):

```
linkadd p1
```

A screen similar to the following displays:

```
Adding Link for Peer ID 1, Link Index 1:
```

- 1) **Description : Link Entry: 2, Peer ID: 1**
{Enter text up to 31 characters}
- 2) **Administrative Status** = Enabled
{(E)nabled, (D)isabled}
- 3) **Link Type** = WSM Port
{(W)SM Port, (I)SDN call}
- 4) **Link Slot** = 0
{Slot number or 0 if not tied to a slot}
- 5) **Link Port** = 0
{Port number or 0 if not tied to a port}

```
(save/quit/cancel)
```

```
:
```

- You must now change the Link Type to ISDN. To do so, enter the following commands:

```
: 3=l
: ?
```

A screen similar to the following displays:

```
1) Link Description :
   {Enter text up to 31 characters}
2) Link Administrative Status ..... = Enabled
   {(E)nabled, (D)isabled}
3) Link Type..... = ISDN Call
   {(W)SM Port, (I)SDN call}
4) Link Slot..... =
   {Slot number or 0 if not tied to a slot}
5) Link Port ..... =
   {Port number or 0 if not tied to a port}

(save/quit/cancel)
:
```

- You must now enter the ISDN slot and port numbers that will be used by this WAN Link. For example, to use slot 4, port 2, you would enter the following commands:

```
: 4=4
: 5=2
```

Note

Incoming and backup ISDN calls may dynamically select and use any available slot and port. However, you *must* specify an ISDN slot and port for the link when you first create its WAN Link.

A screen similar to the following displays:

```
Modify ISDN call record configuration. Peer ID: 1 Link Index: 1
Type: ISDN Call Slot: 4 Port: 2

1) Link Description : Link Entry: 1, Peer ID: 1
   {Enter text up to 30 characters}
2) Link Administrative Status ..... = Enabled
   {(E)nabled, (D)isabled}
3) Inactivity Timer ..... = 0
   {1-9999 seconds or 0 if disabled}
4) Minimum call duration ..... = 0
   {1-9999 seconds or 0 if disabled}
5) Maximum call duration ..... = 0
   {1-9999 seconds or 0 if disabled}
6) Outgoing Calls ..... = Enabled
   {Enable, Disable}
   60) Call Originate Mode ..... = On-Demand
      {(O)n-Demand or (B)ackup}
   61) Carrier Delay Timeout ..... = 0
      {Call completion timeout 1-999 seconds}
   62) Maximum Call Retries ..... = 1
      {Retry call count, 0 if infinite}
   63) Retry Delay ..... = 0
      {Seconds between retry attempts, 0 = retry immediately}
   64) Failure Delay ..... = 0
      {Secs after max calls failed to retry,
       0 = don't retry after max calls failed}
   65) Remote Phone Number ..... =
      {digits 0 through 9}
   66) Desired Calling Speed ..... = 64000
      {56000, 64000}
7) Incoming calls ..... = Enabled
   {Enabled, Disabled}

(save/quit/cancel)
:
```

The fields on this screen have the following meanings:

Link Description

A textual description used to identify this WAN Link.

Link Administrative Status

Sets the Administrative Status of this WAN Link. The options are “**Enabled**,” which will allow the link to operate and “**Disabled**,” which will disable the link without deleting it.

Inactivity Timer

Sets the time period (in seconds) after which the connection will be terminated if it is not carrying useful data. “Useful data” refers to forwarding packets (routing information), but not to encapsulator maintenance frames. An entry of zero (0) specifies no disconnection due to inactivity. The Inactivity Timer is disabled for outgoing backup calls, and should be disabled by the user for incoming calls that are used to backup.

Minimum Call Duration

The minimum duration of a call, in seconds, starting from the time the call is connected until the call is disconnected. If you enable this field by entering a nonzero value, the Inactivity Timer will be disabled until the time set in the Minimum Call Duration field has passed.

Maximum Call Duration

The maximum call duration in seconds. An entry of zero (0) means “unlimited.”

Outgoing Calls

Sets whether outgoing calls can be made by this WAN Link. The option “**Enabled**” will allow the link to make outgoing calls while “**Disabled**” will not allow the link to make outgoing calls. These suboptions further specify the details of the outgoing calls:

Call Originate Mode

Specifies whether the call is to be initiated on demand or only when operating as a backup to another link.

Carrier Delay Timeout

The amount of time, in seconds, allowed for a call to be completed.

Maximum Call Retries

The number of calls to a non-responding address that may be made. An entry of zero (0) means there is no limit to the number of retries. The intent of this parameter is to limit the number of successive calls to an address which is inaccessible or which refuses those calls. Some countries regulate the number of call retries to a given peer that can be made.

Retry Delay

The time, in seconds, between call retries if a peer cannot be reached. An entry of zero (0) means that call retries may be done without any delay.

Failure Delay

The time, in seconds, after which call attempts are to be made again after a peer has been noticed to be unreachable (i.e., after the limit set in **Maximum Call Retries** has been reached). An entry of zero (0) means that a peer will not be called again after the maximum number of unsuccessful call attempts has been made.

Remote Phone Number

The phone number that is to be dialed in order to make the connection. Only one phone number can be associated with a single WAN Link. You can add other WAN Links if you want to use multiple phone numbers.

Desired Calling Speed

The desired calling speed. The options are 56000 and 64000 bits/second. You should set this parameter to the maximum speed supported by the telephone switch to which you will be connecting.

Incoming Calls

Sets whether incoming calls are to be accepted by this WAN Link. “**Enabled**” will allow the link to accept calls. “**Disabled**” will not allow the link to accept calls.

4. You must now enter a value in at least the **Remote Phone Number** field under **Outgoing Calls**. If you do not make an entry in this field, an error will be returned by the system when you attempt to save and exit the screen.

You can also make changes to any of the other fields on this screen if they are needed to provide ISDN call information this WAN Link. The default settings should suit many situations; however, you will need to determine what information will be needed to support your ISDN calls and make the appropriate entries in the fields on this screen.

5. Enter the **save** command when you are ready to create the WAN Link.
The system prompt will then reappear.

Modifying a WAN Link

The **linkmodify** command is used to modify the parameters of an existing WAN Link. Different parameters will be displayed by the command based on the type of link. The first subheading (*Modifying ISDN Links*) below shows the sequence of screens when modifying a link over a WSX port. The second subheading below (*Modifying WSX Links* on page 31-10) shows the sequence of screens when modifying a link over an ISDN port.

Note

The Slot and Port fields in an existing WAN Link record cannot be modified. To change them, you must delete the record then create a new record.

Modifying ISDN Links

- To modify a WAN Link, you must enter its Link Index with the command. For example, to modify Link Index 1 which uses ISDN, you would enter the following command:

```
linkmodify L1
```

A screen similar to the following displays:

```
Modify ISDN call record configuration. Peer ID: 1 Link Index: 1
Type: ISDN Call Slot: 5 Port: 1

1) Link Description : Link Entry: 1, Peer ID: 1
   {Enter text up to 30 characters}
2) Link Administrative Status ..... = Enabled
   {(E)nabled, (D)isabled}
3) Inactivity Timer ..... = 30
   {1-9999 seconds or 0 if disabled}
4) Minimum call duration ..... = 0
   {1-9999 seconds or 0 if disabled}
5) Maximum call duration ..... = 0
   {1-9999 seconds}
6) Outgoing Calls ..... = Enabled
   {Enable, Disable}
60) Call Originate Mode ..... = On-Demand
    (O)n-Demand or (B)ackup}
61) Carrier Delay Timeout ..... = 0
    {Call completion timeout 1-999 seconds}
62) Maximum Call Retries ..... = 0
    {Retry call count, 0 if infinite}
63) Retry Delay ..... = 0
    {Seconds between retry attempts, 0 = retry immediately}
64) Failure Delay ..... = 0
    {Secs after max calls failed to retry,
     0 = don't retry after max calls failed}
65) Remote Phone Number ..... =
    {digits 0 through 9}
66) Desired Calling Speed ..... = 64000
    {56000, 64000}
7) Incoming calls ..... = Enabled
   {Enabled, Disabled}

(save/quit/cancel)
:
```

The fields on this screen are the same as those produced by the **linkadd** command. See *Adding a WAN Link* on page 31-3 for descriptions of each of these fields.

2. Make the desired changes to each of the fields on this screen, then enter the **save** command to implement your changes.

The system prompt will then reappear.

Modifying WSX Links

1. To modify a WAN Link, you must enter its Link Index with the command. For example, to modify Link Index 2 which uses a WSX physical port (serial or Fractional T1/E1), you would enter the following command:

linkmodify L2

A screen similar to the following displays:

Modify Serial Port Link configuration. Peer ID: 2 Link Index: 2
Type: WSM port Slot: 5 Port: 1

- 1) **Link Description : Link Entry: 2, Peer ID: 2**
{Enter text up to 30 characters}
- 2) **Link Administrative Status = Enabled**
{(E)nabled, (D)isabled}

(save/quit/cancel)
:

The fields on this screen are the same as those produced by the **linkadd** command. See *Adding a WAN Link* on page 31-3 for descriptions of each of these fields.

2. Make the desired changes to the fields on this screen, then enter the **save** command to implement the changes.

The system prompt will then reappear.

Deleting WAN Links

The **linkdelete** command is used to delete one or more existing WAN Link records.

Note

Before you can delete a PPP Entity, you must first delete all WAN Links that have been associated with it. See *Deleting a PPP Entity* in Chapter 30 for complete information.

1. To delete an existing WAN Link, for example, Link Index 2, you would enter the following command:

```
linkdelete L2
```

A screen similar to the following displays:

```
This will delete the configuration for Link Peer ID: 3 Link Index: 2  
Continue ? {(Y)es, (N)o} : N
```

2. If you wish to delete this link, enter **y** and press **Enter**. If you wish to abort the deletion, just press **Enter** to accept the default answer of “No.”

The system prompt will then reappear.

Viewing WAN Links

The `linkview` command is used to view information on existing WAN Link records.

Displaying All Existing WAN Links

To view information on all existing WAN Links, enter the following command:

```
linkview
```

A screen similar to the following displays:

```
List of ISDN Port Type:
Peer Link  Link  Link  Link  Outgoing  Incoming  Peer  Inac.  Min/Max  Call
Id   Index Mode  Slot Port  Called Num. Caller Id.  Speed  Timer  Dur.    Retry
=====
   1    1  DEM   4    2    7145555555 8015551212 56000   0    0/0    0
   2    2  BKP   4    2    7145551212 8015555555 64000   0    0/0    0

List of WSM Port Type:
Peer Link  Link  Link
Id   Index Slot  Port
=====
   1    1    5    2
   2    2    5    4
```

The fields on this screen have the following meanings:

Peer ID

The number assigned to the PPP Entity that is related to this WAN Link. You assign this number when you create the PPP Entity (see Chapter 30, *Point-to-Point Protocol*, for more information on creating PPP Entities).

Link Index

The number assigned by the system to this WAN Link; used to identify the link in the table.

Link Mode

Indicates whether this WAN Link is on-demand (“DEM”) or back-up (“BKP”). On-demand links are brought up only when data is ready to be sent. Backup links are brought up when a primary link fails.

Link Slot

The number of the physical switch slot that is to be used for this connection.

Link Port

The number of the physical switch port that is to be used for this connection.

Outgoing Called Number

The phone number that is to be dialed in order to establish the connection.

Incoming Caller ID

The phone number reported by the Caller ID service, if available.

Peer Speed

The specified calling speed for this link. The options are 56000 and 64000 bits/second.

Inactivity Timer

Specifies the time period (in seconds) after which the connection will be terminated if it is not carrying useful data. "Useful data" refers to forwarding packets (routing information) but not encapsulator maintenance frames. Zero (0) specifies no disconnection due to inactivity.

Min/Max Duration

The minimum and maximum duration of a call, in seconds, starting from the time the call is connected until the call is disconnected. Zero (0) means "unlimited."

Call Retry

The number of calls that may be made to a non-responding address. A count of zero (0) means there is no limit to the number of call retries.

Displaying Information for a Specific WAN Link

To view detailed information on a *specific* WAN Link, you must enter its Link Index with the command. Different parameters will be displayed based on the type of link being used.

Example of an ISDN Link

To examine an ISDN link, for example, Link 1, you would enter following command:

```
linkview L1
```

A screen similar to the following displays:

```
View ISDN call record configuration. Peer ID: 1 Link Index: 1
Type: ISDN Call Slot: 5 Port: 2

1) Link Description :
2) Link Administrative Status ..... = Enabled
3) Inactivity Timer ..... = 30
4) Minimum call duration ..... = 0
5) Maximum call duration ..... = 0
6) Outgoing Calls ..... = Enabled
60) Call Originate Mode ..... = On-Demand
61) Carrier Delay Timeout ..... = 0
62) Maximum Call Retries ..... = 1
63) Retry Delay ..... = 0
64) Failure Delay ..... = 0
65) Remote Phone Number ..... = 7145551212
66) Desired Calling Speed ..... = 64000
7) Incoming calls ..... = Enabled
```

The fields on this screen provide the same information as those on the **linkadd** screen. See *Adding a WAN Link* on page 31-3 for descriptions of each of these fields.

Example of WSX Serial or T1/E1 Link

An example of a link over a WSX serial or Fractional T1/E1 port would look like this:

View ISDN Link configuration. Index: 2 Link Peer ID: 3
Type: WSM port Slot: 5 Port: 2

- 1) **Link Description : Link Entry: 1, Peer ID: 1**
{Enter text up to 31 characters}
- 2) **Link Administrative Status = Enabled**
{(E)nabled, (D)isabled}

The fields on this screen provide the same information as those on the **linkadd** screen. See *Adding a WAN Link* on page 31-3 for descriptions of each of these fields.

Displaying Link Status

The `linkstatus` command is used to display the operational status of WAN Links.

Displaying Status for All WAN Links

To view information on all WAN Links, enter the following command:

```
linkstatus
```

A screen similar to the following displays:

Link Idx	Peer Id	Slot/Port	Last Setup Time
1	1	4/2	00:00:00 03/97
2	2	4/2	00:00:00 03/97

Active Session:						
Setup Time	Link Index	Peer Id	Peer Call Addr.	Conn. Time	Call St.	Call Org.
00:00	1	1	8188783500	00:00	CON	ANS
00:00	2	1	8188783500	00:00	CON	ANS

The fields on this screen have the following meanings:

Link Index

The number assigned to identify this WAN Link.

Peer ID

The number assigned to the PPP Entity that is related to a WAN Link (indicated by Link Index).

Slot/Port

The slot and port numbers associated with a given Link Index and Peer ID.

Last Setup Time

The value of "sysUpTime" (the time of day recorded by the switch) when the last call to this peer was started. For ISDN, this will be the time when the setup message was received from or sent to the network. This field will be updated whenever a call is started or answered.

Active Session

The following information is available for the active ISDN session, if one is in progress:

Setup Time

The value of "sysUpTime" (the time of day) when the call to this peer was started.

Peer Id

The Peer ID that is related to this active ISDN session.

Peer Call Address

The number to which this call is connected. Zero (0) means the number is not available.

Connection Time

The value of “sysUpTime” (the time of day) when the call was connected. Zero (0) means the call is not currently connected.

Call State

The current call state. The possible entries are **IDLE** (meaning there is no active call), **CONT** (meaning the call is in the process of connecting), **CONN** (meaning the call is connected), **ACTX** (meaning the call is active), **DISC** (meaning the call has been disconnected), and **UNKN** (meaning that the state is unknown).

Call Origination

The call origin. Possible entries are **OUTG** (meaning the call was outgoing) and **INCM** (meaning the call was incoming).

Displaying Status for a Specific WAN Link

To view detailed status information on a *specific* WAN Link, you must enter its Link Index with the command.

For example, to examine Link 1 (an ISDN link), you would enter following command:

```
linkstatus L1
```

A screen similar to the following displays:

```
      Status for Link Index: 1
Connect Time ..... 0
Success Calls ..... 0
Failed Calls ..... 0
Accepted Calls ..... 0
Refused Calls ..... 0
Last Setup Time ..... 12:56:00 3/96
```

The fields on this screen have the following meanings:

Connect Time

Accumulated connect time to the peer since system start-up. This is the total connect time, i.e., the connect time for outgoing calls plus the time for incoming calls.

Success Calls

The number of completed calls to the Peer ID related to this WAN Link.

Failed Calls

The number of failed call attempts, or any reason, to the Peer ID related to this WAN Link since system start-up.

Accepted Calls

The number of calls from the Peer ID related to this WAN Link accepted since system start-up.

Refused Calls

The number of calls from the Peer ID that were refused, or any reason, since system start-up.

Last Setup Time

The value of “sysUpTime” (the time of day) when the last call to this peer was started. For ISDN, this will be the time when the setup message was received from or sent to the network. This field will be updated whenever a call is started or accepted.

32 Managing ISDN Ports

The WAN Switching Module for the Basic Rate Interface (WSX-BRI) supports 1 or 2 Universal Serial Ports (USP) and 1 or 2 ISDN Basic Rate Interfaces (BRI). The USPs can support Frame Relay or Point-to-Point Protocol (PPP). The BRI interface can support only PPP.

The Universal Serial Port on a WSX-BRI board is operationally identical to the USPs found on the 4- or 8-port WSX-S board. The ISDN BRI port is an RJ-45 connector. The BRI port can be configured either as a “U” interface for the North American market or as an “S/T” interface for international markets. The WSX-BRI board also supports hardware data compression via the STAC 9705 Data Compression Coprocessor.

The ISDN BRI interface supports switched connections, usually through a central office switch. Connections can be established when data is available for a remote peer, referred to as “demand” mode, or when a primary circuit is inactive, referred to as “backup” mode.

Overview of ISDN

Integrated Services Digital Network (ISDN) is a switched network that incorporates a digital connection to the central office (the local loop), instead of the current telephone network’s analog connection. Because the worldwide telephone network is becoming increasingly digital in the trunks between switching centers, the incorporation of ISDN allows for end-to-end switched digital connections. In general, there are three main goals for ISDN:

- provide end-to-end digital connectivity
- support a wide range of services, both voice and non-voice
- access the ISDN by a limited set of standard user-to-network interfaces

Basic Rate Interface (BRI) Versus Primary Rate Interface (PRI)

There are two methods defined for accessing ISDN. The Basic Rate Access (BRA) method, commonly known as the Basic Rate Interface (BRI), was intended for residential subscribers and small offices. The Primary Rate Access (PRA) method, commonly known as the Primary Rate Interface (PRI), was intended for users with greater data-transfer capacity requirements, such as offices with a digital PBX. The Omni Switch/Router WSX-BRI board supports only the BRI interface. Future products may be introduced that include support for PRI interfaces.

The WSX-BRI interface terminates at an ISDN-capable switch in the central telephone office. In order to perform properly, the WSX-BRI board must know to which type of telephone switch it is being connected. You must provide your Omni Switch/Router with this information during configuration of the WSX-BRI board. Also, depending upon the type of telephone switch you will be accessing, you may need to obtain from the telephone company a Service Profile Identification (SPID). The SPID is used in North America for DMS100, ATT 5ESS and Nation ISDN 1switch types.

“U”, “S/T”, and “R” Interfaces

The ISDN specification defines a limited set of user-to-network interfaces, including reference points for the BRI access method. The following are the main BRI reference points:

U Interface. The U interface is a two-wire (single pair) interface that supports full-duplex data transfer from the phone switch. Only a single device can be connected to a U interface. This device is called a Network Termination 1 (NT1) which converts the U interface to the S/T interface (described below). The U interface is used in North America. Elsewhere in the world, telephone companies supply the NT1 service, allowing customers the use of S/T interfaces.

S/T Interface. The S/T interface is a four-wire, bus interface on which multiple (up to eight) ISDN access devices can be attached to gain shared access to ISDN's data channels. The S/T interface is the most commonly-used interface in Europe.

R Interface. The R interface is a general reference point at which non-ISDN devices can gain access to an ISDN network through a device called a Terminal Adapter (TA). A Terminal Adapter typically converts various standard interfaces, such as RS232 and V.35, to the S/T bus.

The “B,” “D,” and “H” Channels

ISDN supports three types of data channels: the “B” channel, the “D” channel and the “H” channel. The line encoding and framing structure for each type of channel varies among the U, S/T, and R interfaces and for different access methods. A brief description of the three channels follows:

B Channel. The B channel is used for the transfer of information, which can be any type of data that the endpoints agree on, such as digitized voice, digitized video or packet data. The B channel operates at 64 kbps on both BRI and PRI interfaces, but is commonly rate-adapted to 56 kbps in North America to accommodate switching system limitations. A single BRI interface consists of one D channel operating at 16 kbps and two B channels operating at 64 kbps (or 56 kbps in North America).

D Channel. The D channel operates at 16 kbps on BRI (64 kbps on PRI) and is used for carrying common-channel signaling. The D channel is used both to establish and maintain circuit-switched calls on the B channels. The D channel can also be used to carry low-speed packet-switched data (the Omni Switch/Router does *not* support such usage).

H Channel. The H channel, supported *only* on PRI interfaces, is used to transfer information at higher bit rates by aggregating B channels. The four implementations of the H channel are: H0 (384 kbps, 6 B channels), H10 (1472 kbps, 23 B channels), H11 (1536 kbps, 24 B channels), and H12 (1920 kbps, 30 B channels). The use of the H channel is *not* supported by the Omni Switch/Router because this channel requires a PRI interface.

The ISDN Submenu

The WAN menu contains a submenu, **ISDN**, containing commands specific to WSX-BRI ISDN ports.

To switch to, and to display, the **ISDN** menu, enter the following commands:

```
ISDN
?
```

A screen similar to the following displays:

Command	ISDN Menu
isdnm	Modify an existing ISDN port's configuration
isdnd	Delete an existing ISDN configuration entry
isdnv	View an existing ISDN configuration entry
isdns	Status for the ISDN configuration entry
Main	File
Interface	Security
	Summary
	System
	VLAN
	Services
	Networking
	Help

Switch Configuration

This section describes how to configure the ISDN ports on WSX-BRI boards. You use the **isdnm** command to modify the configuration of an ISDN port. You must select the correct type of telephone switch to which you will be making your ISDN calls, as well as supply signalling calling addresses (phone numbers) and SPIDS, if required. Configuration is described in the next section, *Modifying an ISDN Configuration Entry* on page 32-4.

The other commands on the ISDN submenu are described in the remaining sections of this chapter.

Modifying an ISDN Configuration Entry

The **isdnm** command is used to modify the parameters for a selected ISDN port. These parameters are typically provided by the telephone carrier or other service provider at the time the ISDN line is installed.

1. To modify a specific port, for example in Slot 4, Port 1, enter the following command:

```
isdnm 4/1
```

A screen similar to the following displays:

```
1) Switch Type ..... ETSI
   {5(ES)S, (D)MS100, (NI)1, (ET)SI}
2) B1 Signalling Calling Address..... 8185551212
   {Phone Number}
3) B1 Service Profile Identifier (SPID) ..... 123456789
   {9-20 Numeric characters}
4) B2 Signalling Calling Address..... 7145551212
   {Phone Number}
5) B2 Service Profile Identifier (SPID) ..... 123456789
   {9-20 Numeric characters}
```

```
(save/quit/cancel)
:
```

The fields on this screen have the following meanings:

Switch Type

Specifies the type of switch to which this ISDN port is to be connected. The options are: AT&T 5ESS (**5ESS**), Northern Telecom DMS100 (**DMS100**), National ISDN-1 Bellcore (**NI1**), and Euro-ISDN ETS 300/British Telecom NET3 (**ETSI**).

B1/B2 Signalling Calling Address

The number assigned to this channel by the carrier. If only one address is supplied by the carrier, assign it to channel B1, and leave channel B2 empty.

B1/B2 Service Profile Identifier (SPID)

The Service Profile Identifier assigned to this channel by the carrier. Normally, this value contains the calling address surrounded by some digits. If only one address is supplied by the carrier, assign it to channel B1, and leave channel B2 empty.

Important Note

When changing the **Switch Type** or adding/deleting **SPIDs**, reboot the switch to implement the changes.

