



ES-4810
Management Module
Operations Guide

MANU0296-01 - Rev. A - March, 1998

Software Version 4.7.x

FORE Systems, Inc.

1000 FORE Drive
Warrendale, PA 15086-7502
Phone: 724-742-4444
FAX: 724-772-6500
URL: <http://www.fore.com>

Legal Notices

Copyright © 1995-1998 FORE Systems, Inc. All rights reserved.

U.S. Government Restricted Rights. If you are licensing the Software on behalf of the U.S. Government (“Government”), the following provisions apply to you. If the Software is supplied to the Department of Defense (“DoD”), it is classified as “Commercial Computer Software” under paragraph 252.227-7014 of the DoD Supplement to the Federal Acquisition Regulations (“DFARS”) (or any successor regulations) and the Government is acquiring only the license rights granted herein (the license rights customarily provided to non-Government users). If the Software is supplied to any unit or agency of the Government other than DoD, it is classified as “Restricted Computer Software” and the Government’s rights in the Software are defined in paragraph 52.227-19 of the Federal Acquisition Regulations (“FAR”) (or any successor regulations) or, in the cases of NASA, in paragraph 18.52.227-86 of the NASA Supplement to the FAR (or any successor regulations).

Printed in the USA.

No part of this work covered by copyright may be reproduced in any form. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

This publication is provided by FORE Systems, Inc. “as-is” without warranty of any kind, either express or implied, including, but not limited to, the implied warranties or conditions of merchantability or fitness for a particular purpose. FORE Systems, Inc. shall not be liable for any errors or omissions which may occur in this publication, nor for incidental or consequential damages of any kind resulting from the furnishing, performance, or use of this publication.

Information published here is current or planned as of the date of publication of this document. Because we are improving and adding features to our products continuously, the information in this document is subject to change without notice.

RESTRICTED RIGHTS LEGEND. Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 (October 1988) and FAR 52.227-19 (June 1987).

FCC CLASS A NOTICE

WARNING: Changes or modifications to this unit not expressly approved by the party responsible for compliance could void this user’s authority to operate this equipment.

NOTE: The ES-4810 has been tested and found to comply with the limits for a 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 instruction manual, may cause harmful interference to radio communications. Operation of the equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

DOC CLASS A NOTICE

This digital apparatus does not exceed Class A limits for radio noise emission for a digital device as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Le present appareil numerique n’emet pas de bruits radioelectriques depassant les limites applicables aux appareils numeriques de la class A prescrites dans le reglement sur le brouillage radioelectrique edicte par le ministere des Communications du Canada.

CE NOTICE

Marking by the symbol **CE** indicates compliance of this system to the EMC (Electromagnetic Compatibility) directive of the European Community and compliance to the Low Voltage (Safety) Directive. Such marking is indicative that this system meets or exceeds the following technical standards:

- EN 55022 - "Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment."
- EN 50082-1 - "Electromagnetic compatibility - Generic immunity standard Part 1: Residential, commercial, and light industry."

CERTIFICATIONS

ETL certified to meet Information Technology Equipment safety standards UL 1950 3rd Edition, CSA22.2, No. 950-95, and EN 60950.

TRADEMARKS

FORE Systems is a registered trademark, and *ForeRunner*, *ForeThought*, and *ForeView* are trademarks of FORE Systems, Inc. All other brands or product names are trademarks or registered trademarks of their respective holders.

Table of Contents

CHAPTER 1 Introduction

1.1	ES-4810 Description	1 - 1
1.1.1	ES-4810 Management Modules	1 - 2
1.1.1.1	NMM-1	1 - 3
1.1.1.2	NMM-2	1 - 4
1.1.1.3	NMM-SEG-1	1 - 5
1.1.2	ES-4810 Backplane	1 - 6
1.1.3	ES-4810 Ethernet Modules	1 - 7
1.1.3.1	CAM Capable Management Module	1 - 7
1.1.3.2	Address Learning	1 - 7
1.1.3.3	Secure Address Learning	1 - 8
1.1.3.4	Virtual LAN (VLAN) Support	1 - 8
1.1.3.5	Station Migration	1 - 8
1.1.3.6	Uplink and VLAN Uplink Capability	1 - 9
1.1.3.7	Statistics Counters	1 - 9
1.1.3.8	Switch Monitoring	1 - 9
1.2	Management Interfaces	1 - 10
1.2.1	MIB Support	1 - 10
1.2.2	Operator Console Interface	1 - 10
1.3	Ethernet Switch Concepts	1 - 11
1.3.1	What is a Switch?	1 - 11
1.3.2	Transparent Bridging	1 - 12
1.3.2.1	Learning Process	1 - 12
1.3.2.2	Forwarding Process	1 - 13
1.3.3	Virtual LANs (VLANs)	1 - 14
1.3.3.1	Using VLANs Within a Switch	1 - 14
1.3.3.2	Using VLANs Between Switches	1 - 16
1.3.4	Switch Port Forwarding Modes	1 - 17
1.3.4.1	Forwarding Broadcasts and Multicasts	1 - 17
1.3.4.2	Forwarding Unknown Unicasts	1 - 17
1.3.4.3	VLAN Tagging Between ES-4810s	1 - 18
1.3.4.4	NMM-SEG-1 Forwarding Modes	1 - 19
1.3.5	Avoiding Loops in Switch Configuration	1 - 20
1.3.5.1	Loops Within a Switch	1 - 20
1.3.5.2	Loops with Multiple Switches	1 - 21
1.3.5.3	Load Balancing with Parallel Cables	1 - 23

Table of Contents

1.3.5.4	VLAN Assignments Between Switches	1 - 24
1.3.5.5	Parallel Uplink Ports and Spanning Tree Protocol	1 - 24
1.3.5.6	Spanning Tree Protocol	1 - 24
CHAPTER 2 Management Module Configuration		
2.1	Overview	2 - 1
2.1.1	Online Help	2 - 2
2.1.2	Logging On and Off the Operator Console	2 - 6
2.1.3	Community String Configuration	2 - 6
2.1.4	Recovering from a Forgotten Login ID.	2 - 9
2.2	Configuring Management Interfaces	2 - 10
2.2.1	Displaying the Current Interface Configuration	2 - 10
2.2.2	Assigning the IP Address and Network Mask	2 - 11
2.2.3	Assigning Interfaces to Packet Buses	2 - 13
2.2.4	Displaying and Setting Packet Bus Memos	2 - 14
2.2.5	Configuring the Frame Acceptance Mode	2 - 15
2.3	Defining the Routing Table.	2 - 16
2.4	Configuring Traps.	2 - 20
2.4.1	Configuring the Network Manager Trap Table	2 - 20
2.5	Enabling CAM	2 - 22
2.6	Setting Up Spanning Tree	2 - 23
2.6.1	Enabling Spanning Tree	2 - 23
2.6.2	Setting Spanning Tree Parameters	2 - 23
2.6.2.1	Priority	2 - 24
2.6.2.2	Hello Time	2 - 24
2.6.2.3	Max Age	2 - 24
2.6.2.4	Forward Delay Time	2 - 25
2.7	Configuring SNMPv2 Security Features	2 - 26
2.7.1	Displaying the Non-volatile Party Database	2 - 27
2.7.2	Configuring Initial Security Parameters	2 - 28
2.7.3	Enabling and Disabling SNMPv1.	2 - 30
2.7.4	Resynchronizing SNMPv2 Agents and Management Stations	2 - 31
2.8	Rebooting or Resetting the Management Module	2 - 33
CHAPTER 3 User Module Configuration		
3.1	Overview	3 - 1
3.2	Common Configuration Functions	3 - 2
3.3	Slot and Port Numbering in a Chassis.	3 - 3
3.4	Configuring Groups	3 - 3
3.5	Assigning Groups to a Packet Bus	3 - 4
3.6	Enabling/Disabling Groups of Ports.	3 - 6

3.7	Resetting Groups Of Ports	3 - 7
3.8	Setting the Aging Time	3 - 8
3.9	Configuring Ethernet Switch Ports	3 - 9
3.10	Enabling/Disabling ports	3 - 11
3.11	Enabling/Disabling Spanning Tree on Ports	3 - 12
3.12	Modifying Pathcost on Ports	3 - 13
3.13	Modifying Spanning Tree Priority on Ports	3 - 14
3.14	Resetting Ports	3 - 15
3.15	Assigning Individual Ports To Groups	3 - 16
3.16	Specifying Port Memos	3 - 17
3.17	Setting Port Priority	3 - 18
3.18	Setting Redundant Port Pairs	3 - 19
3.19	Setting Port Monitoring Mode	3 - 21
3.20	Setting Duplex Mode	3 - 23
3.21	Setting 100BaseTX Port Speed	3 - 24
3.22	Modifying the Address Database For Security	3 - 25
	3.22.1 Displaying the Address Database	3 - 26
	3.22.2 Restricting Addresses in the Database	3 - 27
	3.22.3 Setting the Learning Mode	3 - 28
3.23	Modifying Forwarding of Unknown Unicasts on End-station Ports	3 - 30
3.24	Setting Long and Short History	3 - 31
3.25	Configuring VLANs	3 - 32
3.26	Configuring ATM Uplinks	3 - 33
3.27	Configuring Uplink Ports	3 - 35
3.28	Connecting to Another ES-4810 or non-ES-4810	3 - 36
3.29	Using VLAN Tagging Between ES-4810 Modules	3 - 37

CHAPTER 4 Network Statistics Procedures

4.1	Viewing System Variables	4 - 2
4.2	RMON Configuration	4 - 4
	4.2.1 Enabling or Disabling RMON Table Creation	4 - 5
	4.2.2 Configuring the Size of the History Tables	4 - 6
	4.2.3 Enabling RMON Statistics Collection	4 - 7
4.3	Viewing Ethernet Network Statistics	4 - 8
	4.3.1 Displaying Interface Statistics	4 - 8
4.4	Diagnostics	4 - 9
	4.4.1 Checking the Configuration	4 - 9
	4.4.2 Self-Test-Diagnostic Test	4 - 9