Deleting an ISDN Configuration Entry

The **isdnd** command is used to delete one or more ISDN configuration entries. Deleting the configuration entry is equivalent to returning the ISDN port to its default settings. Although you cannot delete a physical ISDN port from the switch, you can remove the configuration entry that was recorded for a port.

1. To delete a specific ISDN entry, for example, for a board in slot/port 2/2, you would enter the following command:

```
isdnd 2/2
```

A screen similar to the following displays:

```
This will delete Slot 2, Port 2.  
Continue ? {(Y)es, (N)o} : N
```

2. To delete this entry, enter **y** and press **Enter**. To abort the deletion, press **Enter** to accept the default answer of “No.” The system prompt will then reappear.

Important Note

After deleting an ISDN configuration entry, you should reboot the switch to implement any configuration changes you make using the **isdnm** command.

Viewing an ISDN Configuration Entry

The **isdnv** command is used to view the configuration of existing ISDN configuration entries. You can either view a configuration summary for all ISDN ports on a specified slot, or display the configuration for a single ISDN port.

To view configuration information on all ISDN ports on a specific slot, for example, slot 4, enter the following command:

```
isdnv 4
```

A screen similar to the following displays:

```
View ISDN Configuration for Slot: 4, Port: 2.
1) Switch Type ..... ETSI
2) B1 Signalling Calling Address. ....
3) B1 Service Profile Identifier (SPID) .....
4) B2 Signalling Calling Address. ....
5) B2 Service Profile Identifier (SPID) .....

View ISDN Configuration for Slot: 4, Port: 4.
1) Switch Type ..... ETSI
2) B1 Signalling Calling Address. ....
3) B1 Service Profile Identifier (SPID) .....
4) B2 Signalling Calling Address. ....
5) B2 Service Profile Identifier (SPID) .....
```

To view information on a specific ISDN port and slot, for example, slot 4, port 4, enter the following command:

```
isdnv 4/4
```

A screen similar to the following displays:

```
View ISDN Configuration for Slot: 4, Port: 4.
1) Switch Type ..... ETSI
2) B1 Signalling Calling Address. ....
3) B1 Service Profile Identifier (SPID) .....
4) B2 Signalling Calling Address. ....
5) B2 Service Profile Identifier (SPID) .....
```

The fields on this screen are the same as those produced by the **isdnm** command:

Switch Type

Specifies the type of switch to which this ISDN port is to be connected. The options are: AT&T 5ESS (**5ESS**), Northern Telecom DMS100 (**DMS100**), National ISDN-1 Bellcore (**NI1**), and Euro-ISDN ETS 300/British Telecom NET3 (**ETSI**).

B1/B2 Signalling Calling Address

The number assigned to this channel by the carrier. If only one address is supplied by the carrier, it should be assigned to channel B1, and channel B2 should be left empty.

B1/B2 Service Profile Identifier (SPID)

The Service Profile Identifier assigned to this channel by the carrier. Normally, this value contains the calling address surrounded by some digits. If only one address is supplied by the carrier, it should be assigned to channel B1, and channel B2 should be left empty.

Displaying ISDN Configuration Entry Status

The `isdns` command is used to view the operational status of existing ISDN configuration entries. You can select to view the status of all ISDN ports, or select to display the status of a single ISDN port.

Displaying Status of All ISDN Ports

To view status information of the ISDN channels on all ISDN ports, enter the following command:

```
isdns
```

A screen similar to the following displays:

Slot/Port	Type	Oper Status	Call Address	Call Setup Time
=====	=====	=====	=====	=====
5/2(B1)	BRI-U	ACTIVE	7145555555	00:00:00 01/70
5/2(B2)	BRI-U	IDLE	7145555555	00:00:00 01/70
5/4(B1)	BRI-U	ACTIVE	7145555555	00:00:00 01/70
5/4(B2)	BRI-U	ACTIVE	7145555555	00:00:00 01/70

The fields on this screen have the following meanings:

Slot/Port

Identifies the ISDN port and slot numbers and the “B” channel number (in parentheses).

Type

Identifies the type of ISDN port (BRI-U or BRI-S/T). See *Overview of ISDN* on page 32-1.

Operational Status

Identifies the operational status of this port. The possible entries in the table are **Active**, meaning the call is currently in progress, or **Idle**, meaning the interface is currently idle.

Call Address

Identifies the current or last phone number that was called on this ISDN channel.

Call Setup Time

Identifies the value of “sysUpTime” (the time of day recorded by the switch) and the date (in *dd/yy* format) when the last call was established on this channel.

Displaying Status of a Specific ISDN Slot

To view status information on all ISDN channels on a specific ISDN slot, for example, slot 4, enter the following command:

```
isdns 4
```

A screen similar to the following displays:

```
Status for ISDN D channel on slot: 4, Port: 1:  
LAPD OperStatus: Layer 1: Active, Layer 2 DataLink: Established.
```

```
The number of incoming calls . . . . . 0  
The number of incoming calls which were actually connected . . . 0  
The number of outgoing calls . . . . . 0  
The number of outgoing calls which were actually connected . . . 0
```

	Oper Status	Peer Address	Call Origin	Call SetupTime
	=====	=====	=====	=====
B1	Idle	7144509154	Incoming	0:00:00 01/70
B2	Conn	7144509156	Outgoing	0:00:00 01/70

```
Status for ISDN D channel on slot: 4, Port: 2:  
LAPD OperStatus: Layer 1: Active, Layer 2 DataLink: Established.
```

```
The number of incoming calls . . . . . 0  
The number of incoming calls which were actually connected . . . 0  
The number of outgoing calls . . . . . 0  
The number of outgoing calls which were actually connected . . . 0
```

	Oper Status	Peer Address	Call Origin	Call Setup Time
	=====	=====	=====	=====
B1	Idle	7144509154	Incoming	0:00:00 01/70
B2	Conn	7144509156	Outgoing	0:00:00 01/70

The fields on this screen have the following meanings:

The number of incoming calls

Indicates the number of incoming calls received on this interface.

The number of incoming calls which were actually connected

Indicates the number of incoming calls which were actually connected on this interface. The difference between the previous field and this one is the number of calls that were refused.

The number of outgoing calls

Indicates the number of outgoing calls made on this interface.

The number of outgoing calls which were actually connected

Indicates the number of outgoing calls which were actually connected on this interface. The difference between the value of the previous entry and this one is the number of calls that failed.

Oper Status

Indicates the current call control state for this interface. The possible entries are:

- Idle** means the B Channel is idle: no call or call attempt is in progress.
- Connecting** means a connection attempt (outgoing call) is being made.
- Connected** means an incoming call is currently in the process of validation.
- Active** means a call is currently active.

Peer Address

Indicates the ISDN address to which the current or last call is or was connected. In some cases, the format of this information cannot be predicted since it largely depends on the type of switch or PBX to which the device is connected. The switch software supports the display of IA5 ASCII digits and the pound key (0-9 and #), but no space characters.

Call Origin

Indicates whether this call was answered on this channel (denoted as “**Incoming**”) or was originated by this channel (denoted as “**Outgoing**”).

Call Setup Time

Indicates the value of “sysUpTime” (the time of day recorded by the switch) when the ISDN setup message for the current or last call was sent or received. If, since system start-up, there has been no call on this interface, this field will display all zeros.

Displaying Status of a Specific ISDN Port

To view status information of the ISDN channels on a specific ISDN port, for example port 4, slot 1, enter the following command:

isdns 4/1

A screen similar to the following displays:

```

Status for ISDN D channel on slot: 4, Port: 1:
LAPD OperStatus: Layer 1: Active, Layer 2 DataLink: Established.

The number of incoming calls . . . . . 0
The number of incoming calls which were actually connected . . . 0
The number of outgoing calls . . . . . 0
The number of outgoing calls which were actually connected . . . 0

      Oper   Peer       Call       Call
      Status Address   Origin    SetupTime
      =====
B1  Idle   7144509154  Incoming  0:00:00 01/70
B2  Conn   7144509156  Outgoing  0:00:00 01/70
    
```

The fields on this screen were described earlier in this section (see *Displaying Status of a Specific ISDN Slot* on page 32-8).

33 Managing T1 and E1 Ports

T1 and E1 ports are supported on a variety of switching modules. In the Omni Switch/Router, T1 and E1 are used as standard WAN access ports. The following switching modules contain T1 or E1 ports:

- WSX-FT1/E1

Ports on these modules share a common set of physical level attributes and a common set of software configuration commands. T1/E1 configuration options include frame format, line coding, and Facility Datalink Protocol. T1/E1 ports can store up to 24 hours of performance statistics for local and remote ports. These software commands do not configure time slots.

Hardware descriptions of the WSX-FT1/E1 can be found in Chapter 3, “Omni Switch/Router Switching Modules.”

This chapter is divided into two parts. The first part provides an overview of T1/E1 digital services. The second part describes the configuration of physical T1 and E1 ports; this second part starts with the section, *The T1/E1 Menu* on page 33-3.

T1 and E1 Overview

Carrier digital services were designed primarily to support digitized voice over long distances. Digital services are the primary method for carrying voice between two endpoints using two pairs of copper wire. Digital wide-area data networking uses the same digital services that were originally designed for digitized voice.

Analog to Digital Conversion

To improve quality and reliability, long-distance phone networks upgraded their backbones from analog Frequency Division Multiplexing (FDM) to digital Time Division Multiplexing (TDM). In TDM, analog data is converted to digital data using a CODEC device that employs a method called Pulse Code Modulation (PCM).

In Pulse Code Modulation, the CODEC samples the analog signal 8,000 times a second and converts each sample to an 8-bit digital value. These 8,000 8-bit samples yield a total digital data rate of 64,000 BPS for one voice service. This service is also known as Digital Service Zero (DS0), which is the basis for T1 and E1 connections.

These 8,000 8-bits in time are also known as a *time slot*. A *channel* is a time slot that can carry voice or data. Using Time Division Multiplexing, 24 channels (for T1) or 32 channels (for E1) are multiplexed to create a service called Digital Service 1 (DS1). The more common name for DS1 is *T1* or *E1*.

T1 Framing

A T1 frame consists of 24, 8-bit time slots and a 1-bit synchronization and control bit. Twelve (12) T1 frames can be grouped into a *SuperFrame (SF/D4)*, or 24 T1 frames can be grouped into an *Extended SuperFrame*. In each SuperFrame, the 6th and 12th frame may contain “robbed bit” (A, B) signalling, which means the least significant bit is robbed from each time slot in the 6th and 12th frame and used for signalling. In Extended SuperFrames, this robbed-bit signalling (A, B, C, D) occurs in the 6th, 12th, 18th, and 24th frames.

E1 Framing

The E1 frame consists of 32, 8-bit time slots (two of these time slots are used for synchronization and multiframe signalling) for 256 bits per frame at 2.048 megabits per second. Sixteen (16) E1 frames are grouped into a multiframe. An E1 multiframe can use Channel Associated Signalling (CAS) contained in time slot 16. Timeslot 16 in multiframe 0 is used for multiframe synchronization and control. Timeslot 16 of multiframes 1 through 15 are used to carry A, B, C, and D signaling bits.

The T1/E1 Menu

The commands for configuring and monitoring T1 and E1 ports are contained in the **te** submenu. This submenu displays as shown below and may be accessed (when in verbose mode) by entering **te** at a system prompt.

Command	T1/E1 Port Management Menu
tes	View status of a T1/E1 port configuration and statistics
temod	Modify a T1/E1 port configuration
tecls	Clear framer statistics of a T1/E1 port
telts	Display 24-hour period statistics of a local T1/E1 port
telcs	Display current 15-minute statistics of a local T1/E1 port
telis	Display 15-minute interval statistics of a local T1/E1 port
terts	Display 24-hour period statistics of a remote T1/E1 port
tercs	Display current 15-minute statistics of a remote T1/E1 port
teris	Display 15-minute interval statistics of a remote T1/E1 port
tebcfg	Configure BERT test
tebs	Display BERT statistics
tebcls	Clear BERT statistics
tecfg	Configure T1/E1 port type

The commands in this menu are described in the following sections. The first command, **tes**, displays configuration information on ports. This configuration information is configured through the **temod** command. The remaining commands, listed after the **telts** command provide a variety of interval statistics for local and remote T1 and E1 connections.

◆ Note ◆

The **tebcfg**, **tebs**, **tebcls**, and **tecfg** commands apply only to the OmniAccess 408 and 512. For the Omni Switch/Router, these are nonfunctioning commands.

Configuring a T1 Port

The **temod** command configures a T1 port at the physical level and is generic to all such ports regardless of the logical level service, such as circuit emulation, that controls them.

To configure a T1 port, enter the following command

```
temod <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the T1 port number on the board that you want to modify. For example, to modify port number 2 on the board in switch slot 5, enter

```
temod 5/2
```

A screen similar to the following displays:

```

T1 Port Configuration for slot 5, port 2

1) Circuit Identifier (30 chars max)           : Alcatel T1 Circuit
2) Frame Format { ESF (2), SF (3), unframed (8) } : ESF
3) Line Build Out { short(1), long(2) }       : short
   30) Line Length in meters (0-200)         : 30
4) Line Coding { B8ZS (2), AMI (5) }         : B8ZS
5) Facility Datalink { ANSI T1.403 (2), AT&T 54016 (4),
   T1.403-AT&T (6), none (8) }             : none
6) Facility Datalink Port Role { network (1), user (2) } : network
7) Transmit Clock Source { loopTiming (1),
   localTiming (2) }                       : localTiming
8) Loopback Mode { none (1), payload (2), line (3),
   inward (5) }                             : inward
9) Signalling { none (1), CAS (2), CCS (3) } : none
10) Trap Generation { enabled (1), disabled (2) } : disabled
11) Yellow Alarm Detection { enabled (1), disabled (2) } : enabled

Enter (option=value/save/cancel) :
```

1) Circuit Identifier

Enter a textual description of this T1 port, up to 30 characters. This text will be used in other screen displays to identify this T1 port.

2) Frame Format

Specify the frame format to be used on this port. The choices are Extended SuperFrame (**ESF**), SuperFrame or D4 (**SF**), or no special frame format (**unframed**). A T1 frame consists of 24 8-bit time slots and a 1-bit synchronization and control. Twelve (12) T1 frames can be grouped into a SuperFrame, and 24 T1 frames can be grouped into an Extended SuperFrame.

Normally, you should configure a T1 port as ESF (the default) since a T1 port configured as SuperFrame (SF) can produce false yellow alarms if a Layer 2 protocol like High-Level Data Link Control (HDLC) is being used.

If you must set the port as SF, you can disable Yellow Alarm detection with the **Yellow Alarm Detection** option, which is described on page 33-7.

If you choose the **unframed** format, then the framer will not look for Channel Associated Signalling (CAS). Data is treated as a data stream. When used in a circuit emulation application, this option must be chosen when configuring an “unstructured” circuit emulation service.

Important Note

You *cannot* use the unframed format on WAN modules.

3) Line Build Out

Indicate whether the T1 port supports short haul or long haul interfaces. Only T1 ports equipped with Line Interface Unit (LIU) support long haul. Long haul support is necessary if this T1 port is directly connected to a Central Office (CO) and the cable length is greater than 655 feet (200 meters). If this T1 port connects locally (i.e., it is not connected to an external CSV) using less than 655 feet (200 meters) of cable, short haul is adequate.

Note

All T1/E1 ports are equipped with a Line Interface Unit (LIU) chip.

An additional prompt displays for either the line length between this port and the T1 device (short haul configurations) or the attenuation of the cable attaching this port and the T1 device (long haul configurations). Each of these options is described below.

40) Line Length in meters

Displayed only when **short haul** is chosen as the **Line Build Out** option. Specify the distance, in meters, between this T1 port and the attached T1 device.

41) Attenuation

Displayed only when **long haul** is chosen as the **Line Build Out** option. Specify the attenuation of the line between this T1 port and the attached T1 device.

4) Line Coding

The type of physical encoding used on the connection. AMI (Alternate Mark Inversion) is more sensitive. B8ZS (Bipolar 8 Zero Substitution) should be used when possible. In most networks, B8ZS is recommended. In all cases, the Line Coding you select must match that provided by your service provider.

5) Facility Datalink

Facility Datalink (FDL) gathers performance statistics every second and stores them in the 24-hour local statistical database. It also sends local performance statistics to the remote T1 port depending on the type of FDL chosen and the “role” of the FDL (specified in the next field). In order to obtain far-end, or remote, performance statistics (viewed through the **terts**, **tercs**, and **teris** commands), you must enable an FDL protocol.

Note

Facility Datalink requires a T1 port and the frame type must be Extended SuperFrame.

You have the following choices:

- | | |
|------------------------|--|
| ANSI T1.403 | The FDL exchange recommended by ANSI. The FDL method sends Performance Report Messages (PRMs) to the far-end port every second, processes received PRMs, and stores them in a 24-hour far-end statistical database. |
| AT&T 54106 | The operation of this FDL protocol depends on the Facility DataLink Port Role setting (configured in the next field). The FDL protocol will either be active (network) or passive (user) in its sending of PRMs. |
| T1.403-AT&T | In this combination selection, the port supports both the ANSI (ANSI T1.403) and AT&T Extended Superframe (AT&T 54106) protocols at the same time. The port processes ANSI messages as described for the ANSI T1.403 option and responds to AT&T request messages. |
| none | The port does not use Facility Datalink. |

6) Facility Datalink Port Role

Indicates the role of this port in relation to the remote port. This setting only affects configurations where the Facility Datalink field is set to **AT&T 54016**. When set to **network**, far-end historical statistics are updated by periodically sending 24-hour and 1-hour performance statistics requests to the far-end port. When set to **user**, the FDL passively waits for messages from the far-end port.

7) Transmit Clock Source

The source of the transmit clock. Loop timing means the receive clock (recovered from receive data) is used as the transmit clock. Local timing indicates the local clock source (generated from PLLs) is used as the transmit clock.

The transmit clock source is related to the clocking mode used in circuit emulation services. In *synchronous* clocking mode, both sides of the T1 connection will use a local clock source. However, in *SRTS* and *adaptive* clocking, the T1 port receives the clock on one end (loop timing) and regenerates the clock locally (local timing) on the other end. In such a case, the T1 port receiving the clock from the network should be configured as **loop timing** and the other end of the link should be configured as **local timing**.

8) Loopback Mode

The loopback configuration for this port. Loopback configurations describe the relation between the device attached to a T1 port and the framing functionality within the T1 port. Framing functionality assembles T1 frames into SuperFrames and Extended SuperFrames, depending on how the port is configured. Possible values are as follows:

- | | |
|----------------|--|
| none | The port is not in a loopback state. This is the typical live network state for a T1 port. |
| payload | The received signal at this T1 port is looped out of the port after passing through the port's framing functionality. This state should only be used for debugging purposes. |
| line | The received signal at this T1 port does not go through the port's framing functionality, and is looped straight back out the port. This state should only be used for debugging purposes. |
| inward | The transmitted signal from the inward side of this port is looped back internally. The signal passes through the T1 framing functionality before looping back. This state should only be used for debugging purposes. |

9) Signaling

The type of signaling used on this port. Only the **none** and **CAS** (Channel Associated Signaling) options are applicable to a circuit emulation service port. The **CCS** (Common Signal Channeling) option is used with external ISDN Primary Rate ports. If you select the **CAS** option, then you are enabling robbed-bit signalling.

Robbed-bit signalling can be used with SuperFrames or Extended SuperFrames. In each SuperFrame, the 6th and 12th frame may contain "robbed bit" (A, B) signalling, which means the least significant bit is robbed from each time slot in the 6th and 12th frame and used for signalling. In Extended SuperFrames, this robbed-bit signalling (A, B, C, D) occurs in the 6th, 12th, 18th, and 24th frames.

10) Trap Generation

Enables all of the SNMP-based traps related to T1 and E1 ports.

11) Yellow Alarm Detection

Specify the yellow alarm detection state for this port. A T1 port configured as SuperFrame (SF) can produce false yellow alarms if a Layer 2 protocol like High-Level Data Link Control (HDLC) is being used. Therefore, you can disable yellow alarm detection with this option. (A T1 port set to Extended SuperFrame (ESF) will not produce false yellow alarms.)

Configuring an E1 Port

The **temod** command configures an E1 port at the physical level and is generic to all such ports regardless of the logical level service, such as circuit emulation, that controls them. You configure the circuit emulation service that controls this port through the **cemodify** command.

To configure an E1 port, enter the following command

```
temod <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the T1 port number on the board that you want to modify. For example, to modify port number 2 on the board in switch slot 4, you would enter

```
temod 4/2
```

A screen similar to the following displays:

E1 Port Configuration for slot 4, port 2

1) Circuit Identifier { 30 chars max }	: Alcatel E1 Circuit
2) Frame Format { E1 (4), E1-CRC (5), E1-MF (6), E1-CRC-MF (7), unframed (9) }	: E1
3) Not FAS { enabled (1), disabled (2) }	: enabled
4) Line Build Out { short(1), long(2) }	: short
40) Cable Type { 75 Ohm (1), 120 Ohm (2) }	: 75 Ohm
5) Line Coding { HDB3 (3), AMI (5) }	: HDB3
6) Transmit Clock Source { loopTiming (1), localTiming (2) }	: localTiming
7) Loopback Mode { none (1), payload (2), line (3), inward (5) }	: none
8) Signalling { none (1), CAS (2), CCS (3) }	: none
9) Trap Generation { enabled (1), disabled (2) }	: disabled

Enter (option=value/save/cancel) :

1) Circuit Identifier

Enter a textual description of this E1 port, up to 30 characters. This text will be used in other screen displays to identify this E1 port.

2) Frame Format

Specify the E1 frame format to be used on this port. The choices are as follows:

- E1** Standard E1 frame format using the framing bits in time slot 0 for framing.
- E1-CRC** E1 frame using framing bits in both time slot 0 and CRC-4 multiframe for framing.
- E1-MF** E1 frame using framing bits in both time slot 0 and time slot 16 multiframe for framing.
- E1-CRC-MF** E1 frame using framing bits in time slot 0, time slot 16 multiframe, and CRC-4 multiframe for framing.
- unframed** The framing software will not look for framing bits to determine the start of a frame or multiframe. Data is treated as a data stream. When used in a circuit emulation application, this option should be chosen when configuring an “unstructured” circuit emulation service.

Important Note

You *cannot* use the unframed format on WAN modules.

3) Not FAS

Indicates whether you want to add an extra level of frame checking. E1 frames in time slot 0 are composed of alternating bits of FAS (Frame containing Frame Alignment Signal) and NFAS (Frame not containing Frame Alignment Signal). The **Not FAS** option tells the framer to check framing on FAS and NFAS bits. Normally, the framer checks only FAS bits, which contain the frame alignment signal pattern. If you enable **Not FAS**, then framing software will additionally also check NFAS bits, which include remote alarm indication information.

4) Line Build Out

The E1 port supports short haul or long haul interfaces. Only E1 ports equipped with a Line Interface Unit (LIU) chip support long haul. Long haul support is necessary if this E1 port is directly connected to a Central Office (i.e., not connected via an external CSU) and the cable length is greater than 655 feet (200 meters). If this E1 port connects locally using less than 665 feet (200 meters) of cable, then short haul is adequate.

Note

All T1/E1 ports are equipped with a Line Interface Unit (LIU) chip.

An additional prompt displays requesting the resistance type used for this port connection.

40) Cable Type

Indicate the cable resistance type used on the short or long haul interface. The cable resistance type can be 75 ohm or 120 ohm. The resistance is a set via a jumper on the E1 board; it is not configurable through software.

5) Line Coding

The type of physical encoding used on the connection. AMI (Alternate Mark Inversion) is more sensitive. HDB3 (High Density Bipolar 3) should be used when possible.

6) Transmit Clock Source

The source of the transmit clock. Loop timing means the receive clock (recovered from receive data) is used as the transmit clock. Local timing indicates the local clock source (generated from PLLs) is used as the transmit clock.