CHAPTER 5 Operator Console Command Reference

5.1	atmuplink	5 - 2
5.1.1	Format	5 - 2
5.1.2	Description	5 - 2
5.1.3	Options	5 - 2
5.1.4	Example	5 - 3
5.2	card	5 - 4
5.2.1	Format	5 - 4
5.2.2	Description	5 - 4
5.2.3	Options	5 - 5
5.2.4	Example	5 - 5
5.3	chassis	5 - 6
5.3.1	Format	5 - 6
5.3.2	Description	5 - 6
5.3.3	Example	5 - 6
5.3.4	See Also	5 - 6
5.4	check	5 - 7
5.4.1	Format	5 - 7
5.4.2	Description	5 - 7
5.4.3	Examples	5 - 7
5.5	community	5 - 9
5.5.1	Format	5 - 9
5.5.2	Description	5 - 9
5.5.3	Options	5 - 9
5.5.4	Example	5 - 10
5.5.5	See Also	5 - 11
5.6	esbpbus	5 - 12
5.6.1	Format	5 - 12
5.6.2	Description	5 - 12
5.6.3	Options	5 - 12
5.6.4	Example	5 - 12
5.7	escam	5 - 13
5.7.1	Format	5 - 13
5.7.2	Description	5 - 13
5.7.3	Options	5 - 14
5.7.4	Example	5 - 14
5.8	esgroup	5 - 15
5.8.1	Format	5 - 15
5.8.2	Description	5 - 15
5.8.3	Options	5 - 15

5.8.4	Example	5 - 16
5.9	espair	5 - 17
5.9.1	Format	5 - 17
5.9.2	Description	5 - 17
5.9.3	Options	5 - 17
5.9.4	Example	5 - 18
5.9.5	See Also	5 - 18
5.10	esport	5 - 19
5.10.1	Format	5 - 19
5.10.2	Description	5 - 20
5.10.3	Options	5 - 20
5.10.4	Example	5 - 24
5.11	group	5 - 25
5.11.1	Format	5 - 25
5.11.2	Description	5 - 25
5.11.3	Options	5 - 25
5.11.4	See Also	5 - 26
5.12	if	5 - 27
5.12.1	Format	5 - 27
5.12.2	Description	5 - 27
5.12.3	Options	5 - 28
5.13	port	5 - 29
5.13.1	Format	5 - 29
5.13.2	Description	5 - 29
5.13.3	Options	5 - 30
5.13.4	See Also	5 - 30
5.14	rmon	5 - 31
5.14.1	Format	5 - 31
5.14.2	Description	5 - 31
5.14.3	Example	5 - 32
5.14.4	See Also	5 - 32
5.15	rmon host	5 - 33
5.15.1	Format	5 - 33
5.15.2	Description	5 - 33
5.15.3	Example	5 - 34
5.15.4	See Also	5 - 35
5.16	rmon long	5 - 36
5.16.1	Synopsis	5 - 36
5.16.2	Description	5 - 36
5.16.3	Options	5 - 37
5.16.4	Example	5 - 37

Table of Contents

5.16.5	See Also	5 - 38
5.17	rmon macip	5 - 39
5.17.1	Format	5 - 39
5.17.2	Description	5 - 39
5.17.3	Example	5 - 40
5.17.4	See Also	5 - 41
5.18	rmon matrix	5 - 42
5.18.1	Format	5 - 42
5.18.2	Description	5 - 42
5.18.3	Example	5 - 43
5.18.4	See Also	5 - 44
5.19	rmon short	5 - 45
5.19.1	Format	5 - 45
5.19.2	Description	5 - 45
5.19.3	Options	5 - 46
5.19.4	Example	5 - 46
5.19.5	See Also	5 - 47
5.20	rmon stats	5 - 48
5.20.1	Format	5 - 48
5.20.2	Description	5 - 48
5.20.3	Example	5 - 49
5.20.4	See Also	5 - 50
5.21	route	5 - 51
5.21.1	Format	5 - 51
5.21.2	Description	5 - 51
5.21.3	Options	5 - 52
5.21.4	Example	5 - 53
5.22	setup	5 - 54
5.22.1	Format	5 - 54
5.22.2	Example	5 - 54
5.22.3	See Also	5 - 54
5.23	snmp	5 - 55
5.23.1	Format	5 - 55
5.23.2	Description	5 - 55
5.23.3	Options	5 - 56
5.23.4	Example	5 - 57
5.24	stbridge	5 - 58
5.24.1	Format	5 - 58
5.24.2	Description	5 - 58
5.24.3	Options	5 - 59
5.24.4	Example	5 - 61

5.25	stport	5 - 63
5.25.1	Format	5 - 63
5.25.2	Description	5 - 63
5.25.3	Options	5 - 64
5.25.4	Example	5 - 65
5.26	system	5 - 67
5.26.1	Format	5 - 67
5.26.2	Description	5 - 67
5.26.3	Options	5 - 67
5.26.4	Example	5 - 68
5.27	trap	5 - 69
5.27.1	Format	5 - 69
5.27.2	Description	5 - 69
5.27.3	Options	5 - 70
5.27.4	Example	5 - 70
5.28	version	5 - 71
5.28.1	Format	5 - 71
5.28.2	Example	5 - 71

CHAPTER 6 Upgrading the Management Module Firmware

6.1	Determining the Current Firmware Version	6 - 2
6.1.1	Using the Version Command	6 - 2
6.1.2	Using SNMP	6 - 2
6.2	Requirements for the Upgrade Process	6 - 3
6.3	Upgrading the Firmware	6 - 4

Table of Contents

List of Figures

CHAPTER 1 Introduction

Figure 1.1	Management Module Faceplate	1 - 2
Figure 1.2	NMM-1 Interface	1 - 3
Figure 1.3	NMM-2 Interfaces	1 - 4
Figure 1.4	NMM-SEG-1 Interfaces	1 - 5
Figure 1.5	Example of Using VLANs	1 - 15
Figure 1.6	Example of Using VLANs Across FORE Switches	1 - 16
Figure 1.7	Example of a Loop Within a Switch	1 - 20
Figure 1.8	Example of a Loop Between Two Switches	1 - 21
Figure 1.9	Example of a Loop Between Three Switches	1 - 22
Figure 1.10	VLANs Assignments Between Switches Using Parallel Cables	1 - 23
Figure 1.11	Example of Spanning Tree Preventing a Loop	1 - 25

CHAPTER 2 Management Module Configuration

Figure 2.1	The HELP Command on the Management Module (One of Two)	2 - 2
Figure 2.2	The HELP Command on the Management Module (Two of Two)	2 - 3
Figure 2.3	The Setup Command on the Management Module	2 - 4
Figure 2.4	The community help Command	2 - 5
Figure 2.5	Logging On to the Console	2 - 6
Figure 2.6	Displaying Community String Configuration	2 - 7
Figure 2.7	Ethernet Interface Configuration Display	2 - 10
Figure 2.8	Relationship of IP Address to Subnet Mask	2 - 12
Figure 2.9	Defining a Subnet Using Subnet Mask	2 - 12
Figure 2.10	The if show Command on the NMM-SEG-1	2 - 14
Figure 2.11	Defining and Displaying a Backplane Memo	2 - 14
Figure 2.12	Displaying the Routing Table	2 - 16
Figure 2.13	Displaying User-defined Routes in NVRAM	2 - 17
Figure 2.14	Adding Routes to the Routing Table	2 - 18
Figure 2.15	Displaying All Routes	2 - 19
Figure 2.16	Displaying the Trap Table	2 - 20
Figure 2.17	Help for SNMP Commands	2 - 26
Figure 2.18	Displaying the Non-volatile Party Database	2 - 27

List of Figures

Figure 2.19 Adding Initial Parties	2 - 28
Figure 2.20 Resetting a Party's Clock Value	2 - 32

CHAPTER 3 User Module Configuration

Figure 3.1 Slot Numbering in a Chassis	3 - 3
Figure 3.2 Work Groups on ES-4810 Modules	3 - 4
Figure 3.3 Displaying Current Group Configuration	3 - 5
Figure 3.4 Setting the Aging Time for a Group	3 - 8
Figure 3.5 Displaying the Current Ethernet Port Configuration	3 - 9
Figure 3.6 Enabling and Disabling Ports	3 - 11
Figure 3.7 Setting Redundant Port Pairs	3 - 20
Figure 3.8 Setting Port Monitoring Mode	3 - 22
Figure 3.9 Displaying the Address Database for a Port	3 - 26
Figure 3.10 Modifying the Address Database	3 - 28
Figure 3.11 Setting Secure Learning for a Port	3 - 29
Figure 3.12 Setting the Forwarding Mode for End-station Ports	3 - 30
Figure 3.13 Setting the History	3 - 31
Figure 3.14 Setting the VLAN Assignment for a Port	3 - 32

CHAPTER 4 Network Statistics Procedures

Figure 4.1 Chassis show and version Commands	4 - 2
Figure 4.2 System show Command	4 - 2
Figure 4.3 Displaying the Current RMON Configuration	4 - 4
Figure 4.4 Configuring the Size of the History Tables	4 - 7
Figure 4.5 Displaying Ethernet Interface Statistics	4 - 8

CHAPTER 5 Operator Console Command Reference

Figure 5.1 Example of atmuplink show Command	5 - 3
Figure 5.2 Example of card show Command	5 - 5
Figure 5.3 Example of chassis show Command	5 - 6
Figure 5.4 Example of check config Command	5 - 7
Figure 5.5 Example of check config Command	5 - 8
Figure 5.6 Example of community show Command	5 - 10
Figure 5.7 Example of esbpbus Command	5 - 12
Figure 5.8 Example of esc Command	5 - 14
Figure 5.9 Example of esgroup Command	5 - 16
Figure 5.10 Example of espair Command	5 - 18
Figure 5.11 Example of esport Command	5 - 24
Figure 5.12 Example of rmon Command	5 - 32

Figure 5.13 Example of rmon host Command	5 - 34
Figure 5.14 Example of rmon long Command	5 - 37
Figure 5.15 Example of rmon macip Command	5 - 40
Figure 5.16 Example of rmon matrix Command	5 - 43
Figure 5.17 Example of rmon short Command	5 - 46
Figure 5.18 Example of rmon stats Command	5 - 49
Figure 5.19 Example of route add Command	5 - 53
Figure 5.20 Example of setup Command	5 - 54
Figure 5.21 Example of snmp Command	5 - 57
Figure 5.22 Example of stbridge Command (one of two)	5 - 61
Figure 5.23 Example of stbridge Command (two of two)	5 - 62
Figure 5.24 Example of the stport Command (one of two)	5 - 65
Figure 5.25 Example of Command (two of two)	5 - 66
Figure 5.26 Example of the system Command	5 - 68
Figure 5.27 Example of trap Command	5 - 70
Figure 5.28 Example of version Command	5 - 71

List of Figures

Preface

This manual provides information about the ES-4810 management modules. The management modules are used to configure, manage, and control the ES-4810 Ethernet and uplink modules. If you have any questions or problems with the installation, please contact FORE Systems' Technical Assistance Center (TAC) using the information on page ii.

Chapter Summaries

Chapter 1 - Introduction - Describes the management modules and their features, as well as providing general information about Ethernet switching.

Chapter 2 - Management Module Configuration - Contains procedures for the initial set up and configuration of a management module.

Chapter 3 - User Module Configuration - Provides procedures for configuring user modules.

Chapter 4 - Network Statistics Procedures - Describes methods of performance monitoring using a management module.

Chapter 5 - Operator Console Command Reference - Provides an alphabetical reference of all operator console commands used to manage ES-4810 modules.

Chapter 6 - Upgrading the Management Module Firmware - Provides procedures for upgrading the firmware on an ES-4810 management module.

Technical Support

In the U.S.A., customers can reach FORE Systems' Technical Assistance Center (TAC) using any one of the following methods:

1. Select the "Support" link from FORE's World Wide Web page:
<http://www.fore.com/>
2. Send questions, via e-mail, to:
support@fore.com
3. Telephone questions to "support" at:
800-671-FORE (3673) or 724-742-6999
4. FAX questions to "support" at:
724-742-7900

Technical support for customers outside the United States should be handled through the local distributor or via telephone at the following number:

+1 724-742-6999

No matter which method is used to reach the TAC, customers should be ready to provide the following:

- A support contract ID number
- The serial number of each product in question
- All relevant information describing the problem or question

Applicable Documents

For more information about subjects related to the FORE Systems ES-4810, refer to the following documents:

Title	Reference Document
<i>ES-4810 Ethernet Module Operations Guide</i>	MANU0297
<i>ES-4810 ATM Uplink User's Manual</i>	MANU0294
<i>ES-4810 Chassis User's Manual</i>	MANU0295
Structure and Identification of Management Information for TCP/IP-based Internets	RFC 1155, May 1990
A Simple Network Management Protocol	RFC 1157, May 1990
Concise MIB Definitions	RFC 1212
Management Information Base of Network Management of TCP/IP-based Internets: MIB II	RFC 1213
Extensions to the Generic Interface MIB	RFC 1229
Introduction to SNMPv2	RFC 1441
Party MIB for SNMPv2	RFC 1447
MIB for SNMPv2	RFC 1450
Evolution of the Interfaces Group of MIB-II	RFC 1573
Definitions of Managed Objects for the Ethernet-like Interface Types	RFC 1643

Typographical Styles

Throughout this manual, all specific commands meant to be entered by the user appear on a separate line in bold typeface. In addition, use of the Enter or Return key is represented as **<ENTER>**. The following example demonstrates this convention:

```
cd /usr <ENTER>
```

File names that appear within the text of this manual are represented in the following style: “... refer to the `README.TXT` file on the CD...”

Command names and GUI control buttons that appear within the text of this manual are represented in the following style: “Choose the `start` button on the Taskbar.”

Parameter names that appear within the text of this manual are represented in the following style: “The |<range> is an optional part...”

Any messages that appear on the screen during software installation and network interface administration are shown in `Courier` font to distinguish them from the rest of the text as follows:

```
.... Are all four conditions true?
```

Important Information Indicators

To call your attention to safety and otherwise important information that must be reviewed to insure correct and complete installation, as well as to avoid damage to the FORE adapter or your system, FORE Systems utilizes the following *WARNING/CAUTION/NOTE* indicators.

WARNING statements contain information that is critical to the safety of the operator and/or the system. Do not proceed beyond a *WARNING* statement until the indicated conditions are fully understood or met. This information could prevent serious injury to the operator and damage to the FORE adapter, the system, or currently loaded software, and will be indicated as:

WARNING!



Hazardous voltages are present. To lessen the risk of electrical shock and danger to personal health, follow the instructions carefully.

Information contained in *CAUTION* statements is important for proper installation/operation. Compliance with *CAUTION* statements can prevent possible equipment damage and/or loss of data and will be indicated as:

CAUTION



You risk damaging your equipment and/or software if you do not follow these instructions.

Information contained in *NOTE* statements has been found important enough to be called to the special attention of the operator and will be set off from the text as follows:



Steps 1, 3, and 5 are similar to the installation for the computer type above. Review the previous installation procedure before installation in your particular model.

Safety Agency Compliance

This preface provides safety precautions to follow when installing a FORE Systems, Inc., product.

Safety Precautions

For your protection, observe the following safety precaution when setting up your equipment:

- Follow all warnings and instructions marked on the equipment.

Symbols

The following symbols appear in this book.

CAUTION



If instructions are not followed, there is a risk of damage to the equipment.

WARNING!



Hazardous voltages are present. If the instructions are not heeded, there is a risk of electrical shock and danger to personal health.

Modifications to Equipment

Do not make mechanical or electrical modifications to the equipment. FORE Systems, Inc., is not responsible for regulatory compliance of a modified FORE product.

This document describes the following FORE Systems ES-4810 management modules:

- NMM-1
- NMM-2
- NMM-SEG-1

The management modules provide the ability to configure and manage the ES-4810 Ethernet modules.

1.1 ES-4810 Description

The FORE Systems ES-4810 provides wire speed per-port Ethernet switching and ATM uplink capabilities.

An ES-4810 system has the following main components:

- A 12-slot chassis with a switched Ethernet backplane.
- One or more Ethernet modules, allowing a wide variety of full-duplex 10/100 Mbps fiber and copper connections.
- An ATM Uplink module, providing single or dual 155 Mbps connections to an ATM backbone network.
- One or more management modules, allowing you to configure Ethernet and ATM uplink modules that are installed in the ES-4810.

1.1.1 ES-4810 Management Modules

There are several different management modules that can be used to manage ES-4810 Ethernet and ATM uplink modules. Each management module is designed to meet a different need.

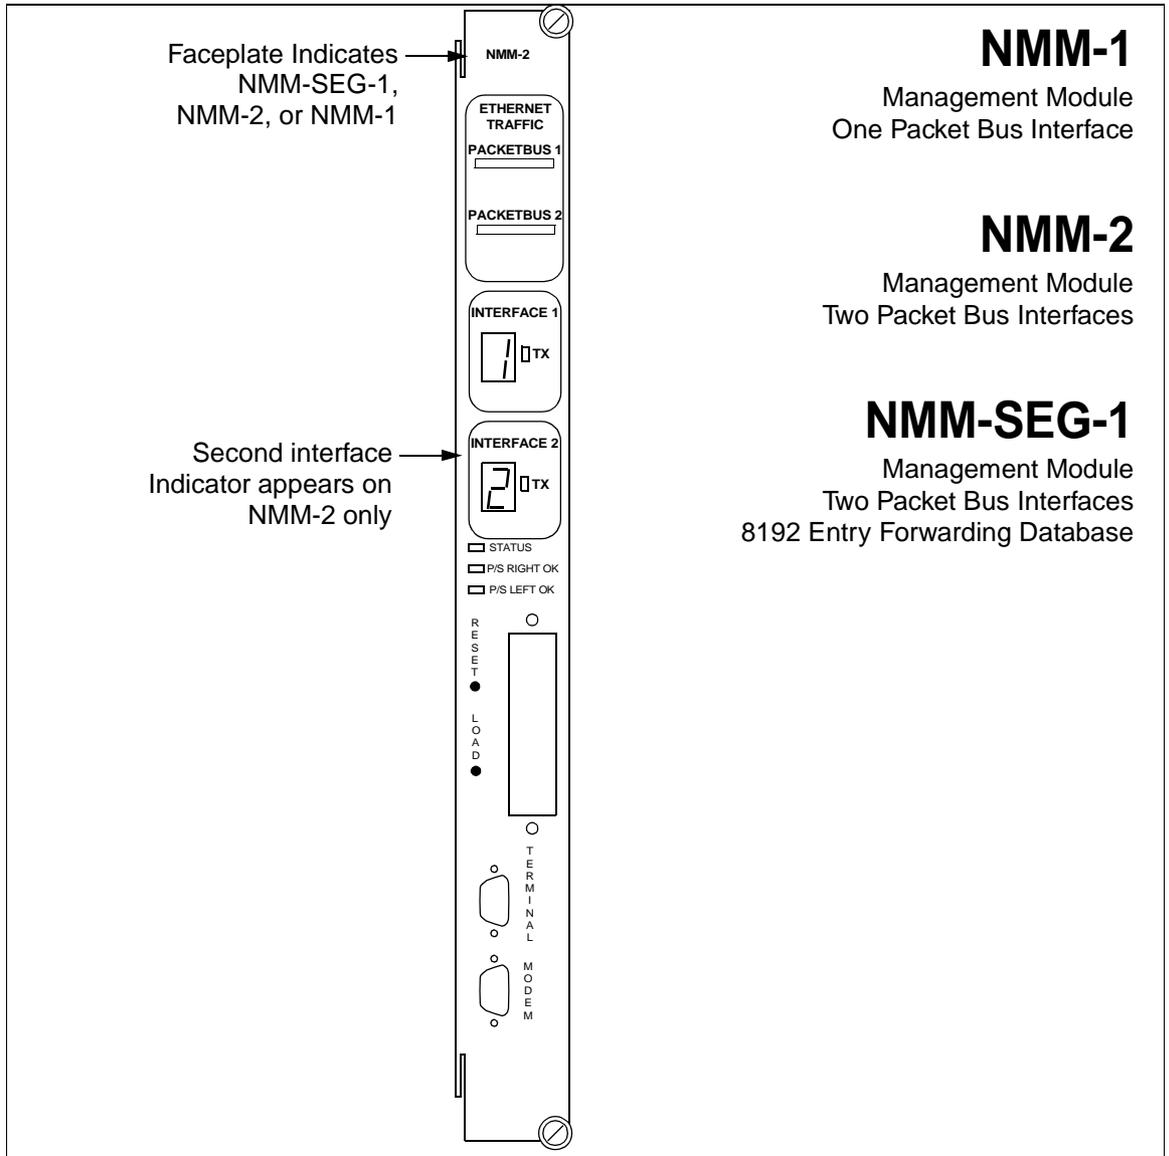


Figure 1.1 - Management Module Faceplate

1.1.1.1 NMM-1

The NMM-1 provides two fixed interfaces, a management interface and an RMON interface that both connect to packet bus 1.