The transmit clock source is related to the clocking mode used in circuit emulation services. In synchronous clock mode, both sides of the E1 connection will use a local clock source. However, in SRTS and adaptive clocking, the E1 port receives the clock on one end (loop timing) and regenerates the clock locally (local timing) on the other end. In such a case, the E1 port receiving the clock from the network should be configured to **loop timing** and the other end of the link should be configured to **local timing**.

7) Loopback Mode

The loopback configuration for this port. Loopback configurations describe the relation between the device attached to an E1 port and the framing functionality within the E1 port. Framing functionality assembles E1 frames into multiframes, depending on how the port is configured. Possible values are as follows:

- | | |
|----------------|--|
| none | The port is not in a loopback state. This is the typical live network state for an E1 port. |
| payload | The received signal at this E1 port is looped out of the port after passing through the port's framing functionality. This state should only be used for debugging purposes. |
| line | The received signal at this E1 port does not go through the port's framing functionality, and is looped straight back out the port. This state should only be used for debugging purposes. |
| inward | The transmitted signal from the inward side of this port is looped back internally. The signal passes through the E1 framing functionality before looping back. This state should only be used for debugging purposes. |

8) Signalling

The type of signaling used on this port. Only the **none** and **CAS** (Channel Associated Signaling) options are applicable to a circuit emulation service port. The **CCS** (Common Signal Channeling) option is used with external ISDN ports. If you select the CAS option, then you are enabling Channel Associated Signaling, which is used with E1 multiframes. In Channel Associated Signaling, timeslot 16 in frame 0 of the multiframe is used for multiframe synchronization and control. Timeslot 16 of frames 1 through 15 are used to carry A, B, C, and D signaling bits.

9) Trap Generation

Enables all of the SNMP-based traps related to circuit emulation service ports.

Viewing T1/E1 Configuration and Alarm Information

You can view all current parameters and alarms for a T1 or E1 port using the **tes** command. These parameters will be either the default parameters or parameters you modified through the **temod** command or network management software.

You have a choice of viewing parameters at the chassis or port level. You receive different displays depending upon which level you choose. The sections below describe all ways to use the **tes** command.

Viewing Information for all T1/E1 Ports in the Switch

To view port parameters for all T1/E1 ports in a chassis, enter the following command

```
tes
```

A screen similar to following displays:

T1/E1 Chassis Status		
Slot/Port	Type	Active Alarms
4/2	E1	NoAlarm
4/3	E1	NoAlarm
5/2	T1	NoAlarm, Loopback
5/3	T1	NoAlarm, Loopback

Slot/Port. The T1 or E1 slot and port for which information is supplied. The slot is listed first, followed by a slash (/), followed by the port number.

Type. The port type. The port will either be a T1 or E1 port.

Active Alarms. Alarms that have occurred on this port. Possible alarms for each port are:

NoAlarm	The port is free of any alarms.
RcvYellow	This port is receiving a yellow alarm from the far-end port. A yellow alarm occurs in SuperFrames when bit 6 of all channels has been zero for at least 425 milliseconds. The yellow alarm will not occur if a Loss of Signal alarm has already occurred. In Extended SuperFrames, an alarm occurs if the yellow alarm pattern is found.
XmtYellow	The port is transmitting a yellow alarm <i>to</i> the far-end port. See the above definition of RcvYellow for a description of a yellow alarm.
RcvAIS	This port is receiving Alarm Indication Signal (AIS) from the far-end port. An AIS occurs when an unframed signal with a high density of 1s (99.9% density) is received for more than 1.5 seconds.
XmtAIS	This port is transmitting Alarm Indication Signal (AIS) <i>to</i> the far-end port. An AIS occurs when an unframed signal with a high density of 1s (99.9% density) is received for more than 1.5 seconds.
RedAlarm	The port is in red alarm state. A red alarm occurs when a T1 port has been in Out-of-Frame (OOF) condition for 2.55 seconds. The red alarm condition will be removed if the OOF condition has been absent for at least 16.6 seconds.

Viewing T1/E1 Configuration and Alarm Information

LossOfSignal	The port has experienced a Loss of Signal (LOS), or Loss of Carrier. An LOS event occurs after 175 contiguous pulse positions with no pulses (10 absent pulses on E1 ports). An LOS failure is cleared after the switch observes a single pulse.
RcvLOMF	This port is receiving loss of multiframe (LOMF) alarms from the far-end port. When a far-end E1 port detects an out-of-multiframe condition, it transmits a frame with the alarm indication bit set (in time slot 16) back to the local E1 port. This error is similar to a yellow alarm on T1 ports.
LocalUA	This port is not available possibly because a cable is not attached.
Loopback	The port is currently in loopback mode. Loopback mode can be configured through the temod command or dynamically activated through Facility Data Link (ANSI T1.403 and AT&T 54106) or through loopback control codes on a T1 port.

Viewing Information for T1/E1 Ports on One Module

To view port parameters, enter the following command

```
tes <slot>
```

where **<slot>** is the slot number where the on which you want to view information resides. For example, to view configuration parameters for the board in slot 5, enter

```
tes 5
```

A screen similar to following displays:

T1/E1 Port Status for slot 5		
Port	Type	Active Alarms
2	T1	NoAlarm, Loopback
3	T1	NoAlarm, Loopback

Explanations of the columns in this table are described in the section, *Viewing Information for all T1/E1 Ports in the Switch* on page 33-11.

Viewing Information For a T1 Port

To view T1 port parameters, enter the following command

```
tes <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the T1 port number on the board on which you want to view information. For example, to view information for Port 2 on the board in slot 5, enter

```
tes 5/2
```

A screen similar to following displays for a T1 port:

T1/E1 Port Status for slot 5, port 2

Circuit Identifier	: Alcatel T1 Circuit		
Frame Format	: ESF	Line Build Out	: 30 (SH)
Facility Datalink	: none	FDL Port Role	: network
Line Coding	: B8ZS	Signalling	: none
Transmit Clock Source	: localTiming	Trap Generation	: disabled
Status Change Time	: 0 days, 00:07:24.69		
Loopback Status	: LocalInwardLoop		
Line Status	: NoAlarm, Loopback		

Framer Statistics

Loss of Signal Events	: 0
Line Code Violation Events	: 431986
Out of Frame Events	: 0
Red Alarm Events	: 1
Squelch Alarm Events	: 0
Frame Bit Error Events	: 2
Alarm Indication Signal Events	: 0
Yellow Alarm Events	: 1
ESF CRC-6 Error Events	: 3

Circuit Identifier, Frame Format, Line Build Out, Facility Datalink, FDL Port Role, Line Coding, Signalling, Transmit Clock Source, Trap Generation. These parameters are described in the section, *Configuring a T1 Port* on page 33-4. Please refer to that section for descriptions.

Status Change Time. The system time when the last change in Line Status (i.e., alarm) parameter occurred.

Loopback Status. The type of loopback mode configured for this port through the **temod** command or activated remotely through FDL. Loopback modes are described in *Configuring a T1 Port* on page 33-4.

Line Status. A list of any alarms that have occurred on his port. The possible items in the list are the same as those for **Active Alarms** described in *Viewing Information for all T1/E1 Ports in the Switch* on page 33-11.

Loss of Signal Events. The total number of Loss of Signal (LOS) events that have been detected on this port. An LOS event occurs after 175 contiguous pulse positions with no pulses (10 absent pulses on E1 ports). An LOS failure is cleared after the switch observes a single pulse.

Line Code Violation Events. The occurrence of either a bipolar violation or an excessive zeros error. A bipolar violation is the occurrence of a pulse of the same polarity as the previous pulse. In B8ZS coded signals, a bipolar violation is a pulse of the same polarity as the previous without being part of the zero substitution code. An excessive zeros error is the occurrence of more than 15 contiguous zeros in an AMI-coded signal; in a B8ZS-coded signal, it is the occurrence of seven (7) or more contiguous zeros.

Out of Frame Events. The total number of out of frame events that have been detected on this port. An out of frame event occurs when two or more framing errors occur within a 3 microsecond period for Extended SuperFrame signals, or when two or more errors occur out of five or fewer consecutive framing bits. The signal will be back in frame when there have been fewer than two frame bit errors within a 3 microsecond period for Extended SuperFrame signals.

Red Alarm Events. The number of times this port has been in a red alarm state, which occurs when a T1 port has been in Out-of-Frame (OOF) condition for 2.55 seconds. The red alarm condition will be removed if the OOF condition has been absent for at least 16.6 seconds.

Squelch Alarm Events. The number of squelch alarm events that have been detected on this port. A squelch alarm occurs when the line signal level of the input pulse is below a threshold level. The threshold level on a T1 line is 0.5V.

Frame Bit Error Events. The number of framing bit error events that have been detected on this port. A frame bit error occurs when an error bit is detected during the framing process.

Alarm Indication Signal Events. The number of Alarm Indication Signal (AIS) events that have been detected on this port. An AIS occurs when an unframed signal with a high density of 1s (99.9% density) is received for more than 1.5 seconds.

Yellow Alarm Events. The total number of yellow alarm events that have occurred on this T1 port. A yellow alarm occurs in SuperFrames when bit 6 of all channels has been zero for at least 335 microseconds. The yellow alarm will not occur if a Loss of Signal alarm has already occurred. In Extended Superframes, an alarm occurs if the yellow alarm pattern is found.

Note

A T1 port that has been configured as a SuperFrame (SF) port can produce false yellow alarms. You can disable yellow alarm detection on a T1 port with the **temod** command, which is described in *Configuring a T1 Port* on page 33-4.

ESF CRC-6 Error Events. The number of times a CRC-6 error has been found in an Extended SuperFrame.

Viewing Information For an E1 Port

To view E1 port parameters, enter the following command

```
tes <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the E1 port number on the board for which you want to view information. For example, to view information for Port 2 on the board in slot 4, enter

```
tes 4/2
```

A screen similar to following displays for an E1 port:

T1/E1 Port Status for slot 4, port 2

Circuit Identifier	: Alcatel E1 Circuit		
Frame Format	: E1	Line Build Out	: 120 Ohm (SH)
Line Coding	: HDB3	Signalling	: none
Transmit Clock Source	: localTiming	Trap Generation	: disabled
Status Change Time	: 0 days, 00:06:34.69		
Loopback Status	: NoLoop		
Line Status	: NoAlarm		

Framer Statistics

Loss of Signal Events	:	1
Line Code Violation Events	:	9
Out of Frame Events	:	2
Red Alarm Events	:	1
Squelch Alarm Events	:	1
Frame Bit Error Events	:	9
Alarm Indication Signal Events	:	3
Out of Sub-multiframe Events	:	0
Out of TS16 Multiframe Events	:	0
Far End Frame Alarm Events	:	2
Far End Multiframe Alarm Events	:	0
Far End Block Error Events	:	0
CRC-4 Error Events	:	0

Circuit Identifier, Frame Format, Line Build Out, Line Coding, Signaling, Transmit Clock Source, Trap Generation. These parameters are described in the section, *Configuring an E1 Port* on page 33-8. Please refer to that section for descriptions.

Status Change Time. The system time when the last change in Line Status (i.e., alarm) parameter occurred.

Loopback Status. The type of loopback mode configured for this port through the **temod** command. Loopback modes are described in *Configuring an E1 Port* on page 33-8.

Line Status. A list of any alarms that have occurred on his port. The possible items in the list are the same as those for **Active Alarms** described in *Viewing Information for all T1/E1 Ports in the Switch* on page 33-11.

Loss of Signal Events. The total number of Loss of Signal (LOS) events that have been detected on this port. An LOS event occurs after the port detects more than 10 consecutive zeros.

Line Code Violation Events. The occurrence of either a bipolar violation or excessive zeros error. A bipolar violation is the occurrence of a pulse of the same polarity as the previous pulse. In HDB3 coded signals, a bipolar violation is a pulse of the same polarity as the previous without being part of the zero substitution code. An excessive zeros error is the occurrence of more than 15 contiguous zeros in an AMI-coded signal; in an HDB3-coded signal, it is the occurrence of seven (7) or more contiguous zeros.

Out of Frame Events. The total number of out of frames events that have been detected on this port. An out of frame event occurs when three consecutive frame alignment signals have been received with an error. The signal will be back in frame when frame alignment signalling is normal for three consecutive frames.

Red Alarm Events. The number of times this port has been in a Red alarm state. A red alarm occurs when a T1 port has been in Out-of-Frame (OOF) condition for 2.55 seconds. The red alarm condition will be removed if the OOF condition has been absent for at least 16.6 seconds.

Squelch Alarm Events. The number of squelch alarm events that have been detected on this port. A squelch alarm occurs when the line signal level of the input pulse is below a threshold level.

Frame Bit Error Events. The number of framing bit error events that have been detected on this port. A frame bit error occurs when an error bit is detected during the framing process.

Alarm Indication Signal Events. The number of Alarm Indication Signal (AIS) events that have been detected on this port. An AIS occurs when an unframed signal with a high density of 1s (99.9% density) is received for more than 1.5 seconds.

Out of Sub-multiframe Events. The number of sub-multiframe events that have been detected on this E1 port. This error occurs when four (4) consecutive CRC-4 multiframe alignment signals have been received in error or when a frame alignment error has been lost.

Out of TS16 Multiframe Events. The number of TS16 multiframe events that have been detected on this E1 port. This error occurs when two (2) consecutive TS16 multiframe alignment signals have been received in error, or all bits in time slot 16 are logic 0 for one TS16 multiframe, or frame alignment has been lost.

Far End Frame Alarm Events. The number of times the remote end has detected an out-of-frame condition. When a far end E1 port detects an out-of-frame condition, it transmits a frame with the alarm indication bit set (in time slot 0) back to the local E1 port. This error is similar to a yellow alarm on T1 ports.

Far End Multiframe Alarm Events. The number of times the remote end has detected an out-of-multiframe condition. When a far-end E1 port detects an out-of-multiframe condition, it transmits a frame with the alarm indication bit set (in time slot 16) back to the local E1 port. This error is similar to a yellow alarm on T1 ports.

Far End Block Error Events. The number of times the remote end has received a frame with a bad CRC-4. When the far end E1 port detects a CRC-4 error in the incoming frame, it transmits the frame with the E bit cleared.

CRC-4 Error Events. The number times a frame has been received with a bad CRC-4.

Viewing T1/E1 Local Statistics

There are a number of commands available for viewing local T1 and E1 statistics. These commands provide statistics for the past 24 hours, the current 15-minute interval, or the past 96 15-minute intervals. The following sections describe these commands.

Viewing Total Local Statistics

You can view statistics occurring during the past 24 hours on a single port by entering the following command

```
telts <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the port number on the board for which you want to view statistics. For example, to view 24-hour statistics for Port 2 on the board in slot 5, enter

```
telts 5/2
```

A screen similar to the following displays:

```

                                Local 24-hour Period Statistics for port 2 on slot 5

Circuit Identifier      : Alcatel T1 Circuit
Valid Intervals        : 1 of 96      Elapsed Time           : 421 of 900

  ES   SES  BES   UAS  SEFS   LES   CSS   PCV   LCV
  ----  ---  ---  ----  ----  ----  ---  ----  ----
    3    1    1    0    1   313    0    2   313

```

Circuit Identifier. The textual description of this T1 or E1 port as configured through the **temod** command.

Valid Intervals. Indicates the number of 15-minute intervals for which valid statistics were gathered during the previous 24 hours. Statistics may be gathered for up to 96 15-minute intervals during a 24 hour period.

Elapsed Time. The number of seconds that have elapsed during this 15-minute interval of gathering statistics. This time will be reset to zero when a 15-minute session of statistics gathering is complete (and stored) and the next 15-minute interval begins.

ES. Errored Seconds. For T1-ESF and E1-CRC conditions, this is a second with one or more Path Code Violations, one or more out-of-frame defects, one or more controlled slip errors, or an AIS error.

SES. Severely Errored Seconds. For T1-ESF frames, this is a second with 320 or more Path Code Violation errors, one or more out-of-frame defects, or an AIS error. For E1-CRC conditions, this is a second with 832 or more Path Code Violation errors, or one or more out-of-frame defects. For E1-noCRC signals, this is a second with 2048 or more Line Code Violation errors. For D4/(SF) frames, this is a second with framing errors, an out-of-frame error, or a second with 1544 or more line code violation errors.

BES. Bursty Errored Seconds. The number of seconds with fewer than 320 but more than one (1) Path Code Violation error (see below for definition), no Severely Errored Frame errors, and no AIS errors.

UAS. Unavailable Seconds. The number of seconds this port was unavailable for transmitting or receiving data. In general, a port is unavailable after 10 consecutive Severely Errored Seconds or after a failure on the interface occurs.

Viewing T1/E1 Local Statistics

SEFS. Severe Errored Framing Second. A second with one or more out-of-frame errors or an AIS error.

LES. Line Errored Seconds. The number of seconds during which one or more Line Code Violation errors have occurred (see also the definition of Line Code Violation below).

CSS. Controlled Slip Seconds. A one-second interval with one or more controlled slip errors. Controlled slip errors are the replication or deletion of the payload bits on a frame. Such an error may occur when there is a difference between the timing of a synchronous receiving terminal and the received signal.

PCV. Path Code Violations. A frame synchronization bit error in EF/D4 and E1-noCRC frames, or a CRC or frame synchronization error in the T1-ESF (Extended Super Frame) and E1-CRC frames.

LCV. Line Code Violations. The occurrence of either a bipolar violation or excessive zeros error. A bipolar violation is the occurrence of a pulse of the same polarity as the previous pulse. In B8ZS and HDB3 coded signals, a bipolar violation is a pulse of the same polarity as the previous without being part of the zero substitution code. An excessive zeros error is the occurrence of more than 15 contiguous zeros in an AMI-coded signal; in a B8ZS-coded signal, it is the occurrence of seven (7) or more contiguous zeros.

Viewing Current Local Statistics

You can view statistics for the current 15-minute interval on a single port by entering the following command

```
telcs <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the port number on the board on which you want to view statistics. For example, to view 15-minute interval statistics for Port 2 on the board in slot 5, enter

```
telcs 5/2
```

A screen similar to the following displays:

```
Local Current 15-minute Measurement for port 2 on slot 5
Circuit Identifier      : Alcatel T1 Circuit
Valid Intervals        : 1 of 96      Elapsed Time      : 431 of 900
=====
  ES   SES  BES  UAS  SEFS  LES  CSS  PCV  LCV
-----
   0    0   0   0   0    0   0   0   0
```

Definitions of the fields and columns in this display are the same as those used for the **telts** command. See *Viewing Total Local Statistics* on page 33-17 for an explanation of these statistics.

Viewing Local Historical Statistics

The **telis** command allows you to display historical statistics for the past 96 15-minute intervals. Enter the following command

```
telis <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the port number on the board on which you want to view statistics. For example, to view historical 15-minute interval statistics for Port 2 on the board in slot 5, enter

```
telis 5/2
```

A screen similar to the following displays:

Local 15-minute Interval Statistics for port 2 on slot 5

Circuit Identifier	: Alcatel T1 Circuit								
Valid Intervals	: 5 of 96			Elapsed Time			: 440 of 900		
Intv#	ES	SES	BES	UAS	SEFS	LES	CSS	PCV	LCV
1	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0

Definitions of the fields and columns in this display are the same as those used for the **telts** command. See *Viewing Total Local Statistics* on page 33-17 for an explanation of these statistics.

Viewing T1 Remote Statistics

To receive and monitor remote statistics on T1 ports you must enable the Facility Datalink (FDL) protocol through the **temod** command. These statistics will not be available if you do not enable FDL.

Note

Because there is no FDL standard for E1 configurations, remote statistics are not supported on E1 ports.

Viewing Total Remote Statistics

You can view remote statistics occurring during the past 24 hours on a single port by entering the following command

```
ters <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the port number on the board on which you want to view statistics. For example, to view remote 24-hour statistics for Port 2 on the board in slot 5, enter

```
ters 5/2
```

A screen similar to the following displays:

Remote 24-hour Period Statistics for port 2 on slot 5

```
Circuit Identifier      : Alcatel T1 Circuit
Valid Intervals       : 1 of 96      Elapsed Time           : 1 of 900

  ES  SES  BES  UAS  DM  SEFS  LES  CSS  PCV  LOFC
  ----  ---  ---  ---  ---  ---  ---  ---  ---  ---
    0    0    0    0    0    0    0    0    0    0
```

Most of the definitions of the fields and columns in this display are the same as those used for the **ters** command. See *Viewing Total Local Statistics* on page 33-17 for an explanation of these statistics. The remaining statistics are described below.

LOFC. Loss of Frame Count. A loss of frame count is the accumulation of the number of times a “Loss of Frame” is declared.

Viewing Current Remote Statistics

You can view remote statistics for the current 15-minute interval on a single port by entering the following command

```
tercs <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the port number on the board on which you want to view statistics. For example, to view remote 15-minute interval statistics for Port 2 on the board in slot 5, enter

```
tercs 5/2
```

A screen similar to the following displays:

```

Remote Current 15-minute Measurement for port 2 on slot 5

Circuit Identifier      : Alcatel T1 Circuit
Valid Intervals        : 1 of 96      Elapsed Time           : 1 of 900

  ES  SES  BES  UAS  DM  SEFS  LES  CSS  PCV  LOFC
  ----  ---  ---  ---  ---  ---  ---  ---  ---  ---
    0    0    0    0    0    0    0    0    0    0

```

Definitions of the fields and columns in this display are the same as those used for the **telts** command. See *Viewing Total Local Statistics* on page 33-17 for an explanation of these statistics.

Viewing Remote Historical Statistics

The **teris** command allows you to display remote historical statistics for the past 96 15-minute intervals. Enter the following command

```
teris <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the port number on the board on which you want to view statistics. For example, to view remote historical 15-minute interval statistics for Port 2 on the board in slot 5, enter

```
teris 5/2
```

A screen similar to the following displays:

```

Remote 15-minute Interval Statistics for port 2 on slot 5

Circuit Identifier      : Alcatel T1 Circuit
Valid Intervals        : 5 of 96      Elapsed Time           : 25 of 900

Intv#   ES  SES  BES  UAS  DM  SEFS  LES  CSS  PCV  LOFC
  ----  ---  ---  ---  ---  ---  ---  ---  ---  ---
    1    0    0    0    0    0    0    0    0    0    0
    2    0    0    0    0    0    0    0    0    0    0
    3    0    0    0    0    0    0    0    0    0    0
    4    0    0    0    0    0    0    0    0    0    0
    5    0    0    0    0    0    0    0    0    0    0

```

Definitions of the fields and columns in this display are the same as those used for the **telts** command. See *Viewing Total Local Statistics* on page 33-17 for an explanation of these statistics.

Clearing the Framer Statistics for a T1/E1 Port

The **tecls** command enables you to clear the accumulated physical-layer (Framer) statistics for a T1 or E1 port. To clear statistics, enter

```
tecls <slot>/<port>
```

where **<slot>** is the slot number where the board is located, and **<port>** is the port number on the board on which you want to clear statistics. For example, to statistics for Port 2 on the board in slot 5, enter

```
tecls 5/2
```

Once the statistics have been cleared, the following message will be displayed:

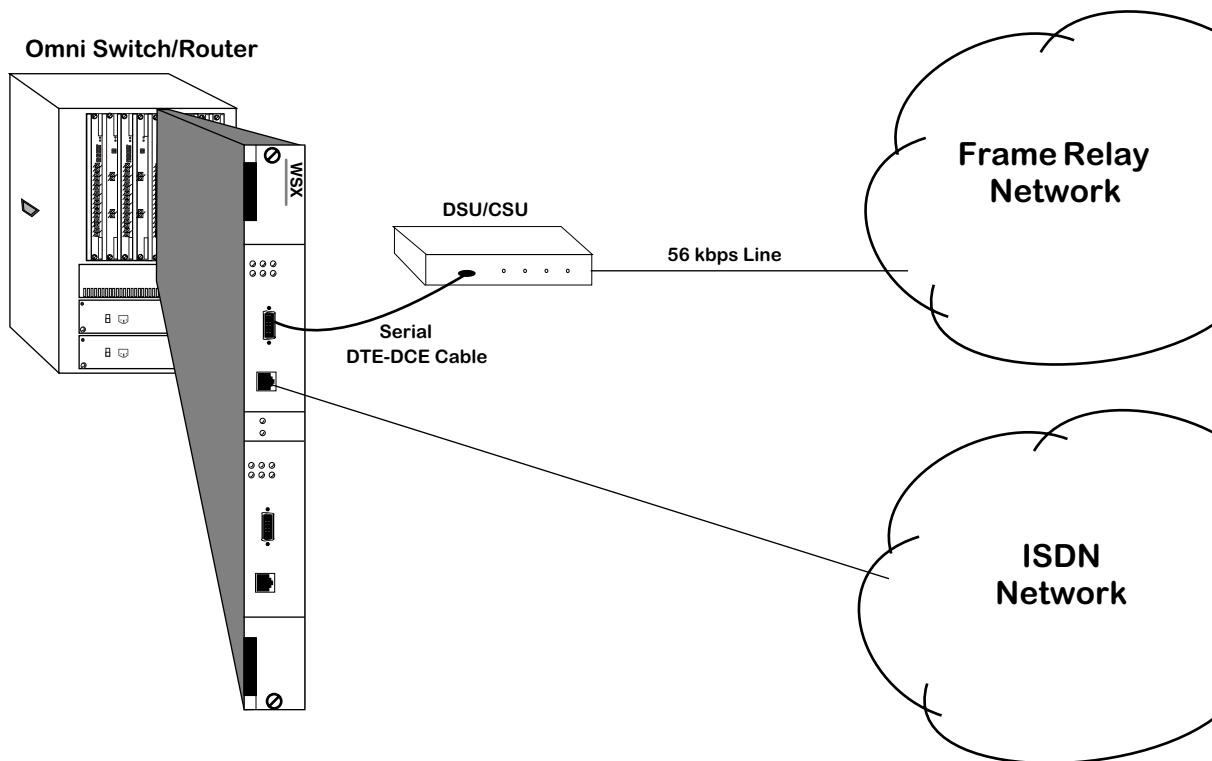
```
Statistics of port 5/2 have been cleared.
```

34 Backup Services

Introduction

Backup Services are intended to be an integral part of a well-designed Wide Area Network (WAN). The purpose of a backup service is to provide an alternate route for data to take in the event of failure of the Primary port or Virtual Circuit. Initially, the primary entity may be either a physical port (any physical port type in the system), or a frame relay Private Virtual Circuit (PVC). The backup is via an ISDN BRI running Point-to-Point Protocol (PPP).

Backup services are configured by specifying information on the primary entity, the backup entity, and timers that control under what conditions the system will switch to backup mode. Both the primary and backup entities must be configured prior to accessing this menu. This menu also does no cross-checking to ensure that the primary being backed up is backed up by an “appropriate” backup entity. This is the responsibility of the user.



Frame Relay to ISDN Backup

Backup Services Commands

Backup services provides commands to view and configure your backup services. All commands start with “bs” for “Backup Service” followed by the function desired. All backup commands may be typed in full, or a three character abbreviation may be used (e.g. **bsadd** or **bsa** may be used to create a backup service).

Accessing the Backup Services Menu

The Backup Services menu is a submenu to the Interface menu. To access the Interface menu, enter **inter**, followed by **<return>**, as shown below.

```
/ % inter <return>
```

To display a command summary for the Interface menu, enter **?**, followed by **<return>**:

```
/inter % ?
```

A screen similar to that shown below will display:

Command	Networking Menu
slipc	Configure SLIP (Serial Line IP) on a TTY Port
atm	Enter the atm Management submenu
eth100	Enter the 100BaseT submenu
10/100	Enter the 10/100BaseT submenu
wan	Enter the Wide Area Networking submenu
backup	Enter Backup networking command submenu

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

To enter the Backup menu, enter **backup**, followed by **<return>**, as shown below:

```
/Interface % backup <return>
```

To display a command summary for the Backup Services menu, enter **?**, followed by **<return>**.

```
/Interface/backup % ? <return>
```

A screen similar to that shown below will display:

Command	Networking Menu
bsadd	Add a Backup Service
bsmodify	Modify a Backup Service
bsview	View Backup Service(s)
bsdelete	Delete a Backup service
bsstatus	Display Backup service status
bsclear	Clear Backup service status

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

Adding a Backup Service

With the **bsadd** command, you can:

- Add a backup for a physical port
- Back up a frame relay PVC

Adding a backup for a Physical Port

To add a backup service for a physical port:

1. Enter the **bsadd** command with no parameters, followed by **<return>**.

```
/Interface/backup % bsa <return>
```

A screen similar to that shown below will display:

```

Adding Backup Service Index          :1
1) Description                        : Backup 1
2) Admin Status { (E)nabled, (D)isabled } : Enabled
3) Primary Type { Physical Port (1),
      Frame Relay PVC DLCI (2) }      : Physical Port
   30) Slot                            :
   31) Port                             :
4) Backup Type { PPP Peer (1) }       : PPP Peer
   40) Peer ID                          :
5) Startup Timer Value { Time in Seconds after
      System Startup to wait
      for Primary to come up
      before activating
      Backup }                          : 300
6) Activate Timer Value { Time in Seconds after
      Primary Failure to
      activate Backup }                  : 30
7) Restore Timer Value { Time in Seconds after
      Primary restoral to
      disable Backup }                   : 30
(save/quit/cancel)                    :

```

2. When you first enter the command, the next unique index is assigned automatically, a default description is created (**Backup** followed by the new index number), and defaults for primary type, backup type, and all backup timers are created (as shown above).
3. To back up a physical port, enter the numbers for the slot and port to be backed up and the PPP peer index (which defines ISDN call and PPP parameters). Optionally, you can modify the timer values (fields 5-7). Below is an example of backing up the port on slot 2, port 1 with PPP peer index 5.

```
      : 30=2
      : 31=1
      : 40=5
      : ?
1) Description: Backup 1
2) Admin Status { (E)nabled, (D)isabled }           : Enabled
3) Primary Type { Physical Port (1),
  Frame Relay PVC DLCI (2) }                       : Physical Port
  30) Slot                                           : 2
  31) Port                                           : 1
4) Backup Type { PPP Peer (1) }                     : PPP Peer
  40) Peer ID                                        : 5
5) Startup Timer Value { Time in Seconds after
  System Startup to wait
  for Primary to come up
  before activating
  Backup }                                           : 300
6) Activate Timer Value { Time in Seconds after
  Primary Failure to
  activate Backup }                                  : 10
7) Restore Timer Value { Time in Seconds after
  Primary restoral to
  disable Backup }                                   : 10
(save/quit/cancel)
:
```

4. Once you are satisfied with the values, enter the **save** command, followed by **<return>**.

```
: save <return>
```

The following will display:

```
Backup Service Index 1 created.
/Interface/backup %
```

Field Descriptions

The following section explains the fields and their corresponding values.

1) Description

Enter a description of the backup service in this field. Your description may consist of a maximum of 30 ASCII characters.

2) Admin Status

The available options for this field are **Enable** and **Disable**. **Enable** allows the backup service to operate. **Disable** will render the backup service inoperative without deleting it.

3) Primary Type

This field sets the type of entity that will be backed up in the case of network failure. The available options are **Physical Port** and **Frame Relay PVC DLCI**.

4) Backup Type

This field sets the entity type to be used as a backup in the event of primary failure. At this time, the only available backup type is **PPP**.

5) Startup Timer Value

This field sets the time after system startup to wait for the primary entity to come up. If the primary entity fails to come up within the defined time after system startup, the backup entity will be activated. Acceptable values are in the range of 0-65535 seconds. The default value is 300 seconds.

6) Activate Timer Value

This field sets the amount of time that the primary entity must remain in a failed state before the backup entity is activated. Acceptable values are in the range of 0-65535 seconds. The default value is 10 seconds.

7) Restore Timer Value

This field sets the amount of time the primary entity returns and remains in an operational state before the backup entity is deactivated. Acceptable values are in the range of 0-65535 seconds. The default value is 10 seconds.

Backing Up a Frame Relay PVC

Adding a backup service for a frame relay PVC is basically the same as for a physical port. The only differences are that you must specify Primary Type as **Frame Relay**, and you must specify a DLCI number. To add a backup service for a frame relay PVC:

1. Enter the **bsadd** command with no parameters, followed by **<return>**, as shown below:

```
/Interface/backup % bsa <return>
```

A screen similar to that shown below will be displayed:

```
Adding Backup Service Index           : 2
1) Description                       : Backup 2
2) Admin Status { (E)nabled, (D)isabled } : Enabled
3) Primary Type { Physical Port (1),
Frame Relay PVC (2) }                 : Physical Port
30) Slot                               :
31) Port                               :
4) Backup Type { PPP Peer (1) }       : PPP Peer
40) Peer ID                             :
5) Startup Timer Value { Time in Seconds after
System Startup to wait
for Primary to come up
before activating
Backup }                               : 300
6) Activate Timer Value { Time in Seconds after
Primary Failure to
activate Backup }                     : 10
7) Restore Timer Value { Time in Seconds after
Primary restoral to
disable Backup }                       : 10
(save/quit/cancel)                   :
:
```

2. When you first enter the command, the next unique index is assigned automatically, a default description is created (“Backup” followed by the created index number), and defaults for primary type, backup type, and all backup timers are created (as shown above).

To backup a frame relay PVC, first change the primary type. Whenever the primary type is changed, the menu will be redisplayed, because different parameters are needed to define the primary type. An example is shown below:}]


```

: 3=2
1) Description : Backup 2
2) Admin Status { (E)nabled, (D)isabled } : Enabled
3) Primary Type { Physical Port (1),
Frame Relay PVC (2) } : Frame Relay PVC
30) Slot :
31) Port :
32) DLCI :
4) Backup Type { PPP Peer (1) } : PPP Peer
40) Peer ID :
5) Startup Timer Value { Time in Seconds after
System Startup to wait
for Primary to come up
before activating
Backup } : 300
6) Activate Timer Value { Time in Seconds after
Primary Failure to
activate Backup } : 10
7) Restore Timer Value { Time in Seconds after
Primary restoral to
disable Backup } : 10
(save/quit/cancel)
:

```

To backup a frame relay PVC, specify the slot (**30=x**), port (**31=x**) and DLCI number (**32=x**) of the PVC to be backed up. Next, enter the PPP peer index (which defines ISDN call parameters and PPP parameters). Optionally, you can modify the timer values. Below is an example of backing up the port on slot 3, port 3, PVC DLCI 32 with PPP peer index 1:

Backup Services Commands

```
      : 30=3
      : 31=3
      : 32=32
      : 40=1
      : ?
1) Description : Backup 2
2) Admin Status { (E)nabled, (D)isabled } : Enabled
3) Primary Type { Physical Port (1),
   Frame Relay PVC (2) } : Physical Port
   30) Slot : 3
   31) Port : 3
   32) DLCI : 32
4) Backup Type { PPP Peer (1) } : PPP Peer
   40) Peer ID : 1
5) Startup Timer Value { Time in Seconds after
   System Startup to wait
   for Primary to come up
   before activating
   Backup } : 300
6) Activate Timer Value { Time in Seconds after
   Primary Failure to
   activate Backup } : 10
7) Restore Timer Value { Time in Seconds after
   Primary restoral to
   disable Backup } : 10
(save/quit/cancel)
:
```

Once you are satisfied with the values, enter the **save** command, followed by **<return>**:

```
: save <return>
```

A screen similar to that shown below will display:

```
Backup Service Index 2 created.
/Interface/backup %
```

Modifying a Backup Service

With the **bsmodify** command, you can modify:

- A backup for a physical port
- A frame relay PVC.

Modifying a backup for a Physical Port

To modify a backup service for a physical port:

1. Enter the **bsmodify** command, followed by the index of the Backup service, followed by **<return>**. An example is shown below:

```
/Interface/backup % bsm 1 <return>
```

A screen similar to that shown below will display:

```
Modify configuration for Backup Service Index 1
1) Description                               : Backup 1
2) Admin Status { (E)nabled, (D)isabled }   : Enabled
   Primary Type                             : Physical Port
     Slot                                    : 2
     Port                                    : 1
   Backup Type                               : PPP Peer
   Peer ID                                   : 5
5) Startup Timer Value { Time in Seconds after
   System Startup to wait
   for Primary to come up
   before activating
   Backup }                                  : 300
6) Activate Timer Value { Time in Seconds after
   Primary Failure to
   activate Backup }                         : 10
7) Restore Timer Value { Time in Seconds after
   Primary restoral to
   disable Backup }                          : 10
(save/quit/cancel)
:
```

The command works in a manner similar to the **bsadd** command, except the parameters that define the backup service may not be changed. These parameters are the:

- index
- primary type
- primary type slot, port, and dlci
- backup type, and
- peer ID.

Only the description and startup, activate, and restore timer fields may be modified.

2. Once you are satisfied with the values, enter the **save** command, followed by **<return>**, as shown below:

```
: save <return>
```

A screen similar to that shown below will display:

```
Backup Service Index 1 modified.
/Interface/backup %
```

Modifying a Frame Relay PVC Backup Service

To modify a backup service for a frame relay PVC:

1. First, enter the **bsmodify** command, followed by the index of backup service, followed by **<return>**, as shown in the example below:

```
/Interface/backup % bsm b2 <return>
```

A screen similar to that shown below will display:

```
1) Description : Backup 1
2) Admin Status { (E)nabled, (D)isabled } : Enabled
   Primary Type : Frame Relay PVC
   Slot : 3
   Port : 3
   DLCI : 32
   Backup Type : PPP Peer
   Peer ID : 1
5) Startup Timer Value {Time in Seconds after
   System Startup to wait
   for Primary to come up
   before activating
   Backup } : 300
6) Activate Timer Value {Time in Seconds after
   Primary Failure to
   activate Backup } : 10
7) Restore Timer Value {Time in Seconds after
   Primary restoral to
   disable Backup } : 10
(save/quit/cancel)
:
```

The command functions in a manner similar to the **create** command, except the parameters that define the backup service may not be changed. These parameters are the:

- index
- primary type
- primary type subparameter
- backup type, and
- backup type subparameters.

Only the Description and Timer fields may be modified.

2. Once you are satisfied with the values, enter the **save** command, followed by **<return>** at the prompt, as shown below:

```
: save <return>
```

A screen similar to that shown below will display:

```
Backup Service Index 2 modified.
/Interface/backup %
```

Viewing Backup Service(s) Configurations

With the **bsview** command, you can view the configuration of either all backup services, or a single backup service.

Viewing the Configurations of All Backup Services

To view the configurations for all backup services, enter the following command, followed by **<return>**, at the prompt:

```
/Interface/backup % bsv <return>
```

A screen similar to that shown below will display:

Backup Table Entries

Idx	Description	Primary Type	Slot/Port/ DLCI	Bkup Type	Peer Id	Strup Time	Act. Time	Rest. Time
1	Backup 1	PHYPORT	2/1	PPP	5	300	10	10
2	Backup 2	FR PVC	3/3/32	PPP	1	300	10	10
3	Backup of PVC to Chicago	FR PVC	3/3/33	PPP	7	300	0	60

Viewing the Configuration of a Single Backup Service (bsview Command)

To view the configuration for a single backup service, enter the **bsview** command followed by the index number of the backup service, followed by **<return>**, as shown in the example below:

```
/Interface/backup % bsv 2 <return>
```

A screen similar to that shown below will display:

Backup Table Entries

Idx	Description	Primary Type	Slot/Port/ DLCI	Bkup Type	Peer Id	Strup Time	Act. Time	Rest. Time
1	Backup 1	Port	3/3/32	Peer	1	300	10	10

Deleting a Backup Service

Use the **bsdelete** command to delete a backup service. Deleting a backup service will delete the backup service configuration record. If a backup is enabled (e.g. due to the primary entity being down), the backup entity will be brought down (e.g., for ISDN the call will be disconnected).

To delete a backup service, enter the **bsdelete** command followed by the index number of the backup service, followed by **<return>**, as shown in the example below:

```
/ % bsdelete 2 <return>
```

A screen similar to that shown below will display.

```
This will bring down Backup (if up) and delete Backup Service Record
Index : 1
Description : Backup 1.
Continue? {(Y)es, (N)o} (N) :
```

Enter **<return>** or **N** (the default value) to cancel the command. Enter **Y** to delete the backup service

Viewing Backup Service Statistics

To view the statistics of a back service, enter the **bsstatus** command in the following manner:

```
bsstatus b<backupIndex>
```

where **b<backupIndex>** is the service index number assigned to the service when it was created. For example, to see the statistics for a backup service with an index number of 1, enter:

```
bsstatus b1
```

A screen similar the following displays:

```
Status for Backup Index: 1.
```

```
Current State                               :Primary Up
Number of Times Primary Port Disconnected   :0
Number of Times Backup Port Disconnected    :0
Number of Times Backup Port Initiated       :0
Number of Times Backup Port Connected      :0
Number of Times Primary Port Connected      :0
```

As a variation of this command, enter the **bsstatus** command without specifying the service index number. A screen displays showing all backup services on the switch, as shown:

Idx	Description	Slot/		Bkp	Peer	Current.
		Primary Port/	Dlci/			
1		Port	5/1/0	Peer	1	Primary Up

Current State. The current state of the backup service. The options for this are **Primary Up**, **Primary Down**, **Backup Up**, **Backup Down**, **Backup Initiated**.

Number of Times Primary Port Disconnected. The number of times the primary port has disconnected since the last clearing of statistics for this service.

Number of Times Backup Port Disconnected. The number of times the backup port has disconnected since the last clearing of statistics for this service.

Number of Times Backup Port Initiated. The number of times the backup port has been activated since the last clearing of statistics for this service.

Number of Times Backup Port Connected. The number of times the backup port has connected since the last clearing of statistics for this service.

Number of Times Primary Port Connected. The number of times the primary port has connected since the last clearing of statistics for this service.

Idx. The index number of the backup service.

Description. Enter a description of the backup service in this field. Your description may consist of a maximum of 30 ASCII characters.

Primary Type. This field shows the type of entity that will be backed up in the case of network failure. The available options are **Physical Port** and **Frame Relay PVC DLCI**.

Slot/Port/Dlci. The slot, port number, and DLCI number (if applicable) attached to this backup service.

Bkp Type. This field shows the entity type to be used as a backup in the event of primary failure. At this time, the only available backup type is **PPP**.

Peer Id. The identification number of the peer that has the backup port for this service.

Current State. The current state of the backup service. The options for this are **Primary Up**, **Primary Down**, **Backup Up**, **Backup Down**, **Backup Initiated**.

Clearing Backup Service Statistics

To clear the statistics for a backup service, enter the **bsclear** command as shown:

```
bsclear b<backupIndex>
```

where **b<backupIndex>** is the service index number assigned to the service when it was created. For example, to clear the statistics for a backup service with an index number of 1, enter:

```
bsstatus b1
```

A prompt similar the following displays:

```
This will reset the statistic for Backup Index: 1  
Continue ? {(Y)es, (N)o} (n) :
```

Enter **y** to clear the statistics.

35 Troubleshooting

This chapter provides information that will help you troubleshoot Omni Switch/Router hardware and software problems. The sections within this chapter describe problems or errors you may encounter during switch hardware and software installation, configuration, or operation. Subsections within these categories reflect unique problems and provide the recommended corrective action(s).

Common problems installing switch software and possible solutions are described on page 35-5. Common network problems and possible solutions are described on page 35-6. Common hardware problems and possible solutions are described on page 35-9. And User Interface (UI) error messages, which can be used to diagnose problems, are described in page 35-11.

◆ Important Note ◆

In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. See Chapter 4, “The User Interface,” for documentation on changing from CLI mode to UI mode.

Detecting Problems

The Omni Switch/Router provides several mechanisms to detect problems. Hardware problems can be detected through:

- LEDs (OK1)
- PING tests using the **ping** command
- Network Management Software (NMS) error reporting
- Diagnostics software
- Command Line Interface (CLI) commands
- UI error messages

This chapter lists UI error messages. Refer to the appropriate hardware chapters for a complete description of LED states. Refer to NMS online documentation for explanations of NMS error messages. Refer to Chapter 25, “IP Routing,” for procedures to use the **ping** command. Refer to Chapter 36, “Running Hardware Diagnostics,” for documentation on diagnostics software. And refer to the *Text-Based Configuration CLI Reference Guide* for documentation on CLI commands.

Software problems can be detected through:

- LEDs (OK2)
- NMS error reporting
- CLI diagnostic commands (e.g., **dump** and **configuration check**)
- UI error messages

This chapter lists UI error messages. Refer to NMS online documentation for explanations of NMS error messages. And refer to the *Text-Based Configuration CLI Reference Guide* for documentation on CLI commands.

Reporting Problems

In some cases, you will not be able to correct the problem that occurs (for instance, a module failure). In such cases, you should contact Alcatel Technical Support at one of the following locations:

West Coast:

Alcatel Technical Support
26801 West Agoura Road
Calabasas, CA 91301

Telephone: 1-800-995-2696 (Domestic) 818-878-4507 (International)

Fax: 818-878-3505

Web: www.ind.alcatel.com/support

Email: support@ind.alcatel.com

East Coast:

Alcatel Technical Support
100 Nagog Park
Acton, MA 01720

Telephone: 1-800-995-2696 (domestic); 818-878-4507 (international)

Fax: (978) 264-3933

Web: www.ind.alcatel.com/support

Email: support@ind.alcatel.com

When reporting problems, you should note hardware and software details, as described in the subsections that follow.

Report Hardware Details

When reporting problems you should be ready to report the following hardware details to Alcatel Technical Support:

- Type of chassis (Omni Switch/Router or OmniAccess) and version of chassis (e.g., OmniS/R-3, OmniS/R-9)
- Serial number of chassis and module(s)
- Type of module that failed
- Hardware revision of module
- Model number of power supply
- UPS or direct connect to power source
- Any dump files on the flash file system

Report Software Details

When reporting problems you should be ready to report the following software details to Alcatel Technical Support:

- Software revision (e.g., 3.4.8, 4.3.2)
- Whether the feature never worked or was intermittent
- Bridging or routing configured
- Multiple groups or VLANs configured
- IP PING access
- Statistics incrementing correctly
- Protocols used
- Any capture file (trace) available
- Any dump files on flash file system

Understanding Problems

The following self-questions can be used to get a better idea on the nature of the problem:

- Has this functionality ever worked?
- What changes have occurred in the network? Was software upgraded? Were device(s) added?
- Are all users affected or are the problems related to a single port, module, or switch?
- Are statistics (as reported by UI commands such as **vs**, **ve**, **bps**, and **rmon**) incrementing on the affected port(s)?
- Are all protocols (routed or switched) failing?
- Can the affected device be successfully pinged via IP/IPX?
- Can a trace be captured on the affected segment(s)?
- Is an external analyzer, such as a Sniffer or Alcatel's Port Mirroring/Port Monitoring, available?

This chapter provides documentation on some common problems and potential solutions for problems with your switch in the sections that follow.

Software Installation Problems

If you encounter problems during software installation, most likely you will see error messages that indicate the problem.

If you cannot install the software, you can use the Boot Line prompt to download files via ZMODEM or a computer attached to a SLIP line. You can also temporarily set boot parameters and load from Boot Line in an attempt to load under different settings (refer to Appendix A, "The Boot Line Prompt"). For more information about loading software via ZMODEM, refer to Chapter 5, "Installing Switch Software."


Operational Problems

The following paragraphs describe operational problems you may encounter.

Deadlocked VLAN

Occasionally, a VLAN may deadlock. This may be a result of the configuration process you used when you set up the VLANs.

If, for example, you have a setup with three switches, as shown in the following figure, the VLAN can enter a deadlock. In this example, there are two switches, one configured with one VLAN (Switch X), another configured with two VLANs (Switch Y), and another switching device that connects to the VLANs (Switch Z).