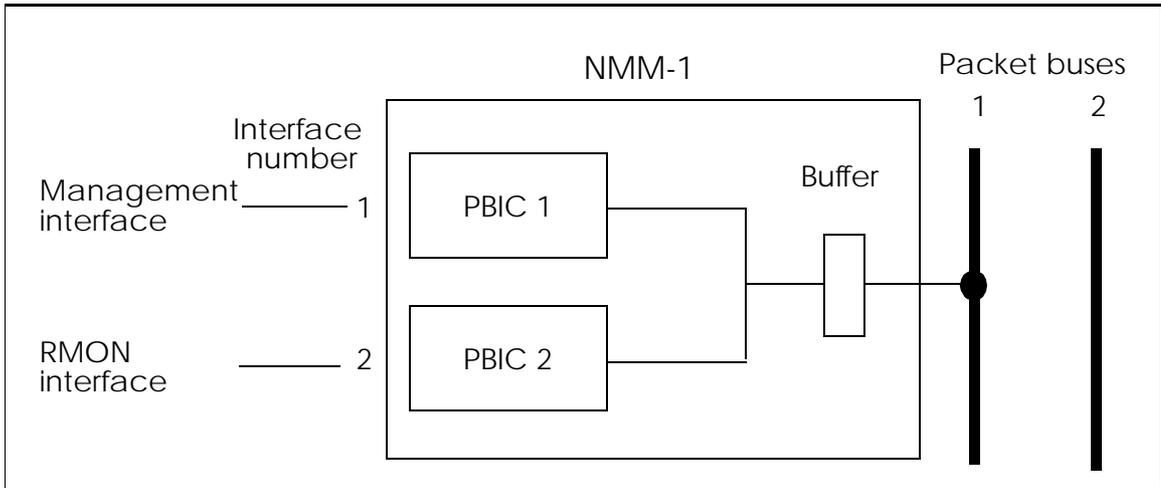


Figure 1.2 - NMM-1 Interface

Introduction

1.1.1.2 NMM-2

The NMM-2 management module provides 4 fixed interfaces that connect directly to the switched Ethernet backplane as shown in Figure 1.3.

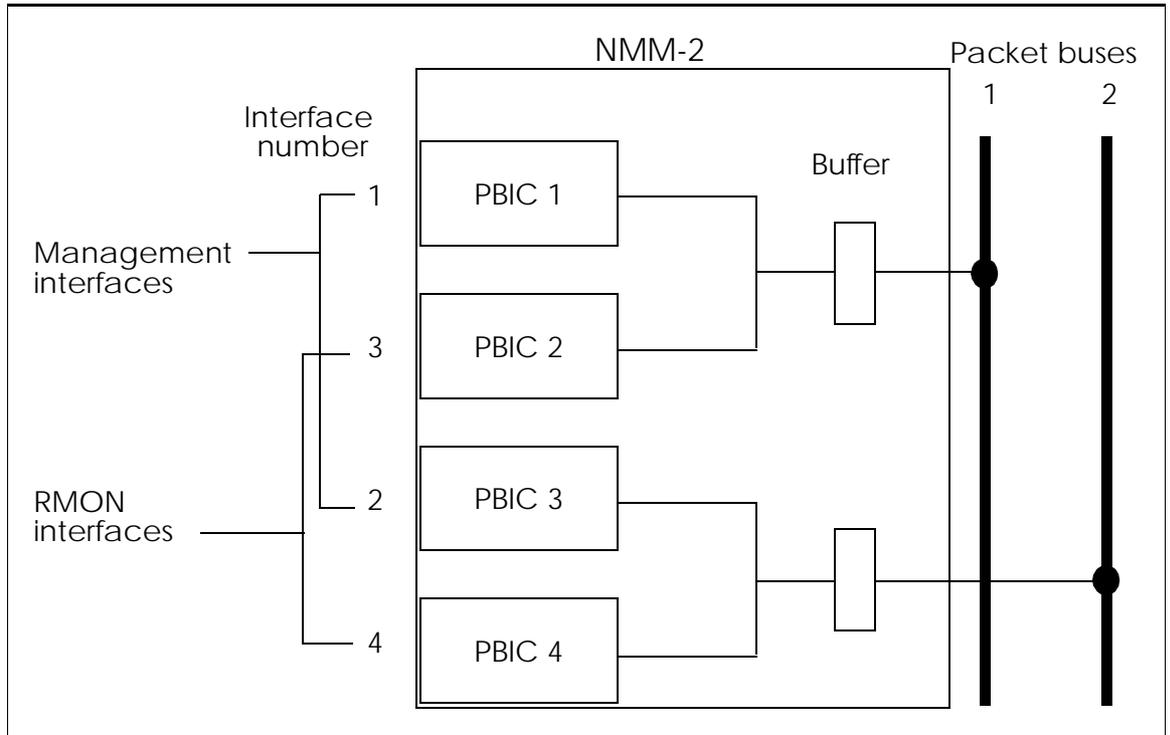


Figure 1.3 - NMM-2 Interfaces

Interfaces 1 and 2 are management interfaces and permanently connect to packet buses 1 and 2 respectively. The management interfaces only operate in normal mode.

Interfaces 3 and 4 are RMON interfaces and connect to packet buses 1 and 2 respectively. These interfaces can be set to operate in either normal or promiscuous mode. When set to normal mode, the interface looks only at sniffed packets. All packets are examined when operating in promiscuous mode.

1.1.1.3 NMM-SEG-1

The NMM-SEG-1 (Segment Switch Manager) management module provides two interfaces that can be connected to either packet bus 1 or packet bus 2 as shown in Figure 1.4.

The management interface and RMON interface are moved between packet buses together along with the content addressable memory (CAM) database.

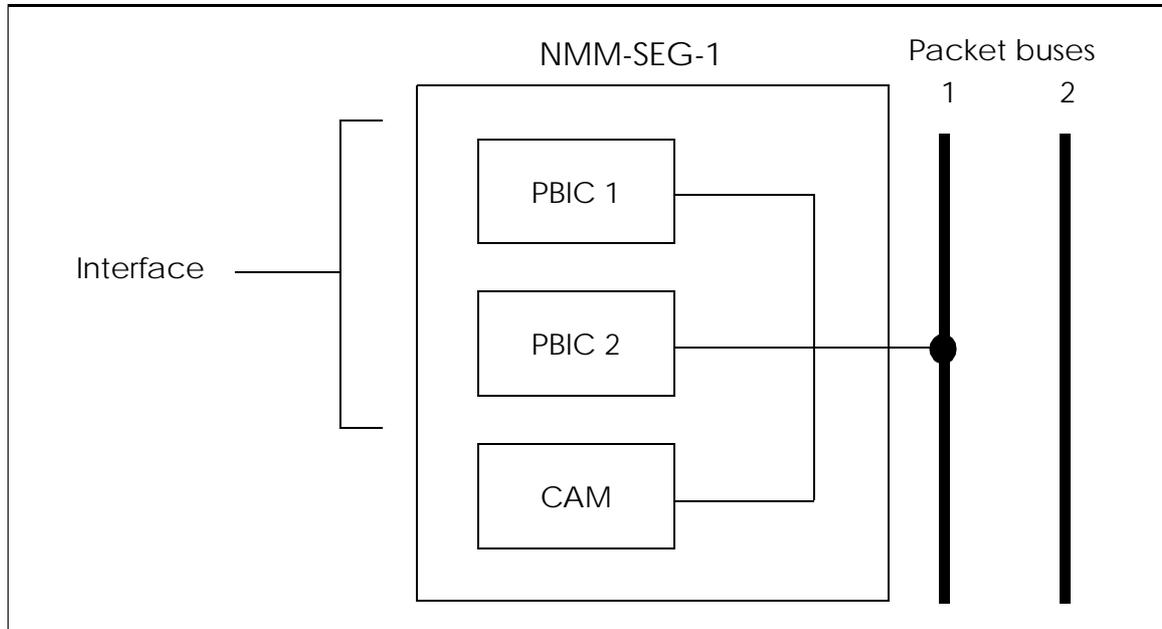


Figure 1.4 - NMM-SEG-1 Interfaces

The NMM-SEG-1 is the only management module that provides a central CAM capable of storing 8192 MAC addresses in its forwarding database. By default the CAM feature is disabled. When operating with the CAM enabled, the central CAM is used instead of the four address CAMs available on each switched user port.

WARNING!



Only one NMM-SEG-1 management module interface should be attached and enabled on a packet bus at any given time. Attaching multiple NMM-SEG-1 interfaces to the same interface will likely result in damage to management module.

The packet bus with CAM enabled is limited to the first 250 switched user ports in the chassis, counting from the first port in the first slot. Ports numbered above 250 will not be able to transmit packets.

The NMM-SEG-1 also is the only ES-4810 management module to support the Spanning Tree Protocol (STP). The Spanning Tree Protocol prevents loops by blocking packets from traveling down multiple paths to the same destination. A single instance of spanning tree is implemented for the packet bus on which the NMM-SEG-1 is connected.

By default STP is disabled. To enable STP the CAM must first be enabled. Individual ports may have STP disabled.

1.1.2 ES-4810 Backplane

ES-4810 switch backplanes have an aggregate capacity of 1.280Gbps split into two packet buses. Switch module ports can be assigned to either packet bus or placed in standalone mode.

The packet bus implements a store-and-forward scheme to filter and forward packets to their destination address. This method of packet transfer, also used in transparent bridging, is described in “Ethernet Switch Concepts” on page 1-11.

1.1.3 ES-4810 Ethernet Modules

There are several versions of ES-4810 Ethernet modules. For descriptions of the switch modules' hardware features, installation instructions, and specifications, refer to the ES-4810 Ethernet Module Operations Guide (MANU0297).

ES-4810 Ethernet modules are available with a combination of 10BaseT, 10BaseFL, 100BaseTX (Fast Ethernet), and 100BaseFX (Fast Ethernet fiber) ports. Ports on the switch module are organized into a single workgroup. Each module can operate as a standalone switch or can be interconnected with other switch modules in the chassis across the packet bus.

Each 10BaseT and 10BaseFL port has a 256K memory buffer. The ports can be configured to operate at half or full duplex and can be designated as uplink ports for packet transfer to remote switches.

Each 100BaseTX and 100BaseFX port has 1M of buffer space. The ports can be configured to operate at half or full duplex, or can be configured to automatically negotiate the duplex mode with the device connected to the port. The 100BaseTX ports can be set to operate at 10Mbps, 100Mbps, or can be configured to automatically negotiate the speed. 100BaseFX ports only operate at 100Mbps.

Other key features of the Ethernet switch ports are briefly described below. These features are described in more detail in "Ethernet Switch Concepts" on page 1-11. See Chapter 2 for information about configuring the features described in this section.

1.1.3.1 CAM Capable Management Module

The NMM-SEG-1 management module will operate in two different modes:

- **CAM disabled** - With CAM disabled it operates like a NMM-1 management module and forwarding decisions are made using the four address CAM database attached to each port.
- **CAM enabled** - With CAM enabled it ignores the four address CAM database and stores the addresses in a central database on the management module.

1.1.3.2 Address Learning

Each port can have up to four addresses stored in an address database that it uses to make filtering and forwarding decisions. This database is created by examining the source address of packets received into the port's memory buffer. If the source address does not match any of the stored addresses, it is added to the address database. When no room is left in the database, the address that has not been seen for the longest period of time is discarded and replaced by the new address.

For the NMM-SEG-1, when the CAM is enabled, the four address database is disabled and all addresses are stored in a central database on the management module capable of storing 8192 addresses.

1.1.3.3 Secure Address Learning

Addresses in the port's address database can be locked to prevent being overwritten by new addresses. This effectively disables address learning, and only those packets whose destination address matches one of those in the secure database are transmitted. Entries in the database can also be "reserved" which means that there is no address or learning capability for that entry.

The learning process also has a secure mode that allows the port to learn addresses, but does not use them until they are validated by being locked.



Even with secure learning enabled, unvalidated addresses are permitted to transmit packets onto the packet bus.



Secure address learning is not supported by the NMM-SEG-1 operating with the CAM enabled.

1.1.3.4 Virtual LAN (VLAN) Support

The hardware for each Ethernet switch port includes support for configuring virtual LANs (VLANs). VLANs can be thought of as "broadcast domains" where broadcast and multicast packets originating in a particular domain will be forwarded only to addresses in the same domain.

Each switch port's address database has a field that specifies to which of the 16 possible VLANs, or broadcast domains, it will forward multicast and broadcast packets. When a port receives a packet, it labels the packet with the VLAN information. The packet will only be forwarded by those ports that share at least one common VLAN with the labeled packet.

1.1.3.5 Station Migration

Station migration is a hardware feature that allows the port to delete invalid addresses from its database. An address in a port's address database becomes invalid when a station is moved from one port to another.

The switch ports handle station migration in this way: when port A receives a packet from the packet bus (originating from port B) it examines the source address and compares it to the addresses in its address database. If the source address matches an address in port A's database, port A knows that the address is now associated with port B and deletes the address from its database.

1.1.3.6 Uplink and VLAN Uplink Capability

The 10BaseT, 10BaseFL, 100BaseTX, and 100BaseFX Ethernet switch ports can be designated as uplink ports. An uplink port is used to forward packets destined for stations that are not local to the switch.

100BaseTX and 100BaseFX ports in uplink mode can also be assigned with VLAN information. The uplink port uses the VLAN information to forward only those packets labeled with the same VLAN information.

When using uplink ports between ES-4810 Ethernet modules, the 100BaseTX and 100BaseFX ports can encapsulate VLAN information within the Ethernet packet, so VLANs can span multiple ES-4810 modules.

1.1.3.7 Statistics Counters

Each switch port's hardware maintains 25 statistic counters. The SNMP agent on the management module uses these statistics to set values for MIB objects related to the Ethernet switch. Values for these MIB objects can be viewed using a MIB browser or SNMP network management application. Refer to "MIB Support" on page 1-10 for the list of standard and enterprise-specific MIBs the Ethernet switch supports.

1.1.3.8 Switch Monitoring

In a traditional shared-media Ethernet hub, a single interface can monitor all stations in that collision domain. Because each port in the Ethernet module has its own separate collision domain, an "analysis port" is provided to allow monitoring of each port.

Ports on the Ethernet switch modules can be configured to be a monitored port or a monitoring port. When a port is configured as a monitored port, all traffic on that port is forwarded to other ports on the same packet bus that are configured as monitoring ports. Monitoring ports receive all traffic from ports that are in "monitored" mode.

1.2 Management Interfaces

Access to the management functions of the ES-4810 are provided through the agent that resides on the management module. Interfaces available for communicating with the agent software—SNMP, the operator console, and modem—are the same for all the Ethernet agents.

Management information specific to the ES-4810 is provided below.

1.2.1 MIB Support

Table 1.1 lists MIBs supported by the ES-4810 Ethernet agent.

Table 1.1 - MIBs supported for the ES-4810

Name	RFC/other	Implementation
<i>Management Information Base of Network Management of TCP/IP-based Internets: MIB II</i>	RFC 1213	Only ifTable
<i>Extensions to the Generic Interface MIB</i>	RFC 1229	Only ifExtnsTable
Evolution of the Interfaces Group of MIB-II	RFC 1573	All except ifTestTable and ifRcvAddressTable
Definitions of Managed Objects for Ethernet-like Interface types	RFC 1643	All except dot3CollTable
ES-4810 Ethernet Switch MIB	Enterprise	All

1.2.2 Operator Console Interface

For information on these Ethernet switch procedures, refer to Chapter 2 or “User Module Configuration” on page 3-1. For console commands, refer to “Operator Console Command Reference” on page 5-15.

1.3 Ethernet Switch Concepts

This section describes basic switch concepts as they apply to the ES-4810 to help you better configure and manage the switch.

1.3.1 What is a Switch?

A switch is basically a multi-port MAC layer bridge. A switch operates the same way as a bridge in that it forwards or filters packets based on the destination MAC layer address in the packet header. The major difference between a bridge and an Ethernet switch is that a switch, with its hardware-based switching capabilities, can economically support many more wire-speed ports. Switches can also have features such as VLANs (described later in this section) to further increase the efficiency and flexibility of the bridged LAN.

In a shared media Ethernet LAN, where only one device can transmit at a time, the physical limitations of the media are quickly met with large numbers of attached workstations and heavy traffic. Segmenting the LAN with a switch or a bridge extends the LAN's physical limits and increases performance by increasing the number of collision domains in the network. The packet filtering process ensures that a packet will only be forwarded out the port with the appropriate destination address, rather than broadcast out all ports until it is received by the intended address.

Because each port in the ES-4810 has its own separate collision domain, each port has dedicated access to the network. This increases the aggregate throughput of the network, and allows any combination of ports in the switch to communicate with each other.

Each ES-4810 Ethernet module must be assigned to one of the two packet buses on the backplane or placed in standalone mode. When placed in standalone mode, a port on an Ethernet module can communicate only with other ports on the same module.

When assigned to a packet bus, all ports in the module communicate across that packet bus. Separate modules, either on different packet buses in the same ES-4810 or in different ES-4810, communicate through an uplink port. There must be an uplink port at each end of an inter-switch link.

1.3.2 Transparent Bridging

The mechanism used in the ES-4810 to forward packets to the appropriate destination address is the same as that used in a transparent bridge. This type of bridge is called “transparent” because it joins separate LANs together, making them appear as one LAN, and is transparent to higher layer protocols passing through the bridge at the MAC layer. Transparent bridging has the following characteristics:

- Handles only one MAC layer protocol at a time; no protocol translation is performed at the MAC layer.
- Does not perform routing.
- Has the ability to learn MAC addresses of hosts on the network.
- Maintains an address database for filtering packets.

1.3.2.1 Learning Process

The learning process is how the bridge builds and maintains the address database. The bridge listens promiscuously and receives all packets. For each packet received, the bridge compares the packet’s source address to the addresses in the address database and takes an action based on the results of the comparison:

- If the address is already in the database, it is marked as having been seen.
- If the address is not already in the database, the new source address is “learned,” that is, added to the database.
- If the database is already full, the address that has not been seen in the longest period of time is overwritten.

1.3.2.2 Forwarding Process

Transparent bridges use one of two methods of forwarding packets: cut-through and store-and-forward. In a cut-through bridge, each packet is forwarded as soon as the packet header containing the destination address is read. In a store-and-forward bridge, the entire packet is placed in a memory buffer before forwarding it to the destination address.

While a cut-through bridge has low latency, store-and-forward bridges check packets for integrity before forwarding and discard packets containing errors. The packet buffering capability of a store-and-forward bridge is also important for networks that have both 10Mbps and 100Mbps connections.

The ES-4810 uses the store-and-forward method of packet transfer. In the store-and-forward process, when a bridge receives a packet it examines the entire packet and stores it before forwarding. The bridge looks at the packet's destination address and makes the decision to forward or filter (drop) the packet by comparing it to the address database:

- If the destination address is not found in the address database, the packet is forwarded to all ports except the one on which it was received.
- If the destination address is found in the address database, the packet is forwarded only to the port specified in the address database. If the specified port is the one from which the packet was received, the packet is dropped.