Deadlocked VLANs Due to Loop

In this situation, VLAN 2 (V2) in the Switch X is in a loop because it has not learned that it has connected to Switch Y with two virtual bridges (V1 and V3), which are inside one switch. Since V2 detects a loop, it invokes blocking at the port connected to V1, which results in a deadlock. V1 and V3, inside Switch Y, can still communicate, and traffic still exits V2 in Switch X, going to V3; however, traffic will not exit V3.

To determine if this problem has occurred in your setup, you can use the **vi** command to display information about a specific port. (See Chapter 19, “Managing Groups and Ports,” for more information on the **vi** command.) The syntax for this command is as follows:

```
vi <slot>/<interface>
```

The system will show the port in Blocking mode and not in Forwarding mode.

Probable Cause

You did not configure the network from the point furthest away from the point of connection.

Solution

To rectify the problem, you should always start configuration from the switch that is the furthest away from the point of connection. In the figure on page 35-6, for example, you would start the change from V2 in Switch X. By configuring this switch first, you would set it up to see the two VLANs in Switch Y, and use two Spanning Trees to looping.

Problems with IP Applications

You may have enabled routing on a VLAN, but have problems with PING and other IP applications.

Probable Cause

When routing is enabled on a VLAN, packets will not be forwarded unless the Spanning Tree Status for the port being forwarded to has progressed from Listening to Learning to Forwarding.

You can determine if Spanning Tree Protocol has entered the Forwarding state for a port by viewing port status with the **sts** command. Refer to Chapter 17, “Configuring Bridging Parameters,” for information on Spanning Tree Status and the **sts** command.

Solution

Spanning Tree algorithms put the ports into the correct state. There may be propagation delays when the Spanning Tree passes protocol information throughout a bridged network. This is normal as bridge ports wait for new topology information and for the lifetime of frames being forwarded using the old topology to expire. Immediate transitions from port state to port state should not be expected.

If the port is in the blocking mode, then the Spanning Tree has detected a loop. Blocking is a desired, preventive measure invoked by the Spanning Tree algorithm.

You should not attempt to alter the port state or remove the Spanning Tree. If you attempt to move a port from non-participation to the forwarding state, you take the risk of introducing data loops.

Once in the Forwarding state, PINGs and other IP applications should function properly.

Protocol Problems

You may notice an abnormal number of errors in a particular protocol. You can view protocol errors by using the networking commands. Refer to Chapter 25, “IP Routing,” for more information on the networking commands.

Probable Cause

Incompatible versions of the protocol are running on stations in the network.

Solution

Check the version of the protocol and verify that you are using the same version on all stations in the network. For example, you may be required to run Spanning Tree, Revision C on all stations.

Also, check the parameter values that you set for the protocol.

Hardware Problems

The following sections describe problems you may encounter with switch hardware.

LEDs Do Not Light on All Modules

You have turned on the power supply to the switch, but the LEDs on the modules do not light.

Probable Cause

The power supply has blown a fuse.

Solution

Call Alcatel Technical Support.

Amber Color in LEDs

During power-up, the switch goes through a Power-On Self Test (POST). Results of the test are reflected in the OK1 and OK2 LEDs on the MPX and switching modules; specifically, OK1 indicates hardware failures, while OK2 indicates software failures.

The first time you start the switch, the OK1 LED will blink in amber once to indicate start mode. The OK2 LED will blink in green rapidly to indicate image loading. Thereafter, OK2 should blink slower in green to indicate operational mode.

Probable Cause

Hardware failure or software failure.

Solution

If the amber LED displays on a switching module, replace the module with a known, good module.

If the amber LED displays on the MPX, or after replacing the switching module the problem persists, shut down the switch and call Alcatel Technical Support.

Non-Blinking OK2 LED

When the switch is operating properly, the OK2 LED blinks in green. When the OK2 LED displays a steady green light, this is an indication of problems.

Probable Cause

The MPX or the software is malfunctioning. Typically the problem cannot be resolved by rebooting.

Solution

Shut down the switch and call Alcatel Technical Support.

TEMP LED is Amber

If the TEMP LED is amber, the internal temperature of the switch has exceeded the operational limit.

Solution

Perform the following steps:

1. Turn off the switch and wait until it has completely cooled down.
2. Check the immediate environment and ensure that the switch is not located in an area where it can be overheated by other heat-producing devices.
3. Ensure that the switch is located in an area where there is ample room for air flow around the chassis.

If the environment is satisfactory, check the internal cooling fans. The switch is shipped with redundant fans that start automatically when you power up the unit. Try powering up and listen for the fan motors. Also, you should feel a slight air flow near the chassis. If the fans are not working, power down and contact Alcatel Technical Support.

STA LED Is Off

There is one status LED per port on Ethernet switching modules. When lit, it indicates that a good cable connection exists to an Ethernet device.

Probable Cause

The LAN cable is not connected properly or is faulty.

Solution

Check all port connections and inspect the cable. If you find a faulty cable, replace the cable.

Switch Does Not Boot When Flash File System Is Full and Trying To Create the `mpm.cnf` File

You may have saved too many files to the flash file system. If the flash file system is full, it will be unable to create the `mpm.cnf` file and it will be unable to complete the boot process.

Probable Cause

Unnecessary image or data files exist in the flash file system.

Solution

Follow the steps below to free up memory in the flash file system.

1. Reboot the switch and enter the Boot Line prompt. (See Appendix A, "The Boot Line Prompt," for more information.)
2. Use the Boot Line **R** command to delete any necessary files. Make sure you have enough room for the switch to create the `mpm.cnf` file.
3. Use the Boot Line **@** command to continue the boot process.

Error Messages

This section provides error messages that you may encounter in the UI.

Understanding Error Messages

Error messages reflect hardware or software problems that the switch encountered during initialization, configuration, or operation.

In some instances, the messages that display on the UI show the C program function name. For example:

cmSetTTY(): Illegal port requested

where **cmSetTTY** represents the function, and **()** indicates that parameters are passed. This information is for internal debugging purposes.

In this section, the phrase **xxx** in error messages represents a value that is specific to that message. For example, in the message **board type xxx**, the specific board type displays in the error message.

Correcting Errors

In most cases, you will not be able to correct error conditions that result because of internal hardware or software malfunctions. You should contact Alcatel Technical Support when you receive these messages. Refer to *Reporting Problems* on page 35-3.

You can correct error conditions that result because incorrect parameter values were entered during configuration. The tables that follow list error messages to which you can respond.

Module Startup/Shutdown Error Messages	
Message	Corrective Action
False Shutdown: restarting to handle queued msgs	This message does not reflect an error condition. No action required.
P3 diags failed...	Message results when the module fails diagnostic tests. Try replacing the module.
Download failed	Try replacing the module.
No reply from VSE driver board-up request	Try replacing the module.
No reply from MBox	Try replacing the module.

Serial Port Configuration Errors	
Message	Corrective Action
Problem deleting SLP port xxx, errno=xx	Reboot the system, then use the Boot Line configuration to force SLIP down at the boot line (refer to Appendix A “The Boot Line Prompt”).
Can't modify SLIP if it's not up! current mode=xxx	Reconnect the SLIP line; reconfigure using the slipc command.
Problem changing SLIP remote IP addr to xxx	Check the remote IP address by using the slipc command at the UI. Refer to Chapter 6, “Configuring Management Processor Modules.”
Couldn't setup SLIP port slxxx on xxx	Reboot the switch.

Module Connection Errors	
Message	Corrective Action
interrupt: Link Error Monitor ALERT on xxx/xxx PHY-xxx	If this message shows up once or twice, it probably means that someone is plugging a new cable in slot/port xxx/xxx, physical connector xxx. If it displays more frequently, then there is probably a bad CDDI or FFDI connection on slot/port xxx/xxx, physical connector xxx, caused by either dirty connectors or bad cabling. Try cleaning the connections or replacing the cabling.

Chassis Error Messages

The slots in the messages within the following table are all zero based. That is, Slot 1 will be displayed as “Slot 0,” Slot 2 will be displayed as “Slot 1,” etc.

Chassis Error Messages Table	
Message	Corrective Action
Problem deleting SLP port xxx, errno=xx	Reboot the system, then use the Boot Line configuration to force SLIP down at the boot line (refer to Appendix A “The Boot Line Prompt”).
Unknown mod type xxx in slot xxx	Remove the module from the slot.
Board xxx needed to be restarted at xxx	The module appears dead. Remove the module from the slot and replace with a known good module.
Chassis mgr discovered xxx has a problem!	The software has discovered a dead task. The system will reboot automatically.
System seems to have (perhaps) recovered. A reboot may not be unwise, however.	The system encountered an unexpected condition. Reboot the switch.
cm_Mod_Event(): the slot wasn't empty	The system is confused. Clear the system by rebooting it.
ERROR: can't read ID info from MPM in slot xxx...shutting down chassis manager	This may indicate a bad MPX. Try power cycling.
Please run cmConfigEPROMxxx and reboot	This may indicate a bad MPX. Try power cycling.
Can't read ID info from slot.xxx fail...	This may indicate a bad module in slot xxx. Try power cycling.
cm_Mod_Event(): slot was already empty!	Reboot the system.
Problem reading ID PROM on module xxx	Try power cycling. If the problem remains, remove the module and try another slot.
ID PROM on module xxx has unknown format number xxx	Try power cycling. If the problem remains, remove the module and try another slot.
Real-Time Clock not set yet! Starting at zero.	Reset the clock by using the uic command.
Unknown modem stop bits=xxx	Change stop bits by configuring boot line (refer to Appendix A “The Boot Line Prompt”).
Couldn't read reset count, returning 0	This message appears only once if the configuration file is removed.

continued on next page...

Chassis Error Messages Table (Cont.)	
Message	Corrective Action
Couldn't read chassis description, setting default	Enter a new chassis description with the syscfg command.
cmSavePortInfo() successful	This message does not indicate an error.

36 Running Hardware Diagnostics

Hardware diagnostics provide you with software tools for diagnosing hardware-related problems on Omni Switch/Router switching modules. These diagnostics allow you to test switching modules off-line during network down time.

The Omni Switch/Router have a variety of switching modules interconnected by a frame backplane and a management backplane. When a hardware failure occurs, the problem may be related to a number of different failures. As part of a systematic troubleshooting procedure, you can use the built-in diagnostic software to test basic connectivity and functionality.

The diagnostic software includes two basic types of tests: static tests and port tests. Static tests verify the basic functions of memory and control/status registers of submodules. Port tests check for data packet processing functions.

You can run the tests individually or sequentially. Diagnostic software also provides an option that allows you to run all the tests in one session (exception: WSX modules require power recycle after static test). The diagnostic tests performed vary, depending on the switching module type under test.

◆ Important Note ◆

For Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. Because Hardware Diagnostics are supported only in UI mode, you must change from CLI mode to UI mode to run Hardware Diagnostics. See Chapter 4, “The User Interface” for information on changing from CLI mode to UI mode.

The following tests are available for the Omni Switch/Router:

- **alpreg** Alpine ASIC Register Test
- **csr** Command Status Register Test
- **gigareg** Giga-Chip ASIC Register Test
- **hrexmem** HRE-X Memory Test
- **hrexport** HRE-X Port Test (MPX only)
- **ifled** Submodule LED Test
- **ilb** Internal Loopback Test (replaces **mloopphy** in Release 3.4 and later)
- **ilbstress** Internal Loopback Stress Test
- **mamcam** Mammoth CAM Test
- **mammem** Mammoth ASIC Register and Memory Test
- **mloopmac** Mammoth MAC Loopback Test
- **morreg** Moriah Register Test

- **mvbus** Mammoth VBUS Test
- **pcam** Pseudo CAM Test
- **port** Port Traffic Test
- **stress** Port Stress Test (available for Ethernet modules)
- **submem** Submodule Local Memory Test
- **sunl** SUNI Register Test
- **tellreg** Telluride Register Test
- **whsreg** Whistler Register Test
- **wsmcable** WSX Cable Connection Test
- **xcam** Alcatel CAM Off-Board Test

Running Diagnostics

You must log in to the **diag** account to access the hardware diagnostics functionality or use the **framefab** and command.

There are several image files used for hardware diagnostics. These files have the following uses:

- **diagx.img** Omni Switch/Router diagnostics image file
- **desx.img** Omni Switch/Router stress test image file

◆ Note ◆

To function properly, hardware diagnostics must be run offline (i.e., the switch should not be connected to a network) or during network downtimes. In addition, spanning tree must be set to **OFF** via the **stc** command. For details on using the **stc** command, see Chapter 17, “Configuring Bridging Parameters.”

The OK2 LED of the module under test will be set to red if a failure is detected by diagnostic testing. The OK2 LED can be restored by resetting the module or by rebooting the chassis.

Diagnostics may not run if the **mpm.cfg** and **mpm.cnf** files are not in their default configurations. In addition, some diagnostics may affect the settings in configuration files. Therefore, any customized **mpm.cfg** and **mpm.cnf** files should be saved prior to testing. Once testing is completed, these files should be restored and the chassis rebooted prior to normal operation.

The default **mpm.cfg** and **mpm.cnf** files are obtained by performing the following steps:

1. Remove these files from flash memory by renaming the files to names besides **mpm.cfg** and **mpm.cnf**. For example, you can rename **mpm.cfg** to **mpm_cfg.old** to highlight the fact that it is the original version of the file.
2. Delete the **mpm.cfg** and **mpm.cnf** files from flash memory.
3. Reboot the system. The MPX will create default **mpm.cfg** and **mpm.cnf** files when these files are missing from flash memory. These default files are the ones to be used with diagnostic software.

Login to Run Diagnostics

You must log in to the **diag** account to access the hardware diagnostics functionality. The **diag** user is a superset of the **admin** user. The **diag** user can run all hardware diagnostics in addition to all of the capabilities available to the **admin** user. The default password for the **diag** user is **switch**.

Once logged in as a **diag** user, the Main Menu will display as follows.

Command	Main Menu
File	Manage system files
Summary	Display summary info for VLANs, bridge, interfaces, etc.
VLAN	VLAN management
Networking	Configure/view network parameters such as routing, etc.
Interface	View or configure the physical interface parameters
Security	Configure system security parameters
System	View/set system-specific parameters
Services	View/set service parameters
Switch	Enter Any to Any Switching Menu
Help	Help on specific commands
Diag	Display diagnostic level commands
Exit/Logout	Log out of this session
?	Display the current menu contents

Note the menu listing for **Diag** underneath the **Help** sub-menu. To access the diagnostics sub-menu, enter **diag** at the prompt. If the display mode is set to verbose, the diagnostics sub-menu will display as follows:

Command	Diagnostic Menu
reset	Reset a module in a slot
maskta	Control masking of temperature alarm led
test	Run tests on one or more slot modules
framefab	Run the Frame Fabric Tests
testdisp	Display test blocks on one or all slot modules
testcfg	Configure test parameters on one or all slot modules

The **test** command is the main interface into the diagnostics functionality; you must log in as **diag** to run this command. The **testdisp** and **testcfg** commands also require being logged in as **diag** to run these commands. The **reset** and **maskta** commands have specialized functionality; you do not have to be logged in as **diag** to use these commands, but you do at least need to be logged in as **admin**. Each of the sub-menu options are described in the sections that follow.

Resetting a Switching Module

The **reset** command initiates a soft reset on the module in a specified slot. Conceptually, resetting a switching module with this command is similar to switching off power to the module; the module will be in the same state after a reset as it is after a power on.

◆ Notes ◆

Some NI modules do not support the **reset** command.

The primary MPX module cannot be reset. To reset the secondary MPX, use the **secreset** command, which is described in Chapter 6, “Configuring Management Processor Modules.”

To reset a switching module, enter the **reset** command followed by the slot number for the module. For example, to reset the switching module in slot 4, enter:

```
reset 4
```

A message similar to the following displays:

```
Resetting slot of type xxxx may crash system
Attempt reset anyway {Y/N}? (N) :
```

Enter a **Y** and press **<Enter>** at this point. The module will be reset and the following message will indicate the reset took place:

```
resetting slot 4 to enable
```

Disabling a Switching Module

The **reset** command can also be used to disable a switching module. When used in conjunction with the **swap** command, this option is useful if you want to hot swap a module. (See Chapter 3, “Omni Switch/Router Switching Modules,” for information on how to hot swap a switching module.)

To disable a switching module, enter the **reset** command followed by the slot number for the module and followed by **disable** at the system prompt. For example, to reset the switching module in slot 4, enter:

```
reset 4 disable
```

To enable the switching module again, enter the reset command followed by the slot number for the module, and followed, optionally, by **enable** (**enable** is the default for the **reset** command). For example, to enable a previously disabled switching module in slot 4, enter:

```
reset 4 enable
```

Temperature Masking

The **maskta** command provides a way of modifying the behavior of the temperature alarm to mask the effect of the temperature sensor. By masking the temperature alarm bits, you can ensure that the MPX's TEMP LED doesn't signal or that it resets after a specified delay time. By default, temperature masking is disabled.

To enable temperature masking, enter

maskta enable

This command masks the temperature alarm completely. The TEMP LED will not signal, even if the temperature exceeds the set ranges. The following message confirms the masking:

Masking of Temperature Alarm enabled

You could also enable temperature alarm masking but not mask the alarm completely. If you enter an integer after the **maskta enable** command, the TEMP LED will still signal, but it will reset after the number of minutes you specified. For example, if you enter the command

maskta enable 5

the temperature alarm will still signal, but it will reset automatically five (5) minutes after the alarm-initiating event occurs.

◆ Note ◆

Once you enter a minute value when enabling temperature alarm masking, that value is saved even if you disable masking. To reset the minute value, you must re-enable temperature alarm masking and set the minute value to zero (i.e., enter the command **maskta enable 0**).

To disable temperature alarm masking, enter:

maskta disable

This is the default setting, so you only need to specify this command if you had previously enabled alarm masking. The following message confirms that you disabled masking:

Masking of Temperature Alarm disabled

Running Hardware Diagnostics

The **test** command initiates one or more test routines on a switching module that you specify. You can also optionally test all switching modules in one test session. Test status, instructions, and a summary of results are provided as output. Start a diagnostic test session using the following command syntax:

```
test <slot_number> [<repeat_count> [<test_name>]]
```

where

- <slot_number>** Indicates the slot number in the Omni Switch/Router for the module on which you want to run tests. If you enter **all** for this parameter, then all switching modules in the chassis will be tested. This parameter is required; if you do not enter a slot number then the test session will not start.
- <repeat_count>** Indicates the number of times to run the specified tests on the module. This value can be an integer between 0 and 999. A value of zero (0) repeats the test infinitely. The default value is 1. This default will be assumed if you do not enter a **repeat_count**.
- <test_name>** Indicates the test to be performed on the module. You can indicate the test name or **all** to run all tests. You can enter only one test name or **all**. The default is **all**. This default will be assumed if you do not enter a **test_name**.

◆ Note ◆

A combination of **repeat_count** set to **0** and **test_name** set to **all** allows the user to run either the port test infinitely or all off-board tests infinitely. If the user chooses to run the port test when prompted, all the static tests (memory and control/status register tests) are run once, followed by an infinite run of the port test. See *Sample Command Lines* on page 36-9 for more information.

Descriptions of each test follows:

- alpreg** Tests the Alpine registers. Test the Alpine control logic, registers, and data/address lines.
- csr** Tests the command/status registers. Includes testing management bus buffers, management bus read/write control logic, reset and LED memory, ID EEPROM, and reset circuitry.
- gigareg** Tests the Giga-Chip registers. Test the Giga-Chip control logic, registers, and data/address lines.
- hrexmem** Tests the HRE-X's local memory. Includes testing the HRE-X read/write functions, data/address, and the memory.
- hrexport** Tests the HRE-X's functions. Packets are generated by the MPX, sent out to the port, and claimed by the HRE-X. The HRE-X will insert additional routing information to the claimed packet and place it back on MVBUS to be claimed and verified by the MPX. This test can be bypassed. See *Running Diagnostics on an Entire Chassis* on page 36-20.
- ilb** Performs a port test using the internal loopback at the PHY or framer interface. Packets are generated by the MPX and sent out to the port and returned through an internal loopback within the PHY or framer. The MPX verifies the packets on a bit by bit basis.

ilbstress	Performs a stress test using the internal loopback at the PHY or framer interface. Packets are generated by the MPX and sent out to the port and returned through an internal loopback within the PHY or framer. The MPX verifies the packets on a bit by bit basis. See the description for stress test on page 36-8. If Ethernet type switch is tested, this test requires the desx.img file to be in the flash memory.
mamcam	Tests the Mammoth CAM. Tests the Mammoth CAM control logic, CAM access, and the data line and buffers.
mammem	Tests the Mammoth registers and memory. Includes testing the Mammoth control logic, registers, internal memory, internal cache, external SDRAM, SRAM, and data/address lines.
mloopmac	Performs a port test using the internal loopback within the Mammoth MAC chip. Packet are generated by the MPX and sent out to the port and returned through an internal loopback within the Mammoth MAC chip. The MPX verifies the packets on a bit by bit basis.
morreg	Tests the Moriah registers. Test the Moriah control logic, registers, and data/address lines.
mvbus	Tests the mammoth VBUS circuitry. Frames are generated within the Mammoth buffer system, sent out the VBUS, and then received on various Mammoth queues. Data integrity is verified.
pcam	Tests the HRE-X Pseudo CAM. Tests the HRE-X Pseudo CAM control logic, CAM access, and the data line and buffers.
port	Functional testing of physical ports with a burst of data packets generated by the MPX. Packets are generated by the MPX, sent out the physical port, looped back through external cables or wrap plugs, and returned to the MPX. The returned packets are verified bit by bit by the MPX. The port test requires the use of external cables or wrap plugs. The system will provide user with instructions for setting up external cables or wrap plugs for port test and prompts the user for input upon completion of setup. This test can be bypassed if cables are not available. For more information on port tests, see <i>Port Tests</i> on page 36-9. For information on cables required for the port test, see <i>Omni Switch/Router Port Test Wrap Cable/Plug Requirements</i> on page 36-10 for the Omni Switch/Router.

◆ **Important Note** ◆

For VSD and VSA submodules, the **port** test requires no outside cabling. It is a combination of multiple static tests for the submodule, rather than a traffic test.

stress	Functional testing of physical ports with continuous full-wire traffic. The data packets are initially generated by the MPX, sent out the physical port, and looped back through external cables or wrap plugs. Once the packets are returned, modifications in the packets' destination address allows the packets to continuously circulate between the NI CPU and the external cables or wrap plugs for a predefined period. Once the predefined period is reached the packets are returned to the MPX. The packets are checked on a bit by bit basis by the MPX. If Ethernet type switch is tested, this test requires the desx.img file to be in the flash memory. Stress test requires the use of external cables or wrap plugs. The system will provide user with instructions for setting up external cables or wrap plugs for stress test and prompts the user for input upon completion of setup. For more information on port tests, see <i>Port Tests</i> on page 36-9. For information on cables required for the port test, see <i>Omni Switch/Router Port Test Wrap Cable/Plug Requirements</i> on page 36-10 for the Omni Switch/Router.
submem	Tests the submodule's local memory. Includes testing local memory control logic, data/address lines, and local memory.
sunl	Tests the SUNI registers. Includes testing the SUNI control logic, registers, and data/address lines.
tellreg	Tests the Telluride ASIC registers. Test the Telluride ASIC control logic, registers, and data/address lines.
whsreg	Tests the Whistler registers. Test the Whistler control logic, registers, and data/address lines.
wsmcable	Tests the detection of DCE and DTE cables by the WSX circuitry. The operator is prompted for the appropriate cable connection.
xcam	Tests the Alcatel CAM. Tests the Alcatel CAM control logic, CAM access, and the data line and buffers.

Sample Command Lines

There are numerous ways to specify a test session through the **test** command. The following are some sample command lines along with a description of what they test. The following command:

```
test all 100 vram
```

would run the VRAM test on all the modules in the chassis that are capable of executing the VRAM test for 100 times. In another example, the following command:

```
test 3 0 all
```

would run either all the static tests or the port test on the module in slot 3 infinitely. Finally, the following command:

```
test 4 5
```

would run all tests (the default) on the module in slot 4 five (5) times.