1.3.3 Virtual LANs (VLANs)

A virtual LAN (VLAN) is a logical grouping of end stations in a switched network. Through software, devices on different LAN segments can be configured to communicate with one another as if they were on the same physical network. This simplifies the task of adding, removing, or moving devices in the network, since the devices' physical location can change without changing their logical grouping. VLANs also make networks more secure and improve performance by limiting broadcast traffic to certain areas of the network.

1.3.3.1 Using VLANs Within a Switch

As explained in “Virtual LAN (VLAN) Support” on page 1-8, the address database for each port on the ES-4810 includes VLAN information that specifies to which of the 16 possible VLANs the port belongs. The port labels the packets it receives with this VLAN information before transmitting the packets onto the packet bus. The packets will be forwarded only by those ports belonging to at least one common VLAN.

For example, suppose you want to restrict broadcasts based on departments within your organization. You could designate engineering as a member of VLAN 1, sales as a member of VLAN 2, and administration as a member of VLAN 3. Any broadcast traffic originating from a port in VLAN 1 will be transmitted only to ports in the engineering VLAN. An external router connected to VLANs 1, 2, and 3 will allow communications between sales, administration, and engineering.

Figure 1.5 illustrates this example of using VLANs.

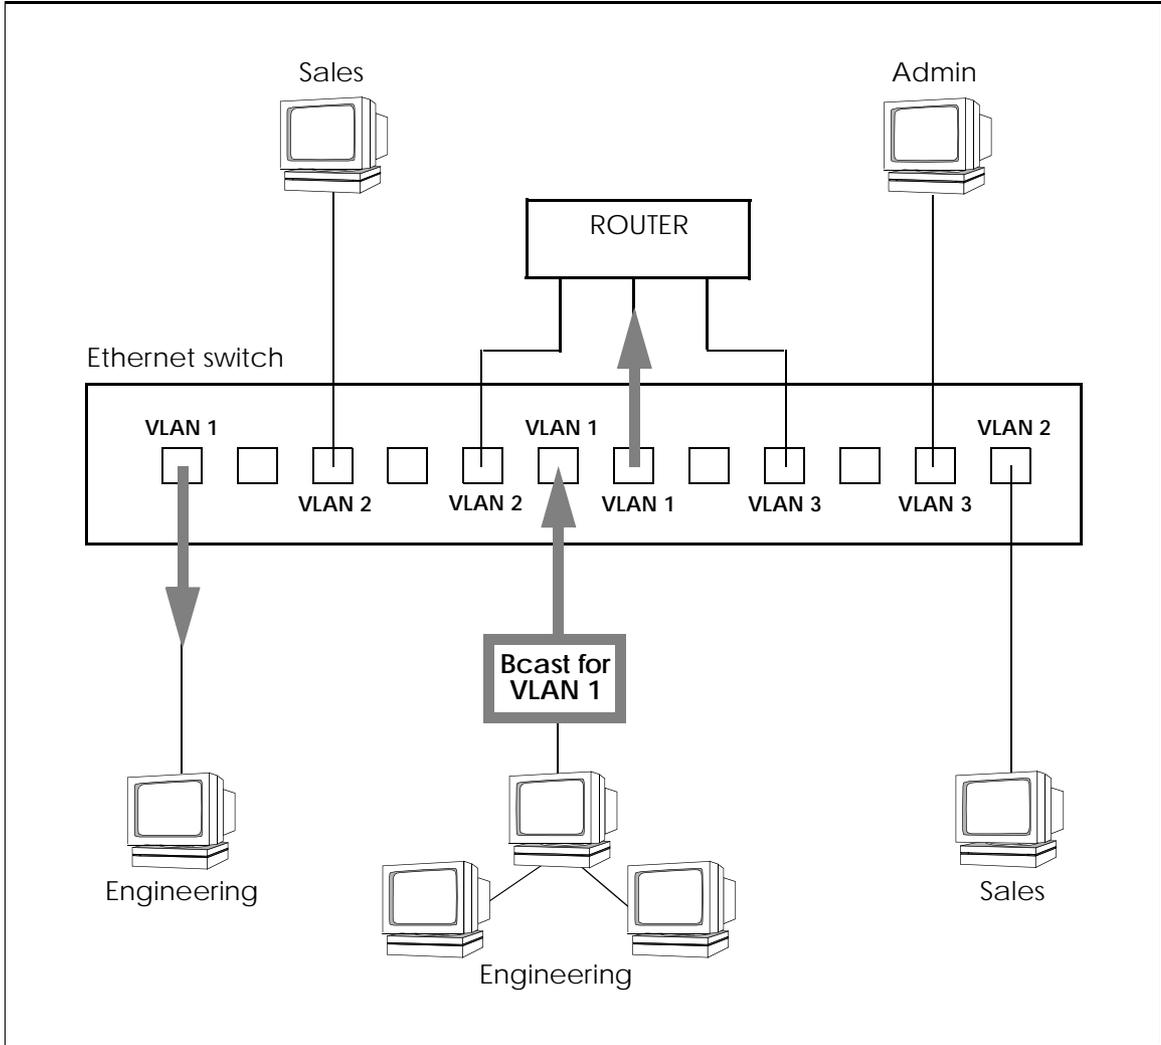


Figure 1.5 - Example of Using VLANs

1.3.3.2 Using VLANs Between Switches

In the example in Figure 1.5, the concept of VLANs applies only to packets transferred within the switch. It is possible to preserve the VLAN information between switches, under these conditions:

- The remote switch is another FORE Systems ES-4810
- The ports connecting the two FORE Systems switches are 100BaseTX or 100BaseFX ports, both configured as uplink ports that perform VLAN tagging (encapsulate the VLAN information in the Ethernet packet)
- VLANs are set up on both switches with the intent to create a VLAN between the switches

Figure 1.6 shows how the VLANs could span FORE Systems switches.

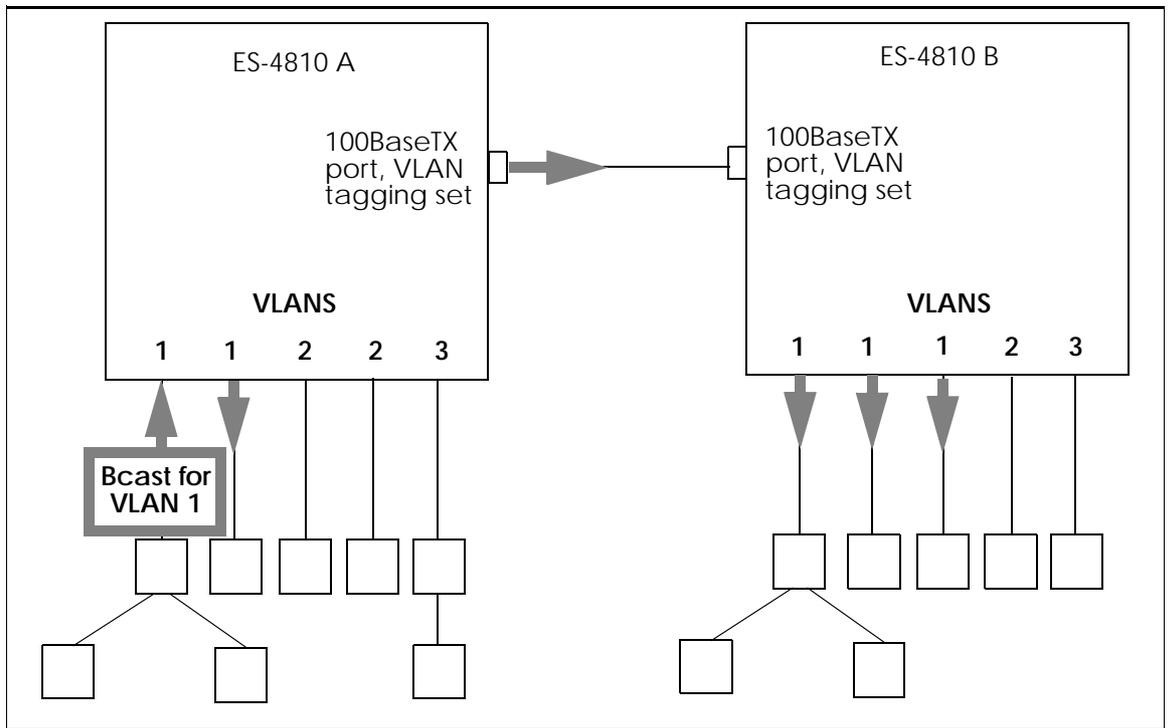


Figure 1.6 - Example of Using VLANs Across FORE Switches

In the example in Figure 1.6, both FORE Systems switches have been configured the same way, so that a broadcast packet originating on switch A for VLAN 1 is received by ports belonging to VLAN 1 on switch B.

1.3.4 Switch Port Forwarding Modes

There are several factors that affect the packet forwarding behavior of ES-4810 ports:

- The type of management module (NMM-SEG-1, NMM-1, or NMM-2)
- The switch port's forwarding mode:
 - Forward packets to another switch (**uplink**, **vlanuplink**, **uplinktag**, or **vlanuplinktag**)
 - Monitor other ports (**sniffer**)
 - Flood or filter unknown unicasts to an end station (**normal** or **limited**)
- Port type (10BaseT, 10BaseFL, 100BaseTX, or 100BaseFX)
- Type of packet being forwarded
- VLAN configuration

1.3.4.1 Forwarding Broadcasts and Multicasts

Uplink ports and end-station switch ports (except ports in **sniffer** mode) forward only broadcast and multicast packets belonging to the same VLANs as the port. Therefore, how an uplink port forwards broadcasts and multicasts is modified by setting its VLAN assignment. Ports in **sniffer** mode forward only packets that have been labeled by ports in **sniffed** mode.

1.3.4.2 Forwarding Unknown Unicasts

A unicast packet is called “unknown” when its destination address is not recognized by any of the ports in the switch. An end-station port can be configured to forward unknown unicasts to all ports (forwarding mode **normal**) or to filter unknown unicasts (forwarding mode **limited**).

There are four ways an uplink port can be configured to forward unknown unicasts. Two of these modes involve a technique called VLAN tagging (described below) and are for use between FORE Systems ES-4810s only.

The other two modes, **uplink** and **vlanuplink**, can be used to connect the ES-4810 to any other type of switch. The **uplink** forwarding mode, available on both 10BaseT, 10BaseFL, 100BaseTX, and 100BaseFX ports, causes the uplink port to forward all unknown unicasts.

In **vlanuplink** mode, the uplink port forwards only those unknown unicasts belonging to the same VLAN(s) as the port. This mode is available only on 100BaseTX ports.

1.3.4.3 VLAN Tagging Between ES-4810s

In addition to forwarding all unknown unicasts or filtering them by VLAN, a 100BaseTX or 100BaseFX uplink port can be configured to preserve the VLAN information added by the port that originally received the packet before transmitting it onto the packet bus. The uplink port encapsulates this VLAN information in the packet (VLAN tagging) so that it can be used by the ES-4810 at the other end of the link.

Table 1.2 lists the forwarding modes available for each port type.

Table 1.2 - Forwarding Modes for Ethernet Switch Ports

Port type	Forwarding mode	Description ¹	Notes
end-station	normal	Forward all unknown unicasts.	Default forwarding mode.
	limited	Filter all unknown unicasts.	Uplink ports can not filter all unknown unicasts.
	sniffer	Forward packets that have been labeled by ports in sniffed mode.	Uplink ports can not monitor other ports, but they can be monitored.
uplink	uplink	Forward all unknown unicasts.	Secure learning and modifying the address database does not apply to uplink ports.
uplink, 100BaseTX and 100BaseFX only	uplinktag	Forward all unknown unicasts, and encapsulate VLAN information in the packet.	For use between ES-4810s only. Secure learning and modifying the address database does not apply to uplink ports.
	vlanuplink	Forward only those unknown unicasts belonging to the same VLAN as this port.	Secure learning and modifying the address database does not apply to uplink ports.
	vlanu-plinktag	Forward only those unknown unicasts belonging to the same VLAN as this port, and encapsulate VLAN information in the packet.	For use between ES-4810s only. Secure learning and modifying the address database does not apply to uplink ports.

¹. Multicasts and broadcasts are forwarded according to the VLAN assignment of the port except in sniffer mode.

Each of the 7 possible forwarding modes for switch ports are mutually exclusive; you can choose only one forwarding mode for a port.

When a port is in any of the four uplink modes or sniffer mode, the entries in its address database are locked or reserved for use by the hardware. This means that you can not manually modify the address database for uplink ports or ports in sniffer mode. Also, secure learning is not applicable to ports in these forwarding modes.

Refer to “Modifying Forwarding of Unknown Unicasts on End-station Ports” on page 3-30 and “Configuring Uplink Ports” on page 3-35 for the commands to configure switch port forwarding modes.

1.3.4.4 NMM-SEG-1 Forwarding Modes

When a NMM-SEG-1 management module is being used with CAM enabled, the forwarding decisions are made by the management module and are therefore a little simpler than other ES-4810 management modules.

Table 1-3 shows three distinct modes. In both normal and tagged modes all unknown unicasts, broadcasts, and multicasts are forwarded to all ports in the same VLAN. In tagged mode, VLAN information is appended to the packet. This allows VLANs to be maintained between different chassis or packet buses



On the NMM-SEG-1, tagging only works if there are four addresses or less per port.

Table 1.3 - Forwarding Modes for Switched Ethernet (NMM-SEG-1 with CAM enabled)

Forwarding mode	Equivalent	Description
normal	normal, limited, uplink, vlanuplink	Forward all unknown unicast, broadcast, and multicast packets to all ports in the same VLAN as the source port.
uplinktag	uplinktag, vlanuplinktag	Forward all unknown unicast, broadcast, and multicast packets to all ports in the same VLAN as the source port and encapsulate VLAN information in the packet.
sniffer	sniffer	Forward packets that have been labeled by ports in sniffed mode.

1.3.5 Avoiding Loops in Switch Configuration

It is possible to accidentally create loops or multiple paths in a network that can degrade network performance. To prevent this, there must always be, at most, one route between any end stations. To avoid creating loops when configuring the switch, follow these guidelines:

- Do not connect two ports in the same switch together with a single cable.
- Do not connect switches together such that there are multiple routes between any two switches.
- Use a NMM-SEG-1 management module with Spanning Tree Protocol to block multiple paths

These guidelines are explained below.

1.3.5.1 Loops Within a Switch

The simplest example of a loop occurs when two ports in the same switch and in the same VLAN are connected. Figure 1.7 illustrates this example.

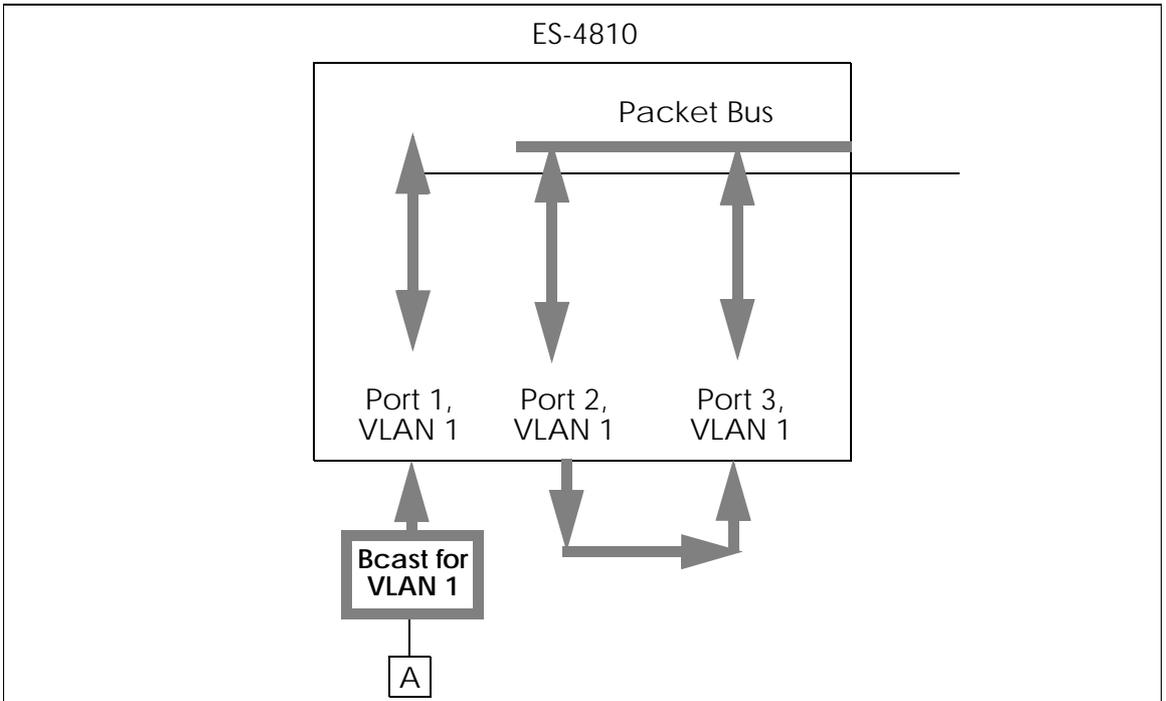


Figure 1.7 - Example of a Loop Within a Switch

In Figure 1.7, Port 2 and Port 3 are connected by a cable on the front of the switch. Station A sends out a broadcast and Port 1 forwards it onto the packet bus. Ports 2 and 3 receive the packet because they are both in VLAN 1.

Because ports 2 and 3 are connected together, both ports will attempt to retransmit the packet back onto the packet bus, causing a loop and unnecessary duplication of broadcast packets in the network.

1.3.5.2 Loops with Multiple Switches

Other simple loop examples involve multiple switches being connected together in a ring, as shown in Figure 1.8 and Figure 1.9.