Halting Diagnostic Tests in Progress

Depending on how many tests and repeat iterations you specify, a test session could take some time to complete. If you need to halt in-progress tests, enter **CTRL-C**. This key sequence pauses the testing and provides a test summary report. You will be prompted to resume or terminate the testing after the pause.

◆ Note◆

During certain phases of diagnostic testing, the **CTRL-C** will not be immediately processed. This delay may last several seconds, or longer.

Port Tests

Because port-to-port cabling is required, port tests may not be available on some modules with only one port, one daughtercard, or on some modules with mismatched daughtercards. (For example, 100BaseTx modules cannot run port tests with single or mismatched daughtercards.) When a port test is run, packets are generated in the MPX and sent out to the switching module, externally looped, and sent back to the MPX. The MPX then inspects the packets. The tables on the following pages provide specific cable/plug information.

◆ Important Note ◆

For VSD and VSA submodules, the **port** test requires no outside cabling. It is a combination of multiple static tests for the submodule, rather than a traffic test.

The table below provides specific cable/plug information for Omni Switch/Router switching modules.

Omni Switch/Router Port Test Wrap Cable/Plug Requirements	
Module Type	Cable Type
GSX-K-FM-2W	Port/Stress (Full Duplex) test: Multi-mode fiber optic wrap plug with SC connectors. Port/Stress (Half Duplex) test: Multi-mode fiber optic cable with SC connectors.
GSX-K-FS-2W	Port (Full and Half Duplex) and Stress tests: Single-mode fiber optic cable with SC connectors.
ESX-K-100C-32W	Port/Stress (Full Duplex) test: ESX Wrap Plug. Refer to <i>ESX Wrap Plug – RJ-45 Connector</i> on page 36-22. Port/Stress (Half Duplex) test: Ethernet Crossover Wrap Cable Refer to <i>Ethernet Crossover Wrap Cable — Category 5 UTP Copper Cable with RJ-45 Connectors</i> on page 36-22.
ESX-K-100FM-16W	Port/Stress (Full Duplex) test: Multi-mode fiber optic wrap plug with MT-RJ connectors. Port/Stress (Half Duplex) test: Multi-mode fiber optic cable with MT-RJ connectors.
ESX-K-100FS-16W	Port/Stress (Full Duplex) test: Single mode fiber optic wrap plug with MT-RJ connectors. Port/Stress (Half Duplex) test: Single mode fiber optic cable with MT-RJ connectors.

continued on next page...

Omni Switch/Router Port Test Wrap Cable/Plug Requirements (cont.)	
Module Type	Cable Type
WSX-S-2W (no compression)	Twisted pair 28GA serial cable with HD50-26 connectors – DCE to DTE.
WSX-SC-4W	Twisted pair 28GA serial cable with HD50-26 connectors – DCE to DTE.
WSX-SC-8W	Twisted pair 28GA serial cable with HD50-26 connectors – DCE to DTE.
WSX-BRI-SC-2W	BRI S/T Crossover Wrap Cable. Refer to <i>BRI S/T Crossover Wrap Cable — Category 5 UTP Copper Cable with RJ-48 (RJ-45) Connectors</i> on page 36-23. Twisted pair 28GA serial cable with HD50-26 connectors – DCE to DTE.
WSX-FE1-SC-2W	T1/E1 Crossover Wrap Cable. Refer to <i>T1/E1 Crossover Wrap Cable — Category 5 UTP Copper Cable with RJ-45 Connectors</i> on page 36-22. Twisted pair 28GA serial cable with HD50-26 connectors – DCE to DTE.
WSX-FT1-SC-2W	T1/E1 Crossover Wrap Cable. Refer to <i>Figure T1/E1 Crossover Wrap Cable — Category 5 UTP Copper Cable with RJ-45 Connectors</i> on page 36-22. Twisted pair 28GA serial cable with HD50-26 connectors – DCE to DTE.
VSD-128M-12CH VSD-128M-24CH VSD-128M-36CH VSD-128M-48CH VSD-128M-60CH VSA-FXO VSA-FXS VSA-4	No Cable Required.

Sample Test Session: Ethernet Module

Test sessions and results will vary among the various switching modules. This section shows the output from a test session on an ESX-C-12. The module is in slot 3 and all tests were requested to be run one time. The command to start this test is:

test 3

After you enter the **test** command line, the following displays:

**Port Tests are available for the selected slot(s).
These tests require external cabling.**

Do you wish to run the Port Tests (y/n) (y)

Enter **y** to run port tests or **n** to skip them. If you select to run the port tests, you will be instructed on how to cable the ports. This cabling will vary depending on the test configuration, module type, number of ports and cable type. In this example, the following displays:

Connect the following cables on Slot 3:
Port 1 to Port 2
Port 3 to Port 4
Port 5 to Port 6
Port 7 to Port 8
Port 9 to Port 10
Port 11 to Port 12

Press <Enter> when finished.

Cable the ports according to the instructions. For Ethernet tests, you should use cross-over cable to connect the ports. Press **<Enter>** when you have finished the cabling.

The module is reset, and then the rest of the tests will run.

Testing Slot 3 - Ether/12
 Resetting slot 3...
 Test In Progress: CSR Test
 OK1, OK2 LEDS will display the following pattern: OFF RED OFF GREEN OFF
 AMBER OFF - Passed
 Test In Progress: VRAM Test - Passed
 Test In Progress: CAMOFFBRD Test(1K) - Passed
 Loading dni.img...
 Test In Progress: BOARDUP Test - Passed
 Test In Progress: CAMONBRD Test(1K) - Passed
 Test In Progress: VBUS Test - Passed
 Restoring slot 3...
 Test In Progress: PORT Test (3-0)
 Wait for ports to come up . Done.
 Error - Frame #1 not found - Failed
 FAILED - PORT TEST: Tx Port1 -> Rx Port2 at Test Number 95001
 Expected Data: 1
 Measured Data: 0

Test Summation:

Started: WED DEC 17 10:48:13 2000

Slot 3	Passes	Fails
Ether/12 (3-0)		
CSR	1	0
VRAM	1	0
CAMOFFBRD	1	0
BOARDUP	1	0
CAMONBRD	1	0
VBUS	1	0
PORT	0	1

Failure Summation:

Ether/12 (3-0)

Test	Fail No.	Test No.	Exp. Data	Meas. Data	Iter. No.	Time	Temp (C)
PORT	1	95001	00000001	00000000	1	10:49:47	30.5

Completed: WED DEC 17 10:49:47 2000

Disconnect the following cables on Slot 3:

- Port 1 to Port 2
- Port 3 to Port 4
- Port 5 to Port 6
- Port 7 to Port 8
- Port 9 to Port 10
- Port 11 to Port 12

Press <Enter> when finished.

The tests are complete at this point. A summary of the test results and failures is displayed at the end of the test sequence. In this example, the module passed all tests except the port test. The ESX-K-C-32 module in slot 3 should have a red OK2 LED to indicate diagnostics failure. And the **Failure Summation** section displays only the first three failures when you request multiple test iterations.

You should now disconnect the cables used in the external loopback tests. Press **<Enter>** and the module will be restored to its normal, pre-testing state. The OK2 LED will remain red until the module is reset or the chassis is rebooted.

The main system command prompt re-displays.

Displaying Available Diagnostic Tests

The **testdisp** command provides the user with a display of applicable tests for a particular slot or for the entire chassis configuration. To display available diagnostic tests for a switching module, enter the **testdisp** command followed by the slot number for the module. The slot number is an integer ranging from 1 to the number of slots in the chassis (3 for 3-slot Omni Switch/Routers; 5 for 5-slot Omni Switch/Routers and 9 for 9-slot Omni Switch/Routers).

No default value is set and input must be provided at the time of entering the command. For example, to display available diagnostic tests for the switching module in slot 3, enter:

```
testdisp 3
```

at the system prompt. The following is a sample display.

```
Ether/12 (3-0)  
CSR - Tests the Command/Status Registers  
VRAM - Tests the VRAM  
CAMOFFBRD - Tests the CAM  
BOARDUP - Basic NI Tests  
CAMONBRD - Tests the CAM  
VBUS - Tests the VSE/SAM  
PORT - Tests the Ports
```

To display all available diagnostic tests for the entire chassis, excluding slot(s) occupied by an MPX without an HRE, enter:

```
testdisp all
```

at the system prompt. The tests are displayed per slot module starting from slot module 1.

Configuring the Diagnostic Test Environment

The **testcfg** command allows the user to tailor diagnostic testing characteristics per slot module. To configure diagnostic tests for a switching module, enter the **testcfg** command followed by the slot number for the module. The slot number is an integer ranging from 1 to the number of slots in the chassis (3 for 3-slot Omni Switch/Routers; 5 for 5-slot Omni Switch/Routers and 9 for 9-slot Omni Switch/Routers).

The **testcfg** command allows the user to bypass testing individual slots when running the **test all** command. In addition, the **testcfg** command allows the user to configure the port speed and port mode for applicable Ethernet modules for tailoring of individual slots during diagnostic testing.

No default value is set and input must be provided at the time of entering the command. For example, to configure applicable diagnostic tests for the switching module in slot 4, enter:

```
testcfg 4
```

at the system prompt. The following is a typical example.

```
Test Configuration for slot 4
```

```
1) Skip this slot during test { No (1),  
                               Yes (2) } : No
```

```
Enter (option=value/save/cancel)      :
```

Note that for all switching modules other than modules, the **Skip this slot during test** option is the only available one. See *Configuring Tests for Ethernet Modules* on page 36-17 for information on using the **testcfg** command with Ethernet modules.

Skip this slot during test. Allows the user to select to bypass this slot when the **test all** command is issued. The default is **No**. If you want the **test all** command to skip this module, enter

```
1=2
```

The following will then be displayed.

```
Test Configuration for slot 4
```

```
1) Skip this slot during test { No (1),  
                               Yes (2) } : Yes
```

```
Enter (option=value/save/cancel)      :
```

Enter **save** if you want to make this change. If you enter **save**, the change will be made and the following will be displayed.

```
Configuration Saved
```

If you want to cancel this change, enter **cancel** and the **testcfg** command will terminate and the following will be displayed.

```
Exiting menu - Test Configuration not modified
```

Configuring Tests for Ethernet Modules

Tailoring of applicable Ethernet modules includes selection of Port Speeds and of Port Modes. To configure applicable diagnostic tests for an Ethernet 10/100 switching module in slot 3, enter:

```
testcfg 3
```

The following is a sample display of the test configuration for an Ethernet 10/100 switching module.

Test Configuration for slot 3

```
1) Skip this slot during test { No (1),
                               Yes (2) } : No
2) Port Speed { 10/100 (1),
                100   (2),
                10    (3) }           : 10/100
3) Port Mode { Full Duplex (1),
              Half Duplex (2) }       : Full Duplex
Enter (option=value/save/cancel)     :
```

To change any of the values above, enter the line number, followed by an equal sign, and followed by the new value. For example, to change the **Port Mode** field to half duplex, enter

```
3=2
```

The configurable fields displayed by the **testcfg** command for an Ethernet module are described below.

Skip this slot during test. Allows the user to select to bypass this slot when the **test all** command is issued. The default is **No**.

Port Speed. Allows the user to select module port speed during the diagnostic port test. Selection includes 10/100BaseT, 100BaseT, or 10BaseT. The default is **10/100BaseT**, which alternates the speed of the port test from 10 to 100 on each pass of the port test.

Port Mode. Allows the user to select module port mode during diagnostic port test. Selection includes Full Duplex or Half Duplex. The default value is **Full Duplex**.

Enter **save** if you want to make this change. If you want to cancel this change, enter **cancel** and the **testcfg** command will terminate.

Running Frame Fabric Tests on Omni Switch/Routers

You can test the Omni Switch/Router Multi VBUS (MVBUS) backplane and the frame fabric ASIC of every switching module with the **framefab** command. The syntax for this command is as follows:

```
framefab [<repeat_count> | ilb <repeat_count>]
```

The **<repeat_count>** option lets you set the number of times to run the test, which can be from 0 to 999. If you enter **0**, the **framefab** test will continue indefinitely. If you do not use the **<repeat_count>** option, then the **framefab** test will be executed once.

Using the **<repeat_count>** option requires the use of external cables or wrap plugs for the first physical port of every switching module in the chassis. The external cables or wrap plugs used in this test are identical to the one listed in the full duplex port test. See *Omni Switch/Router Port Test Wrap Cable/Plug Requirements* on page 36-10 for more information.

The **ilb** option, which can be used with the **<repeat_count>** option, performs an internal loop-back. Using this option performs the **framefab** test without the use of external cables or wrap plugs.

The chassis should be fully loaded (i.e., Omni Switch/Router modules in all slots) to achieve a thorough testing of both the frame fabric ASICs and the Omni Switch/Router backplane. In addition, an MPX should be installed in Slot 1.

To execute the **framefab** test indefinitely, for example, enter

```
framefab 0
```

at the system prompt. A screen similar to the following will be displayed.

```
Testing All Slots
Test In Progress: FABRIC Test
```

```
Test Summation:
```

```
Started: TUE OCT 27 18:40:31 2000
```

All Slots	Passes	Fails
FABRIC	1199	18

```
Failure Summation:
```

Test	Fail No.	Test No.	Exp. Data	Meas. Data	Iter. No.	Time	Temp (C)
FABRIC	1	110402	00004cec	00000000	6	18:50:34	43.0
FABRIC	2	110504	0000a9e4	00000000	13	18:56:45	43.0
FABRIC	3	110307	0008a6ff	00000000	159	21:26:05	43.0

```
First 3 Failure(s) Detail:
```

```
Fail No. 1 - FRAME FABRIC TEST: Slot 5 failed. No packet Received from slot: 3
Fail No. 2 - FRAME FABRIC TEST: Slot 6 failed. No packet Received from slot: 5
Fail No. 3 - FRAME FABRIC TEST: Slot 4 failed. No packet Received from slot: 8
```

— Output continues on next page —

Test Coverage:

All Fabric Inputs/Outputs not tested:

Fabric in slot 2 (ESX-C12) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Fabric in slot 3 (ESX-C12) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Fabric in slot 4 (ESX-C12) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Fabric in slot 5 (ESX-C12) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Fabric in slot 6 (ESX-C12) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Fabric in slot 7 (ESX-C12) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Fabric in slot 8 (ESX-C12) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Fabric in slot 9 (ESX-C32) has 9 inputs (0-8) and 1 output (0)
All inputs tested
All Outputs tested

Completed: WED OCT 28 16:24:04 2000

If you need to halt the **framefab** tests, press **CTRL-C**. This key sequence pauses the testing and provides a test summary report. You will be prompted to restart the testing after the pause.

◆ Note ◆

During certain phases of diagnostic testing, the **CTRL-C** will not be immediately processed. This delay may last several seconds, or longer.

If your chassis is not fully loaded, the **framefab** test will report that the frame fabric in the empty slot was not tested.

Running Diagnostics on an Entire Chassis

The **testcfg** command allows you to tailor diagnostic testing characteristics by module or for an entire chassis. (Please refer to *Configuring the Diagnostic Test Environment* on page 36-16 for configuring tests for a single module.)

For example, to configure diagnostic tests for an entire chassis, enter:

```
testcfg all
```

A screen similar to the following will be displayed.

```
Test Configuration

1) Diagnostic Mode { Normal          (1),
                  { Diagnostic       (2) } : Normal
2) Stop on Failure { Disable        (1),
                  { Enable          (2) } : Disable
3) Port Test Bypass { Disable        (1),
                   { Enable          (2) } : Disable
4) Port Test Type  { Port           (1),
                   { ILB            (2),
                   { STRESS         (3),
                   { ILBSTRESS      (4) } : Port
5) HRE-X Test Mode { Do not test HRE-X (1),
                  { Test HRE-X      (2) } : Test HRE-X

Enter (option=value/save/cancel) :
```

Select the number of the item you want to change. To change any of the values listed above, enter the line number, followed by an equal sign, and then the new value. For example, to change the port test type to **STRESS**, enter:

```
4=3
```

To update the values you have changed, enter **save**. If you do not want to save the changes enter **quit** or **cancel**, or press **Ctrl-D**. If you enter **save**, the change will be made and the following message will be displayed.

```
Configuration Saved
```

If you cancel the **testcfg** command, it will terminate and the following will be displayed.

```
Exiting menu - Test Configuration not modified
```

The fields displayed by the **testcfg** command with the **all** option are described below.

1) Diagnostic Mode

Enter **1** (the default) to set to normal diagnostics testing or **2** for a more detailed version of diagnostic testing. However, setting this field to **2** requires more user intervention during a test.

2) Stop on Failure

Enter **2** to halt diagnostics in an active state when a failure occurs or **1** (the default) to exit diagnostics and display the **Test Summation** and **Failure Summation** sections of the **test** command output. Setting this field to **2** can be used to further troubleshoot problems. However, setting this field to **2** requires more user intervention during a test.

3) Port Test Bypass

Enter **2** to complete testing of all ports regardless of port test failures or **1** (the default) to stop testing at the first port failure. Setting this field to **2** can be used to further troubleshoot problems.

4) Port Test Type

Enter **1** (the default) for a port test, **2** for an Internal Loopback (ILB) test, **3** for a stress test, or **4** for an ILB stress test. External cables are required for the port and stress tests but not for the ILB test. In addition, the stress test requires a special image file (see *Running Diagnostics* on page 36-2) and is only available for Ethernet (ESX and GSX) modules on the Omni Switch/Router.

◆ Note ◆

Option 5, **HRE-X Test Mode**, is for the Omni Switch/Router only.

5) HRE-X Test Mode


This option lets you configure port tests for HRE-Xs installed on Omni Switch/Router switching modules. It does not affect the port test for HRE-Xs installed on MPXs. Currently, the port test on HRE-Xs installed on switching modules runs in conjunction with the normal port test.

Each physical port is tested with the normal port test path and then through the HRE-X port test path before testing the next physical port. Subsequent physical ports are tested with only the normal port test path.


Enter **1** to bypass testing of the HRE-X when the port test is run or **2** to perform the test as described above.

Diagnostic Test Cable Schematics


The figures below and on the following pages provide information on port test cables and plugs.




Ethernet Crossover Wrap Cable — Category 5 UTP Copper Cable with RJ-45 Connectors



T1/E1 Crossover Wrap Cable — Category 5 UTP Copper Cable with RJ-45 Connectors



ESX Wrap Plug – RJ-45 Connector



**BRI S/T Crossover Wrap Cable — Category 5 UTP Copper Cable
with RJ-48 (RJ-45) Connectors**

A The Boot Line Prompt

When the switch boots, it requires basic information so that it can configure itself. The switch is delivered with factory default configuration parameters that provide basic information; however, you can change or customize the configuration parameters using the Boot Line prompt. You can only access the Boot Line configuration through an ASCII terminal.

Customizing parameters can be helpful when troubleshooting your system. Changing configuration items in the boot process allows you to:

- Stop the boot process
- Boot from a SLIP device
- Boot from a ZMODEM connection
- Revert back to factory default settings
- Boot/load with a different set of parameters

In addition, you can use the Boot prompt to configure an IP address for the Ethernet management port or you can use the **ethernetc** command (which is described in Chapter 6, “Configuring Management Processor Modules”). You can use the Ethernet management port to Telnet into the UI, perform high-speed software loads, or as a connection to a boot device. See *Configuring a Switch with an MPX* on page A-7 for more information on configuring the Ethernet management port with the Boot prompt.

To enter the Boot line prompt, see the section that follows. See *Boot Prompt Basics* on page A-3 for documentation on basic Boot prompt commands. If you are configuring an Omni Switch/Router see *Configuring a Switch with an MPX* on page A-7.

◆ Important Note ◆

In Release 4.4 and later, the Omni Switch/Router is factory-configured to boot up in CLI (Command Line Interface) mode, rather than in UI (User Interface) mode. Chapter 4, “The User Interface,” includes documentation on changing from CLI mode to UI mode.

Entering the Boot Prompt

Perform the following steps to reach the Boot prompt.

1. Connect an ASCII terminal (or computer with a terminal emulator) to the console port on the MPX. The default communication parameters are:
 - 9600 bps
 - 8 data bits
 - 1 stop bit
 - no parity
 - no hardware flow control (Microsoft Windows 95)
2. Turn on the switch.
3. You should see text scrolling on the terminal, indicating that the boot is starting. If you do not see any text within a few seconds of turning on the switch press the **<Enter>** key. If you still do not see any text on the screen, verify your connections, turn off the switch, and turn it back on again.
4. Once the boot process starts you have approximately two (2) seconds to interrupt the boot. Press any key during this time to enter the Boot prompt.

◆ Note ◆

MPXs in redundant configurations should not be stopped during the boot process. If you must do this, remove one of the MPXs while configuring the other.

The following screen prompt displays.

[Boot]:

See the following section for documentation on basic Boot prompt commands. If you are configuring an Omni Switch/Router see *Configuring a Switch with an MPX* on page A-7.

Boot Prompt Basics

To get a list of commands enter a question mark (?). A screen similar to the following is shown:

```

?                - print this list
@                - boot (load and go)
p                - print boot params
c                - change boot params
l                - load boot file
g adrs           - go to adrs
d adrs[,n]       - display memory
m adrs           - modify memory
f adrs, nbytes, value - fill memory
t adrs, adrs, nbytes - copy memory
e                - print fatal exception
n netif          - print network interface device address
L                - list ffs files
P                - Purge system: removes ALL ffs files
R file [files]   - remove ffs file(s)
S                - save boot configuration
V                - display bootstrap version
$dev(0,procnum)host:/file h=# e=# b=# g=# u=usr [pw=passwd] f=#
                 tn=targetname s=script o=other

Boot flags:
0x02             - load local system symbols
0x04             - don't autoboot
0x08             - quick autoboot (no countdown)
0x20             - disable login security
0x40             - use bootp to get boot parameters
0x80             - use tftp to get boot image
0x100           - use proxy arp
0x1000          - factory reset

available boot devices: sl ffs zm
[Boot]:

```

The commands for this menu are described in the sections below.

◆ Important Note ◆

The Boot prompt is case sensitive. Always enter letters in lowercase or uppercase as indicated in the menus.

Resuming Switch Boot (@)

If you wish to continue the boot process, enter the @ command at the prompt. This loads the last saved configuration.

Displaying Current Configuration (p)

To display the current configuration, enter a **p** at the Boot prompt. A screen similar to the following will be displayed.

```
Boot device           : ffs
Boot file             : /flash/mpx.img
Eth IP addr[:mask]   : 192.168.11.1
Startup script       : /flash/mpx.cmd
Console params       : 9600,n81c
Modem params         : 9600,n81
Boot flags           : 0xb
Other                 : dvip:no-name,192.168.10.1,255.255.255.0,192.168.10.255;
```

For information on modifying these screens, see *Configuring a Switch with an MPX* on page A-7.

To change the configuration of the boot parameters, enter **c** at the prompt. For more information, see *Configuring a Switch with an MPX* on page A-7.

Loading the Last Configured Boot File (l)

To load the last configured boot file, enter the **l** command. A screen similar to the following is shown:

```
Boot device           : ffs
Boot file             : /flash/mpx.img
Eth IP addr[:mask]   : 172.22.2.20
Startup script       : /flash/mpx.cmd
Console params       : 9600,n81c
Modem params         : 9600,n81d
Boot flags           : 0xb
Other                 : dvip:TECHPUB-
120,172.22.2.120,255.255.0.0,172.22.255.255;

Loading /flash/mpx.img...25320 + 2163504 + 314792
entry = 0x40e00000
```

Listing Available Files in the Flash Memory (L)

To list all of the available files in the flash memory that you could load onto the switch, enter the **L** command. A screen similar to the following is shown:

```
Files available in "/flash":
  mpx.cmd
  mpm.log
  esx.img
  mpx.img
  mpm.cnf
  mpm.cfg
  switch.ascii
[Boot]:
```

Deleting All Files in the Flash Memory (P)

To delete all flash memory files, enter the **P** command at the prompt. The following message is displayed:

```
WARNING: This will remove ALL the files in the system.
Do you want to do this? ->
```

Enter **y** at the prompt to continue. The following message is shown

```
Erasing Flash File System...Done...Rebooting...
```

The switch will automatically reboot at this point. Since there are now no files in the flash memory, you are returned to the boot prompt.

Deleting Specific Files in the Flash Memory (R)

To delete a specific file from the flash memory, use the **R** command followed by the file name. You can delete a single file or multiple files with a single command. For example, to delete the **mpx.cmd** file, you would enter **R** followed by a space, and then **mpx.cmd**, as shown:

```
R mpx.cmd
```

To delete the **mpx.cmd** and the **mpm.log** files, you would enter **R**, a space, **mpx.cmd**, a space, and then **mpm.log**, as shown:

```
R mpx.cmd mpm.log
```

Saving Configuration Changes (S)

To save any changes to the configuration parameters, enter the **S** command at the prompt. The following message appears to confirm when the process is complete:

```
Saving boot information...done  
[Boot]:
```

Viewing Version Number (V)

To view the version number of the bootstrap shell, enter the **V** command at the prompt.

◆ Important Note ◆

Some of the options within the Boot Line configuration menu are for programmer's internal debugging purposes or for Customer Service diagnostics. Alcatel does not recommend that you invoke any menu options not described in this section.

Configuring a Switch with an MPX

Perform the following steps to configure an Omni Switch/Router (MPX). You can press **Ctrl-D** at any time to return to the Boot prompt.

1. At the Boot prompt, enter a lowercase **c** to begin configuring parameters. A prompt similar to the following displays.

```
'.' = clear field;      '.' = go to previous field;    ^D = quit
Boot device             : ffs
```

2. To change the switch's boot device, (i.e., the device it will read the boot file from) enter **ffs** for the flash file system (the default), **pcn** for the Ethernet management port, **sl** for a SLIP device, or **zm** for ZMODEM.

A screen prompt similar to the following displays.

```
Boot file               : /flash/mpx.img
```

3. Enter the boot file name or press the **<Enter>** key to accept the default (**mpx.img**). For FTP downloads, the path you should enter is relative to the log-in (i.e., remote) directory. A prompt similar to the following displays.

```
Eth IP addr[:mask]     :
```

4. Enter an IP address for the Ethernet management port in dotted decimal notation. As an option, you can also enter an IP subnet mask in hexadecimal notation. If no mask is provided, the switch will try to determine the mask using Internet Control Message Protocol (ICMP) requests.

A screen prompt similar to the following displays.

```
Local hostname         :
```

5. Enter a name for the MPX here.

◆ Note ◆

Steps 6 through 10 are only important if you are booting your switch from a network.

6. A screen prompt similar to the following displays.

```
Remote IP addr[:mask] :
```

You can enter an IP address for a remote host. In addition, you can also enter an IP address mask in hexadecimal notation. If no mask is provided, it will infer it from the IP address class.

A screen prompt similar to the following displays.

```
Remote hostname       :
```

7. You can enter a remote host name. A screen prompt similar to the following displays.

```
Gateway IP addr       :
```

8. You can enter an IP address for the first hop router to a remote host (if the host is on a different IP net). A screen prompt similar to the following displays.

User :

9. You can enter a log-in name for a remote host. A screen prompt similar to the following displays.

Remote password :

10. You can enter a password for a remote host.

11. A screen prompt similar to the following displays.

Startup script : /flash/mpx.cmd

Enter the command file name or press the <Enter> key to accept the default (**mpx.cmd**). A prompt similar to the following displays.

Console params : 9600,n81c

12. You can change the parameters for the console port. To change the value, enter the baud rate (**1200**, **9600**, or **19200**, or **38400**), the parity (**n** for none, **e** for even, or **o** for odd), data length (**7** or **8**), stop bits (**0**, **1**, or **2**), and port type (**c** for console, **s** for SLIP, or **d** for down).

For example, **19200n81c** sets the console port to 19,200 baud, no parity, 8-bit data length, 1 stop bit, and console mode.

◆ **Note** ◆

If the default baud rate shunt (E1) has not been removed, any changes to the baud rate you enter will be ignored and a message to that affect is displayed during the boot process.

A screen prompt similar to the following displays.

Modem params : 9600,n81d

13. You can change the parameters for the modem port. To change the value, enter the baud rate (**1200**, **9600**, or **19200**, or **38400**), the parity (**n** for none, **e** for even, or **o** for odd), data length (**7** or **8**), stop bits (**0**, **1**, or **2**), and port type (**m** for modem, **s** for SLIP, or **d** for down).

For example, **19200n81m** sets the modem port to 19,200 baud, no parity, 8-bit data length, 1 stop bit, and modem mode.

A screen prompt similar to the following displays.

Boot flags : 0xb

14. To accept the default (**oxb**) and perform a normal boot, press the **<Enter>** key. To restore the factory-configured boot process, enter **0x1000**. The following flags should only be used for internal debugging or Customer Service diagnosis:

- **0x02** Load the local system symbols.
- **0x04** Do not autoboot.
- **0x08** Quick autoboot (no countdown).
- **0x20** Disable login security.
- **0x40** Use **bootp** to get the boot parameters.
- **0x80** Use **tftp** to get the boot image.
- **0x100** Use proxy arp.

A screen prompt similar to the following displays.

```
Other      : dvip:no-name,192.168.10.1,255.255.255.0,192.168.10.255;
```

15. You can enter the default VLAN IP parameters by entering them in the following format:

```
dvip:<host name>,<IP address>[,<IP mask>[,<IP broadcast address>]]
```

16. The following screen prompt displays.

```
[Boot]:
```

Enter an uppercase **S** to save any parameters you changed. The following screen prompt displays.


```
[Boot]:
```

17. Enter an **@** to boot your switch.

B Custom Cables

This appendix provides detailed information, including illustrations and pin diagrams, for the cables that can be used with Omni Switch/Router Submodules. These custom cables are available from Alcatel, but you can use the following information to manufacture them.

V.35 DTE Cable (For WSX-to-DCE Device Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.


J2300		V35-M	
DTE		V35-M	
B	AB		
A	SHIELD		
P	BA-A		
S	BA-B		
R	BB-A		
T	BB-B		
Y	DB-A		
AA	DB-B		
V	DD-A		
X	DD-B		
U	DA-A		
W	DA-B		
C	CA-A		
D	CB-A		
E	CC-A		
F	CF-A		
H	CD-A		
N	RL		
NN	TM		

air

r

1, CTP0, respectively

V.35 DCE Cable (For WSX-to-DTE Device Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.


V3100		V35-F	
DCE		V35-F	
B	AB		
A	SHIELD		
R	BB-A		
T	BB-B		
P	EA-A		
S	EA-B		
Y	DB-A		
AA	DB-B		
U	DA-A		
W	DA-B		
V	DD-A		
X	DD-B		
F	CF-A		
D	CB-A		
H	CD-A		
C	CA-A		
E	CC-A		
MN	TM		
N	RL		

air

r

l, CTP0, respectively

RS232 DTE Cable (For WSX-to-DCE Device Connection)

















The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.


The table on the right shows the pinouts for the connectors.

<u>12002400</u>	<u>LATE DTE</u>	<u>DB25-M</u>
	7	AB
	1	SHIELD
	2	BA-A
	3	BB-A
	15	DB-A
	17	DD-A
	24	DA-A
	4	CA-A
	5	CB-A
	6	CC-A
	8	CF-A
	20	CD-A
	21	RL
	25	TM

≅d-pair

≅TP1, CTP0, respectively

RS232 DCE Cable (For WSX-to-DTE Device Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.


The table on the right shows the pinouts for the connectors.

12005200	
ATE DCE	DB25-F
7	AB
1	SHIELD
3	BB-A
2	BA-A
15	DB-A
24	DA-A
17	DD-A
8	CF-A
5	CB-A
20	CD-A
4	CA-A
6	CC-A
25	TM
21	RL

ed-pair

CTP1, CTP0, respectively

RS530 DTE Cable (For WSX-to-DCE Device Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.

12500	DTE	DB25-M
7	AB	
1	SHIELD	
2	BA-A	
14	BA-B	
3	BB-A	
16	BB-B	
15	DB-A	
12	DB-B	
17	DD-A	
18	DD-B	
24	DA-A	
11	DA-B	
4	CA-A	
19	CA-B	
5	CB-A	
13	CB-B	
6	CC-A	
22	CC-B	
8	CF-A	
10	CF-B	
20	CD-A	
23	CD-B	
21	RL	
25	TM	


pair

r

r

l, CTP0, r respectively

RS530 DCE Cable (For WSX-to-DTE Device Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.

RS530	DCE	DE25-F
7	AB	
1	SHIELD	
3	BB-A	
16	BB-B	
2	BA-A	
14	BA-B	
15	DB-A	
12	DB-B	
24	DA-A	
11	DA-B	
17	DD-A	
9	DD-B	
8	CF-A	
10	CF-B	
5	CB-A	
13	CB-B	
20	CD-A	
23	CD-B	
4	CA-A	
19	CA-B	
6	CC-A	
22	CC-B	
25	TM	
21	RL	


Pair

r

r

1, CTP0, respectively

X.21 DTE Cable (For WSX-to-DCE Device Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.

12600	DTE	DE15-M
8	SIG GND	
1	SHIELD	
2	T-A	
9	T-B	
4	R-A	
11	R-B	
6	S-A	
13	S-B	
7	E-A	
14	E-B	
3	C-A	
10	C-B	
5	I-A	
12	I-B	


air

r

r

1, CTP0, r respectively

X.21 DCE Cable (For WSX-to-DTE Device Connection)




The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.




pair

r

r

1, CTP0, r respectively

RS449 DTE Cable (For WSX-to-DCE Device Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26 Pin HD50 Connector-male
- AMP 750850-6 26 Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.

12700	DTE	DB37-M
19	AB	
1	SHIELD	
4	SD-A	
22	SD-B	
6	RD-A	
24	RD-B	
5	ST-A	
23	ST-B	
8	RT-A	
26	RT-B	
17	TT-A	
35	TT-B	
7	RS-A	
25	RS-B	
9	CS-A	
27	CS-B	
11	DM-A	
29	DM-B	
13	RR-A	
31	RR-B	
12	TR-A	
30	TR-B	
14	RL	
18	TM	


air

r

r

l, CTP0, respectively

RS-449 DCE Cable Assembly (For WSX-to-DTE Device 75Ω Connection)



The following parts are recommended for the end of the cable connected to the WSX.

- AMP 750833-1 26-Pin HD50 Connector-male
- AMP 750850-6 26-Pin HD50 Backshell

Parts for the customer end of the cable can be of any industry-standard manufacturer. Use of a shielded-type connector is recommended.

Cable should be constructed with data-comm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.

The table on the right shows the pinouts for the connectors.

15500	
DCE	DB37-F
19	AB
1	SHIELD
6	RD-A
24	RD-B
4	SD-A
22	SD-B
5	ST-A
23	ST-B
17	TT-A
25	TT-B
8	RT-A
26	RT-B
13	RR-A
31	RR-B
9	CS-A
27	CS-B
12	TR-A
30	TR-B
7	RS-A
25	RS-B
11	DM-A
29	DM-B
18	TM
14	RL


pair

r

r

l, CTP0, r respectively


RJ-45 to DB15F Cable Assembly (For T1/E1 Port 120Ω Connections)




The following parts are recommended for the ends of the cable:

- For the switch side of the cable assembly (P1): 8-conductor RJ-45 round connector (MTP-88U or equivalent)
- Parts for the customer end of the cable (P2) can be of any industry-standard manufacturer. Use of a shielded-type DB-15 female connector is recommended.

Cable should be constructed with datacomm-quality cable that has an overall mylar foil shield and braided-shield, terminated to the appropriate pins and connector shell at each end of the cable. Twisted-pair 28GA cable is preferred, with any of the pairs used for non-paired signals.



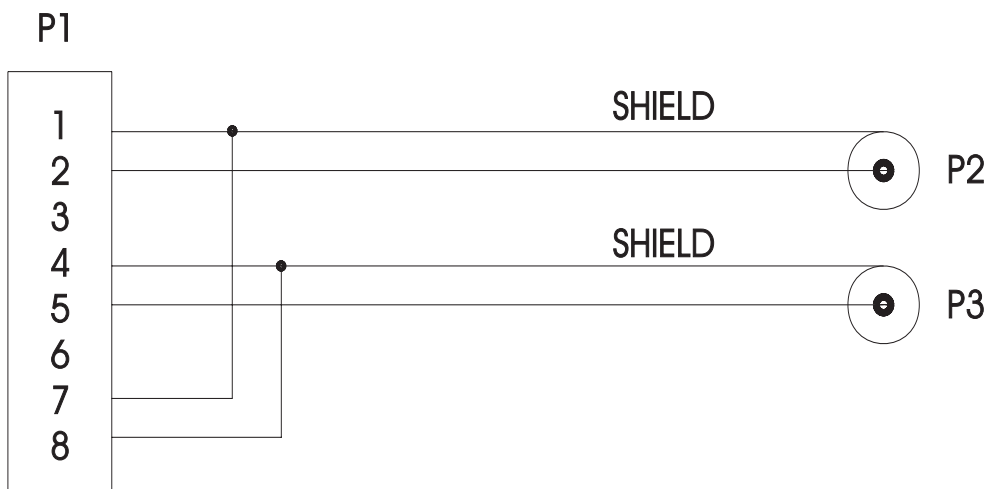
RJ-45 to BNC Cable Assembly (For E1 75Ω Port Connections)



The following parts are recommended for manufacturing the cable:

- For the switch side of the cable assembly (P1): 8-conductor RJ-45 round connector (MTP-88U or equivalent)
- For the cable: RG-187A coaxial cable (Belden 83267 or equivalent)
- For the customer end of the cable assembly (P2 and P3): Coaxial BNC connector, 75Ω (Amp 413760-8, or equivalent).

The figure below shows the pinouts for the cable assembly.



Index

! commands 4-26
+ or - commands 17-7
? command 4-16, 4-25
10/100 16-1, 16-8
10/100 command 15-4
10/100 ports 15-5, 15-8
10/100cfg command 3-15, 15-5, 15-7
10/100vc command 15-8
802.1Q 15-1, 16-1
802.2 pass through 18-17, 19-32

A

ab command 7-8
access rate 28-16, 28-22, 28-32, 28-35
actfstps command 17-38, 17-39
addprtchnl command 15-13
addvp command 19-5, 19-17, 19-28, 19-44, 19-60
 translations 18-17
adjacency
 definition for XMAP 21-2
admin login 4-33, 8-2
aipxsr command 27-12
aisr command 25-17
alert command 4-31
any-to-any switching 18-1
ARP protocol 25-3
ASCII terminal commands 4-17
at command 19-17
atvl command 20-23, 22-27
Authenticated Groups 19-1
 as mobile group 19-5
 configuring 19-27
Authentication 19-1
Authentication Management Console (AMC)
 software 19-5
autoencaps command 18-40
auto-switch
 diagram 19-30
 timer 19-30
Auto-Switch bridge mode 19-30
AutoTracker
 application examples 24-1
 configuring policies 20-4
 DHCP MAC address policy 20-3

DHCP policy 20-29
DHCP policy example 20-27
DHCP port policy 20-3
MAC address policy 20-2
menu 22-2
network address policy 20-2
policies 20-1, 22-3
port binding policy 20-2
port policy 20-2
protocol policy 20-2
user-defined policy 20-2
AutoTracker policies
 in mobile group 19-34
avlbootpmode command 26-2

B

backbones 15-9
 Ethernet 15-9
backplane threshold 11-4
Banyan Vines
 translations for 18-13
 VLAN for 22-31
BECN 29-11
boot configuration A-1
Boot prompt A-1
 basic commands A-3
 logging in A-2
 MPX A-7
BOOTP relay 26-4
 and authentication 26-5
br command 17-4, 19-17
BRI. See ISDN, Basic Rate Interface
bridge 17-7
bridge forwarding table 17-8
bridge mode 19-29
 non-Ethernet default 19-29
 optimized 19-29
 Spanning Tree 19-29
bridge port
 statistics 17-14
bridging
 Frame Relay 29-14, 29-57
bsadd command 34-3
bsdelete command 34-11
bsmodify command 34-9

C

- cacheconfig** command 9-33
- CAM
 - configuring 9-25
- CAM threshold 11-5
- camcfg** command 9-25
- camstat** command 9-24
- caplog** command 10-11
- cas** command 16-8, 16-11
 - Frame Relay bridging 29-57
 - Frame Relay routing 29-60
 - Frame Relay trunking 29-62
- cats** command 19-2
- cb** command 7-7
- cd** command 7-2, 19-62
- chassis
 - grounding 1-21
- chngmac** command 14-6
- chnlinfo** command 15-14
- Class B 1-7
- clearstat** command 9-16
- cmdlog** command 10-9
- command families 8-13, 8-17
- command history 4-26
- Command Line Interface 4-1
- Committed Information Rate (CIR)
 - Frame Relay 29-8
- communications
 - see also serial port
- configsyc** command 6-16
- configuration file
 - switch fails to create 35-10
- configuration files 7-2
- conlog** command 10-10
- console port
 - configuring 2-6, 6-2
 - speed 6-2
- consumable resources 11-2
- cp** command 7-6
- CPU threshold 11-5
- cratvl** command 22-4, 22-16, 22-31
- crechnl** command 15-11
- credit
 - transmit 19-31
- crgp** command 19-5, 19-19, 22-4
 - translations 18-17
- crmvcvl** command 23-5

D

- das** command 16-18
 - Frame Relay 29-66
- date 9-8
- Daylight Savings Time (DST) 9-8, 9-12
- db** command 7-9
- dbmap** command 17-19
- DC power supplies 1-24, 1-27
- debuglog** command 10-13
- def_group variable 19-12
- default group 19-12
- default VLAN 22-4, 22-5
 - routing 24-8
- defvl** command 22-5
- delechnl** command 15-13
- delprtchnl** command 15-14
- DES 12-4
- DHCP 26-4
 - and source routing 26-5
 - application example 20-27
 - overview 26-4
 - policies 20-3
 - with non-DHCP clients 26-8
- DHCP client 20-28
- DHCP server 20-28
- diag login 4-33, 8-2
- diag** user login 36-2, 36-3
- diagnostic tests 3-9
 - chassis 36-20
 - frame fabric test 36-18
- diagnostics
 - configuring 36-15, 36-16, 36-20
 - hardware 36-1
 - login 36-2, 36-3
 - running 36-2, 36-6
- diagnostics sub-menu 4-16, 36-3
- digital services 33-2
- Discard Eligibility
 - Frame Relay 29-9
- displaying Ethernet switch statistics 18-31
- DLCI 29-6
- DLCMI 29-24
- domain bridging
 - mapping table 17-19
- domain name servers 14-1
- DST 9-8
- dt** command 9-8
- duplicate MAC addresses 9-23
- dynamic port assignment 19-2

E

E1
 configuring 31 timeslots 28-45
 fractional 28-6
 framing 33-2

eb command 7-9

echo command 4-31

edit command 7-7

edit commands tutorial 7-11, 7-12

encapsulation 18-6
 IP 18-6
 IPX 18-6, 18-8

error messages 35-11

ESX-K-100C-32W 3-15, 15-5, 15-8

ESX-K-100FM/FS-16W 3-19, 15-8

eth100 command 15-4

ethdef command 18-26

Ethernet
 10/100 ports 15-5
 802.1Q 15-1
 auto-sensing ports 15-5
 auto-switch bridge mode 19-30
 backbones 15-9
 default translations 18-18, 18-26
 duplex mode 15-8
 Ethertype translation 18-20
 high-density ports 15-7
 link mode configuration 15-7
 LLC translation 18-23
 OmniChannel 15-1
 path MTU discovery 25-42
 port mapping 19-66
 port mirroring 19-57
 port monitoring 19-61
 SNAP translation 18-21

Ethernet 10/100 ports 15-8
 auto-negotiation 15-8
 configuring 15-5

Ethernet management port
 configuring 6-5
 MPX 2-5
 redundancy 6-7

Ethernet modules 3-15, 3-19
 configuring tests 36-17
 displaying switch statistics 18-31
 ESX-K-100C-32W 3-15
 ESX-K-100FM/FS-16W 3-19
 optimized ports 15-3
 pinouts 3-15

port partitioning 15-3
 three generations 15-2

ethernetc command 6-5

event 14-3

events command 14-5

F

facility datalink 33-6

Fast Ethernet 16-1
 configuring 15-8

Fast Spanning Tree
 description 17-34
 displaying port parameters 17-36
 enabling port parameters 17-38, 17-39

FDDI
 default translations 18-18, 18-27
 LLC translation 18-23
 SNAP translation 18-21

FDDI raw 18-27

fddidef command 18-27

FECN 29-12

fiber optic
 proper handling of cables 3-10

file command 4-15, 7-1

files 2-7
 configuration 7-2
 flash memory 7-3
 image 2-7

Filter Command. See UI Table Filtering

filter points 8-7

flash memory 2-7

flc command 17-21

flood limits 19-31
 configuring 17-21
 displaying 17-22

fls command 17-22

fping command 25-23

fr command 29-21

fradd command 29-32

frame flooding 22-15

Frame Relay
 back-to-back configurations 29-3
 BECN 29-11
 bridging 29-14, 29-57
 cables 29-4, B-1
 CIR 29-8
 compression 29-5
 congestion control 29-8

- control signals 29-46
- Discard Eligibility 29-9
- DLCI 29-6
- DLCMI 29-24
- errors 29-48
- FECN 29-12
- IP routing 29-15
- IPX routing 29-18
- polling 29-25
- port configuration 29-22
- Routing Group 29-59
- self-configuration 29-7
- split clocking 29-40
- statistics 29-38
- translations 29-13
- trunking 29-19, 29-62
- viewing parameters 29-33
- Virtual Circuit 29-6, 29-14
- Virtual Circuit configuration 29-32
- virtual ports 29-14
- Frame Relay boards
 - ipxsap** command with, 27-11
- framefab** command 36-2, 36-18
- frclear** command 29-54
- frdelete** command 29-36
- frmodify** command 29-22
- frstatus** command 29-38
- frview** command 29-33
- fsck** command 7-14, 9-21
- fstps** command 17-36
- FTP
 - commands 5-3
- FTP client 5-1
- ftp** command 5-3
- FTP commands
 - ? 5-3
 - ascii 5-3
 - binary 5-3
 - bye 5-3
 - cd 5-3
 - delete 5-3
 - dir 5-3
 - get 5-3
 - hash 5-3
 - lpwd 5-3
 - ls 5-3
 - put 5-3
 - pwd 5-3
 - quit 5-3
 - remotehelp 5-3

- user 5-3
- FTP servers 5-1, 5-2
- fwtlv** command 20-26, 22-30

G

- generic service relays 26-19
- Gigabit 16-1, 16-11
- Gigabit Ethernet modules 3-12
 - GSX-K-FM/FS/FH-2W 3-12
- global commands 8-18
- global commands** 8-18
- GMAP
 - configuring 21-11
 - gap time 21-11
 - update time 21-12
- gmapgaptime** command 21-11
- gmapholdtime** command 21-12
- gmaps** command 21-13
- gmapst** command 21-11
- gmapupdttime** command 21-12
- gmcfg** command 19-12, 22-2
- gmstat** command 22-2
- gp** command 19-17, 19-38
- Group 19-1
 - authenticated 19-1, 19-27
 - changing parameters 19-40
 - creating 19-18
 - deleting 19-43
 - flood limits 17-22
 - IP address 19-21
 - IP routing in 19-21
 - mobile 19-1, 19-5, 19-27
 - non-mobile 19-1, 19-15
 - port assignment to 19-2
 - viewing 19-38
 - WAN routing 19-19, 19-35
- Group Mobility 19-2
 - enabling switch-wide 19-12
- GSX-K-FM/FS/FH-2W 3-12

H

- hardware diagnostics 36-1
- hdcfg** command 11-2
- hdstat** command 11-6, 11-7
- Health MIB 11-1
 - management menu 11-1
 - resource thresholds 11-2

Hello messages
 and XMAP 21-2

help command 4-15

history command 4-26

hmstat command 11-7

hot swapping 1-11, 1-14, 3-7

hpstat command 11-8

hreset command 11-8

HRE-X 1-22
 router register limitations 1-23
 valid configurations 1-23

hrex command 9-27

hrexassign command 9-27

hrexdisplay command 9-27

hrexhashopt command 9-29

hrexutil command 9-29

I

ib command 7-9

ICMP protocol 25-3

ICMP statistics and errors 25-20

icmps command 25-20

image files 2-7, 3-3, 7-3

imgcl command 7-5

imgsync command 6-16

info command 9-6

interface command 4-15

Interswitch Protocols (XIP) 21-1
 submenu 21-1

Inverse ARP 29-15

IP
 abbreviated address format 4-28
 address 5-2
 BOOTP relay 26-4
 DHCP 26-4
 framing type 19-23
 problems with 35-7
 RIP mode 19-22

IP address
 changing in group 19-41

ip command 25-7

IP protocol 25-3

IP RIP Filters
 adding "global" filters 25-33
 adding specific filters 25-34
 configuring 25-33
 deleting filters 25-36
 displaying all filters 25-37

displaying global filters 25-38
 displaying specific filters 25-38
 filter precedence 25-35

IP Routing 19-21
 adding an IP address to ARP table 25-9
 adding IP static routes 25-17
 Address Resolution Protocol 25-3
 flushing the RIP Routing Table 25-32
 Internet Control Message Protocol 25-3
 Internet Protocol 25-3
 Open Shortest Path First Protocol 25-2
 PING command 25-22
 PINGing a host 25-22
 removing IP static routes 25-19
 Routing Information Protocol 25-2
 Simple Network Management Protocol
 25-3
 TELNET protocol 25-3
 tracing an IP route 25-31
 Transmission Control Protocol 25-3
 User Datagram Protocol 25-3
 viewing ICMP statistics and errors 25-20
 viewing IP statistics and errors 25-12
 viewing RIP statistics and errors 25-26
 viewing TCP statistics 25-27
 viewing the Address Translation Table
 25-8
 viewing the IP routing table 25-15
 viewing the IP-to-MAC Table 25-39
 viewing the TCP Connection Table 25-29
 viewing the UDP listener table 25-25
 viewing UDP statistics and errors 25-24

ipdirbcast command 25-41

ipf command 25-37

ipmac command 25-39

ipr command 25-15

ips command 25-12

IP-to-MAC Table
 displaying all entries 25-39
 displaying specific entries 25-40
 flushing entries 25-40

IPX
 address mapping 18-10
 Token Ring 18-15
 triggered RIP and SAP 19-37
 VLANs 24-4, 24-7

ipx command 27-4

IPX RIP
 description of protocol 27-2

- IPX RIP/SAP Filtering
 - adding global filters 27-19
 - adding specific filters 27-20
 - configuring NetWare for WAN links 27-33
 - default setting of filters 27-18
 - deleting filters 27-22
 - displaying all filters 27-23
 - displaying global filters 27-24
 - displaying specific filters 27-24
 - filter precedence 27-25
 - uses for filters 27-18
 - IPX routing
 - adding an IPX static route 27-12
 - configuring IPX Serialization Packet Filtering 27-26
 - configuring IPX Watchdog Spoofing 27-28
 - configuring NetWare for WAN links 27-33
 - configuring SPX Keepalive Spoofing 27-30
 - disabling IPX Router Complex 27-14
 - displaying IPX Routing Table 27-5
 - enabling IPX Router Complex 27-14
 - flushing RIP/SAP tables 27-15
 - GNS Output filters 27-18
 - PINGing an IPX node 27-16
 - removing IPX static routes 27-13
 - RIP Input filters 27-18
 - RIP Output filters 27-18
 - RIP/SAP Filters
 - configuring 27-18
 - SAP Input filters 27-18
 - SAP Output filters 27-18
 - the IPX submenu 27-4
 - viewing IPX statistics 27-8
 - viewing SAP Bindery 27-10
 - IPX Serialization Packet Filtering
 - configuring 27-26
 - IPX static routes
 - removing 27-13
 - IPX Watchdog Spoofing
 - configuring 27-28
 - ipxdrt** command 27-38
 - ipxext** command 27-37
 - ipxf** command 27-23
 - ipxfilter** command 27-19
 - ipxflush** command 27-15
 - ipxoff** command 27-14
 - ipxping** command 27-16
 - ipxr** command 27-5
 - ipxs** command 27-8
 - ipxsap** command 27-10
 - ipxserialf** command 27-26
 - ipxspooft** 27-28
 - ipxt** command 27-36
 - ipxtimer** command 27-35
- ISDN
- Basic Rate Interface (BRI) 28-6
- isdn** command 32-3
- ISDN Ports
- accessing the ISDN menu 32-3
 - deleting an ISDN configuration entry 32-5
 - displaying ISDN configuration entry status 32-7
 - modifying an ISDN configuration entry 32-4
 - viewing an ISDN configuration entry 32-6
- isdnd** command 32-5
- isdnm** command 32-4
- isdns** command 32-7
- isdnv** command 32-6
- ## K
- kill** command 4-35
- Kodiak Ethernet Modules 15-3
- ## L
- LAN Emulation (LANE)
- default translations 18-19
 - Ethertype translation 18-20, 18-24
 - SNAP translation 18-22
- lb** command 7-8
- leak monitor 9-19
- leakdumpall** command 9-19
- leakstart** command 9-19
- learning 18-41
- LEDs
- amber 35-9, 35-10
 - OK2 35-9
 - STA 35-10
 - TEMP 35-10
- Line Interface Unit (LIU) 33-5, 33-9
- link** command 31-2
- linkadd** command 31-3, 31-4
- linkdelete** command 31-11
- linkmodify** command 31-9
- linkstatus** command 31-15
- linkview** command 31-12
- LLC 18-7, 18-23

load command 5-4
 login accounts 4-33, 8-2
 login alert banner 4-31
logout command 4-16
lookup command 4-25
ls command 5-4, 7-3

M

MAC 17-16, 17-17
 MAC addresses
 configuring 14-6
 restoring 14-6
 MAC devices
 VLAN membership 20-26, 22-30
 main menu 4-15
mas command 16-12
 Frame Relay 29-65
maskta command 36-5
mcvl command 23-13
 MD5 12-4
 media access control - see MAC
 memory management 9-20
 memory threshold 11-5
 memory utilization statistics 9-19
memstat command 9-20
 Mobile Groups 19-1, 19-5
 aging out devices 19-12
 AutoTracker policies 19-34
 configuring 19-27
 def_group 19-12
 default group 19-12
 dynamic port assignment 19-5
 Ethernet and Token Ring ports 19-2
 move_from_def 19-13
 move_to_def 19-12
 multiple 19-2
 policies 20-1, 20-24
 ports in 19-5
 primary group 19-13
 static port assignment 19-5
 viewing 20-23
modatvl command 20-6, 20-7, 20-21, 20-22,
 22-24
 modem port 6-3
modmcvl command 23-9, 23-12
 modules
 removing 3-6
modvl command 19-22, 19-40, 22-4, 22-19

modvp command 9-32, 19-45, 19-60
 translations 18-17
 move_from_def variable 19-13
 set in mpx.cmd 19-13
 move_to_def variable 19-12
mpm command 6-9
mpmget command 6-17
mpmload command 6-12
mpmreplace command 6-12
mpmrm command 6-13
mpmstore command 6-11
 MPX 1-10, 1-13, 2-7
 Boot prompt configuration A-7
 configuring 6-1
 file corruption 2-2
 hot swap warning 2-2
 LEDs 35-9
 OK1 LED 2-2
 OK2 LED 2-2
 power down warning 2-2
 redundancy 2-9, 6-9
 resetting a secondary 6-19
 MTU 25-42
 multicast claiming 9-32
 Multicast VLANs 22-1, 23-1
 creating 23-4
 deleting 23-11
 device assignment in 23-2
 frame flooding in 23-3
 modifying 23-9
 multicast addresses 23-6
 multicast claiming, compared 23-2
 policies 23-12
 recipients 23-1, 23-7
 viewing 23-13
 multiple spanning tree 16-4
 multiple user sessions 4-33

N

names command 14-1
nb command 7-10
net command 34-2
 NetBIOS relays 26-11
 network address policy 22-7
 precedence 20-2
 Network Time Protocol 12-1
networking command 4-15, 25-6
newfs command 7-15, 9-22

- nisuf** command 6-14
- noecho** command 4-31
- non-Ethernet ports 19-29
- Non-mobile Groups 19-1, 19-15
- ntaccess** command 12-5, 12-36
- ntadmin** command 12-5, 12-33
- ntconfig** command 12-5
- ntinfo** command 12-5, 12-15
- NTP
 - advertised precision 12-14
 - client/server 12-8
 - client/server authentication 12-9
 - current leap second 12-30
 - event timer subsystem 12-28
 - I/O subsystem 12-27
 - key ID 12-34
 - key type 12-35
 - list of peers 12-15
 - local server information 12-21
 - local server statistics 12-23
 - loop filter information 12-26
 - packet count statistics 12-29
 - peer associations 12-12, 12-14
 - peer memory usage 12-26
 - peer summary information 12-16
 - primary receive timeout 12-33
 - reset subsystem counters 12-28
 - server statistics 12-24
 - specify password 12-34
 - system flag 12-35
 - trusted list 12-37
 - version number 12-20
- ntpaddpeer** command 12-12
- ntpaddsrv** command 12-13
- ntpauth** command 12-38
- ntpbcast** command 12-13
- ntpckey** command 12-37
- ntpcres** command 12-39
- ntpctlk** command 12-37
- ntpctlstat** command 12-29
- ntpctrap** command 12-41
- ntpdelay** command 12-33
- ntpdisable** command 12-35
- ntpdkey** command 12-38
- ntpdres** command 12-41
- ntpdtrap** command 12-42
- ntpenable** command 12-35
- ntpiconfig** command 12-6, 12-8
- ntpinfo** command 12-21
- ntpio** command 12-27

- ntpkeyid** command 12-34
- ntpkeytype** command 12-35
- ntpleap** command 12-30
- ntploop** command 12-26
- ntplpeers** command 12-15
- ntpmem** command 12-26
- ntpmlist** command 12-31
- ntpmon** command 12-31
- ntpmres** command 12-41
- ntppasswd** command 12-34
- ntppeers** command 12-16
- ntpprec** command 12-14
- ntppreset** command 12-28
- ntppstat** command 12-24
- ntpreqk** command 12-36
- ntpreset** command 12-28
- ntpshowpeer** command 12-18
- ntpstat** command 12-23
- ntptimeo** command 12-33
- ntptimer** command 12-28
- ntpunconfig** command 12-14
- ntpvrs** command 12-20
- ntpvkey** command 12-37
- ntpvres** command 12-40
- ntpvtrap** command 12-42
- ntstats** command 12-5, 12-23

O

- OK2 LED 35-9
- Omni Switch/Router 1-1, 2-7
 - HRE-X 1-22
 - see also OmniS/R
- OmniChannel 15-1, 15-9
 - creating 15-11
 - Ethernet 15-9
 - ports 15-13
 - primary/secondary ports 15-14
- OmniS/R 1-1
 - DHCP 26-4
- OmniS/R-3 1-8
- OmniS/R-5 1-10
- OmniS/R-9 1-13
- OmniS/R-9P 1-13
- optimized bridge mode 19-29
- Optimized Ports
 - ESX-K Series Modules 15-3
- OSPF protocol 25-2
- output translations 18-42

P

- partition management 8-11, 8-19
- password 4-33, 8-2
 - changing 8-2
- path MTU discovery 25-42
- ping** command 25-22
- pinouts
 - Ethernet modules 3-15
 - WAN modules 3-22
- pmapcr** command 19-17
- pmapdel** command 19-17
- pmapmod** command 19-17
- pmapv** command 19-17
- pmcfg** command 19-62
- pmdelete** command 19-64
- pmon** command 19-63
- pmpause** command 19-64
- pmstat** command 19-65
- PMTU 25-42
- Point-to-Point Protocol (PPP)
 - accessing the PPP menu 30-6
 - adding a PPP entity 30-9
 - deleting a PPP entity 30-21
 - displaying PPP entity status 30-18
 - modifying a PPP entity 30-15
 - setting global parameters 30-7
 - viewing PPP entity configurations 30-16
- policies
 - AutoTracker 20-1
 - configuring 20-4
 - DHCP 20-29
 - DHCP example 20-27
 - DHCP MAC address 20-3
 - DHCP port 20-3
 - IP 22-7
 - IPX 22-7
 - MAC address 20-2
 - network address 20-2, 22-7
 - port 20-2, 22-9
 - port binding 20-2
 - protocol 20-2
 - user-defined 20-2
- port mapping 19-17
 - example 19-66
 - operation of 19-67
 - relationship to policies 19-66
 - subset of ports 19-67
- port mirroring 19-33, 19-57
 - disabling 19-60
 - enabling 19-60
 - operation of 19-57
 - RMON probe 19-58
- port monitoring 19-61
 - menu 19-61
 - starting a session 19-63
 - statistics 19-65
- port monitoring resources 19-62
- Port partitioning
 - Ethernet modules 15-3
- port policies 22-9
 - backbone connections 22-12
 - inactive VLANs 22-12
 - silent stations 22-12
 - usefulness 22-12
 - with VAP 21-9
- port tests 36-9
 - cabling 36-22
 - cabling requirements 36-10
- ports 19-50
 - assignment to Group 19-2
 - Frame Relay 29-1
 - information 19-50
 - optimized (Ethernet modules) 15-3
 - spanning tree 17-30
 - statistics 19-53
 - translations 18-30
- power supply 1-8, 1-11, 1-14
 - connecting a DC power source 1-24, 1-27
 - replacing (9-slot chassis) 1-30
- ppp** command 30-6
- pppadd** command 30-9
- pppdelete** command 30-21
- pppglobal** command 30-7
- pppmodify** command 30-15
- pppstatus** command 30-18
- pppview** command 30-16
- probes 14-3
- probes** command 14-4
- prty_disp** command 19-17
- prty_mod** command 19-17
- pw** command 4-25, 8-2
- pwd** command 7-2

Q

- quit** command 4-16

R

rb command 7-8
 reboot 8-3
 - see also boot
reboot command 8-3
 receive threshold 11-3
 redundancy
 MPX 1-7, 1-10, 1-13, 2-9, 6-9
 power supply 1-11, 1-14
 re-executing commands 4-26
 reg_port_rule variable 22-9
relayc command 26-2, 26-3
relays command 26-23
 remote trunking stations 17-18
renounce command 6-14
res command 14-1
reset command 36-4
 resource thresholds 11-2
 RIF stripping
 and UDP relay 26-1
 RIP protocol 25-2
ripflush command 25-32
rips command 25-26
ripxsr command 27-13
risr command 25-19
rm command 2-8, 5-4, 7-4
rmatvl command 22-4, 22-26
rmgp command 19-43
rmmcvl command 23-11
 RMON 14-3
rmvp command 19-46
 routing 18-1
 default VLAN 24-8
 Frame Relay 29-59

S

sampling interval 11-6
 SAP
 description of protocol 27-2
 SAP Bindery
 viewing the, 27-10
saveconfig command 9-33
 SC connectors
 proper handling 3-10
 Search Command. See UI Table Filtering
secapply command 8-7
secdefine command 8-4
seclog command 8-10, 10-13

secreset command 6-19
 security 8-1
security command 4-15
selgp command 17-7
ser command 5-4, 6-2, 6-3
services command 4-15
 single spanning tree 16-4
 SLIP 6-3
slot command 2-9, 9-14
 slot table 9-14
sls command 6-11
 SNAP 18-7, 18-21
 SNMP protocol 25-3
 SNMP statistics 13-8
 SNMP traps 13-8
snmpc command 13-2
snmps command 13-8
 software
 installation problems 35-5
 switch 2-7
 source routing
 and DHCP 26-5
 Spanning Tree 17-28, 19-29
 non-mobile group 19-16
 parameters 17-25, 17-28, 17-32
 ports 17-30
 Speedy Tree Protocol
 description 17-35
 split clocking 28-18, 28-23
 SPX
 description of protocol 27-2
 SPX Keepalive Spoofing
 configuring 27-30
 SPX-Packet tolerance counting 27-30
spxspoo 27-30
 STA LED 35-10
 static bridge address 17-13
 configuring 17-10
 static port assignment 19-2
 static routes
 adding IP 25-17
 removing IP 25-19
 statistics
 module level 11-7
 port level 11-8
 port monitoring 19-65
 resetting 11-8
 switch level 11-6
 STATUS ENQUIRIES
 Frame Relay 29-7, 29-25, 29-26

- stc** command 17-25, 17-38, 17-39
 - sts** command 17-28, 17-38, 17-39
 - summary** command 4-15, 9-1
 - swap** command 6-20
 - swch** command 18-30, 18-31
 - switch
 - software 2-7
 - switch** command 4-15, 18-25
 - switch menu 18-25
 - switch software
 - Boot prompt 5-5
 - loading with FTP 5-2
 - loading with ZMODEM 5-4
 - switching modules 2-7, 3-12, 3-15, 3-19, 3-22
 - disabling 36-4
 - hot swapping 3-7
 - power consumption 1-19, 1-20
 - removing 3-6
 - resetting 36-4
 - swlogc** command 10-6
 - syncctl** command 6-15
 - syscfg** command 9-2, 9-23
 - syslog** command 10-2
 - sysstat** command 9-15
 - system boot A-2
 - commands 5-6
 - see also boot
 - system** command 4-15, 9-5
 - system description 9-23
 - system info 7-13
 - system menu 7-13, 9-5
 - system prompt 4-18
 - system statistics 9-15
- T**
- T1
 - fractional 28-6
 - framing 33-2
 - T1/E1 menu 33-3
 - T1/E1 ports 33-1
 - alarms 33-11
 - configuring 31 timeslots on a WAN E1 port 28-45
 - Extended Superframe 33-2, 33-4
 - facility datalink 33-6
 - line coding 33-5, 33-9
 - Line Interface Unit (LIU) 33-5, 33-9
 - loopback 33-7, 33-10
 - remote statistics 33-20
 - signaling 33-7, 33-10
 - statistics 33-17
 - Superframe 33-2, 33-4
 - yellow alarms 33-7, 33-14
 - Table Filtering. See UI Table Filtering
 - takeover** command 6-18
 - task utilization statistics 9-17
 - taskshow** command 9-17
 - taskstat** command 9-17
 - TCP protocol 25-3
 - tcpc** command 25-29
 - tcps** command 25-27
 - te** command 33-3
 - technical support 35-3
 - telcs** command 33-18
 - telis** command 33-19
 - TELNET command
 - using 25-30
 - telnet** command 25-30
 - telts** command 33-17
 - temod** command 33-4, 33-8
 - configuring a T1 port 33-4
 - configuring an E1 port 33-8
 - TEMP LED 35-10
 - temperature sensor 9-15
 - tercs** command 33-21
 - teris** command 33-21
 - terts** command 33-20
 - tes** command 33-11, 33-13, 33-15
 - viewing a T1 port 33-13
 - viewing an E1 port 33-15
 - test** command 36-6
 - testcfg** command 36-16, 36-20
 - testdisp** command 36-15
 - time 9-8
 - time slot 33-2
 - time zone 9-8
 - configuring 9-9
 - Token Ring
 - default translations 18-19, 18-28
 - traceroute** command 25-31
 - translations 18-1
 - ATM LANE 18-19
 - automatic 18-40
 - default options 18-17
 - Ethernet 18-18
 - Ethertype 18-20
 - FDDI 18-18
 - LLC 18-23

- SNAP 18-21
- Token Ring 18-19
- transmission states
 - XMAP 21-3
- transmit credit 19-31
- transmit/receive threshold 11-3
- traps
 - configuring 13-2
- trdef** command 18-28
- troubleshooting 35-1
- Truncating Tree Timing
 - description 17-35

U

- UDP 25-3
- UDP relay 26-1
- udpl** command 25-25
- udps** command 25-24
- UI Table Filtering 4-38
 - Filter Command 4-41
 - combining Search Command with 4-42
 - more** mode and 4-38
 - Search Command 4-39
 - combining Filter Command with 4-42
 - more** mode and 4-38
 - renewing a search 4-40
 - wildcards and 4-44
- uic** command 4-17-4-24, 4-30, 7-1, 9-5
- Universal Serial Port 28-6
- Universal Time Coordinate (UTC) 9-8
- User Interface 4-1, 4-16
- user login 4-33, 8-2
- useradd** command 8-12
- userdel** command 8-20
- usermod** command 8-16, 8-20
- userview** command 8-12
- USP. See Universal Serial Port.
- UTC 9-8

V

- VAP
 - configuring 21-9
 - databases 21-8
 - relation to port policies 21-9
- vas** command 16-16
 - Frame Relay 29-64

- ve** command 19-55
- verbose 4-22
- vi** command 18-24, 19-50
- via** command 19-47
- viatrl** command 20-24, 22-28
- view** command 7-6
- vigl** command 19-14
- vimcrl** command 23-14, 23-15
- viqs** command 16-17
- virtual ports 19-15
 - adding 19-44
 - deleting 19-46
 - errors 19-55
 - format 19-31
 - information on 19-47
 - modifying 19-45
 - statistics for 19-50, 19-53
 - VLAN membership 22-29
 - VLAN/group membership 20-25
- virtual router ports 22-19
 - creating 19-21
- vivl** command 20-25, 22-29
- VLAN Advertisement Protocol
 - see VAP
- vlan** command 4-15, 19-17
- VLAN policies
 - deleting 22-25, 23-10
 - modifying 22-26, 23-11
 - viewing 20-24, 22-28, 23-14
- VLANs 19-15, 22-1
 - application examples 24-1
 - backbone 24-10
 - Banyan Vines 22-31
 - bridges 17-4
 - creating 22-16
 - deadlocked 35-6
 - default 22-4
 - deleting 22-26
 - deleting policies in 22-25, 23-10
 - frame flooding in 22-15
 - IP routing in, 25-4
 - IPX networks 24-4, 24-7
 - logical policies 24-2
 - modifying 22-24
 - multicast 23-1
 - network address policies 24-2
 - policies 20-2, 22-3
 - port policy 22-9, 22-18
 - router ports 22-15
 - router traffic in 22-7

- secondary traffic 22-6
- translated frames 24-7
- viewing 20-23, 22-27
- vlap** command 21-9
- vs** command 19-53

W

- wan** command 28-14
- WAN Links
 - accessing the LINK menu 31-2
 - adding a link record 31-3
 - deleting link records 31-11
 - displaying link status 31-15
 - modifying a link record 31-9
 - viewing link records 31-12
- WAN modules 3-22
 - cables B-1
 - pinouts 3-22
 - WSX-BRI-SC 3-36
 - WSX-FT/E1-SC 33-1
 - WSX-FT1/E1-SC 3-32
 - WSX-S-2W 3-27
 - WSX-SC 3-29
- WAN routing 19-19, 19-35
- warning
 - hot swapping and file corruption 2-2
 - power down and file corruption 2-2
- wb** command 7-10
- who** command 4-34
- wpmodify** command 28-14
- wpstatus** command 28-38
- wpview** command 28-27
- write** command 4-35
- WSX
 - back-to-back configuration 28-8
 - back-to-back configurations 28-8
 - cables 28-11
 - data compression 28-12
 - port configuration 28-14
 - statistics 28-38
 - viewing parameters 28-27
- WSX-BRI-SC 3-36
- WSX-FT1/E1-SC 3-32, 33-1
- WSX-S-2W 3-27
- WSX-SC 3-29

X

- xlat** command 25-8
- XMAP
 - adjacency 21-2
 - and remote switches 21-4
 - common transmission time 21-7
 - configuring 21-5
 - discovery transmission time 21-6
 - transmission states 21-3
 - well-known MAC address 21-3
- xmapcmntime** command 21-7
- xmapdisctime** command 21-6
- xmaps** command 21-5
- xmapst** command 21-5

Z

- ZMODEM 5-1, 5-4, A-1