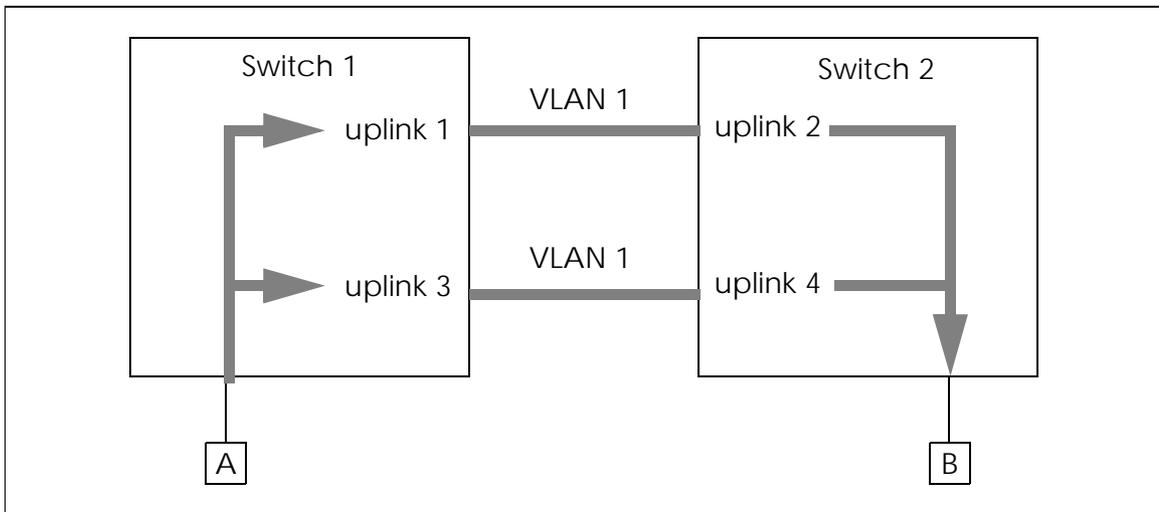


Figure 1.8 - Example of a Loop Between Two Switches

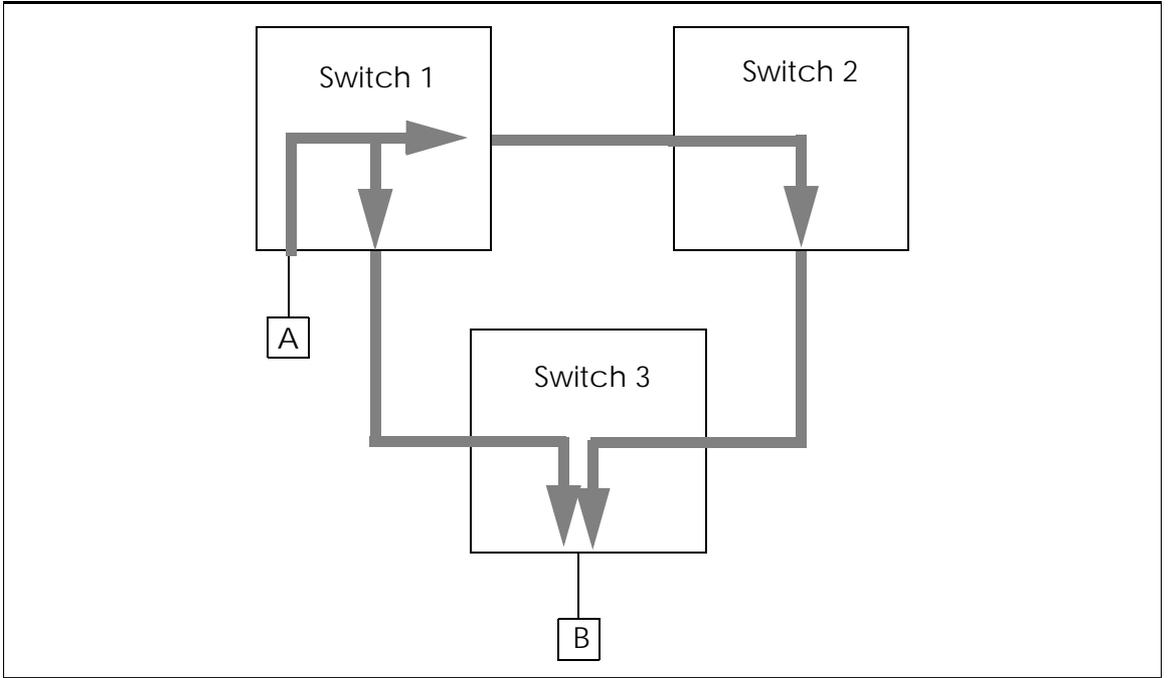


Figure 1.9 - Example of a Loop Between Three Switches

In Figure 1.8, the switches are connected by two uplinks. When station A sends a packet destined for packet B, the packet is sent out over both uplinks and station B receives the packet twice.



Even if both uplinks between the switches are using different VLANs, a loop would exist when a new station is learned and is temporarily a member of all 16 VLANs.

In Figure 1.9, each switch has two uplink ports connecting it to the other two switches. When station A sends a packet destined for station B, the packet is sent out over both wires, and station B receives the packet twice.

1.3.5.3 Load Balancing with Parallel Cables

You can use parallel cables to increase bandwidth when connecting FORE Systems switches. For load balancing to work properly, the following rules apply:

1. All switches (packet buses) connected with parallel cables must have NMM-SEG-1 management modules with CAM enabled connected to it.
2. VLAN tagging is turned off on the ports to which the parallel cables are connected.
3. Only one VLAN is carried on each cable.

Without VLANs, parallel cables between switches would create a loop because there would be two identical paths for packets destined for the other switch.

Using VLANs with parallel cables ensures that packets originating in a particular VLAN will always use the same wire, so that there are no loops or duplicated packets. Configuring switches this way balances the traffic load between the switches, thereby increasing the bandwidth between switches. Each end of the link must be configured with the same VLAN and forwarding mode. Figure 1.8 illustrates an inter-switch link using VLANs and parallel cables.

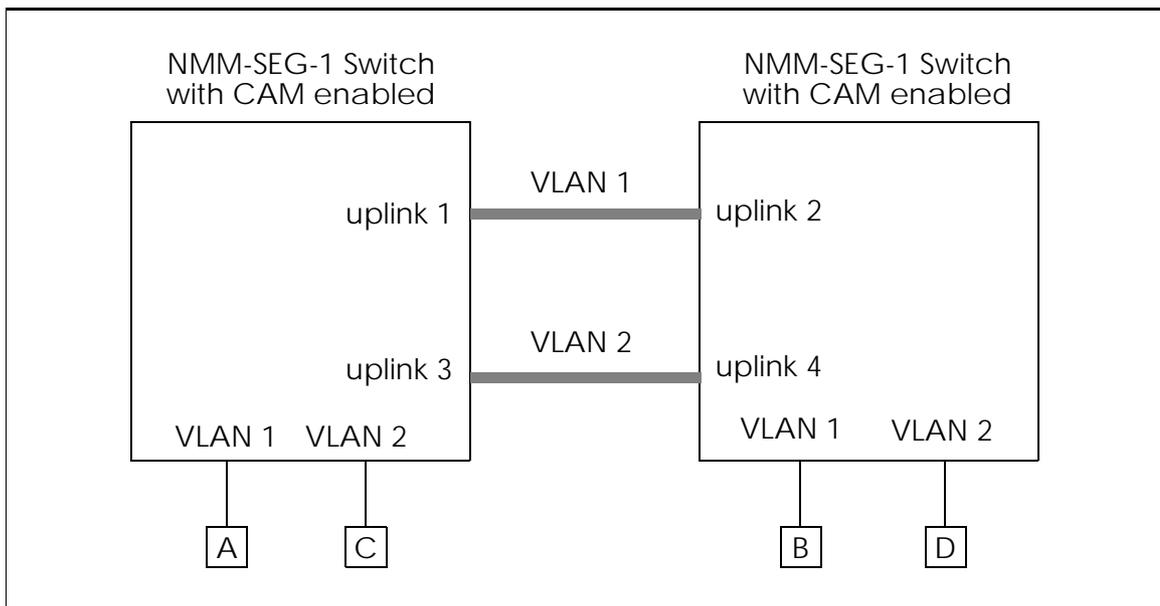


Figure 1.10 - VLANs Assignments Between Switches Using Parallel Cables

1.3.5.4 VLAN Assignments Between Switches

When setting up a parallel cable connection between two switches, VLANs must be set up correctly to avoid creating loops or multiple paths between the switches. Parallel cables must be configured with:

- Forwarding mode set to uplink
- One VLAN for each link between switches

In Figure 1.8, uplinks 1 and 2 have their forwarding mode set to the **uplink** option. Both ends of the uplink are in VLAN 1. Uplinks 3 and 4 are also using the **uplink** option and are in VLAN 2. Stations A and B are in VLAN 1; stations C and D are in VLAN 2.

Broadcast traffic between stations A and B uses the VLAN 1 uplink and is not received by stations C or D. Traffic between C and D uses the VLAN 2 uplink and is not received by A or B.

In Figure 1.8, each pair of uplink ports has a unique VLAN assignment. If uplink 2 were assigned to VLAN 2, there would be two paths for packets destined for VLAN 2, causing unnecessary duplication of packets.

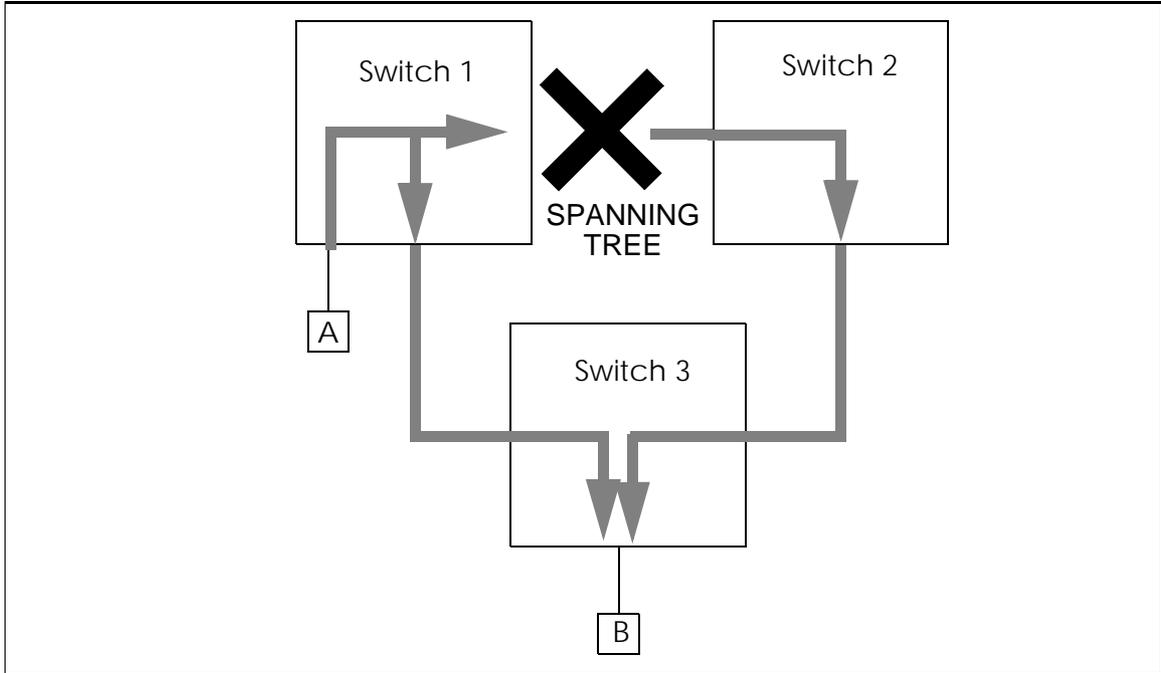
1.3.5.5 Parallel Uplink Ports and Spanning Tree Protocol

When configured properly, parallel cables do not create a loop. However, spanning tree will still detect a loop and block one of the paths. For this reason, spanning tree must be disabled on one of the four ports shown in Figure 1.10.

1.3.5.6 Spanning Tree Protocol

Spanning tree detects accidental and intentional loops in a network. When a loop is detected, spanning tree blocks a port from forwarding packets. Spanning tree only runs on the NMM-SEG-1 management module.

The example in Figure 1.9 shows a loop between three switches. If spanning tree was running, one of the paths would have automatically been blocked as shown in Figure 1.11.



Introduction

Figure 1.11 - Example of Spanning Tree Preventing a Loop

Introduction

CHAPTER 2

Management Module Configuration

This chapter describes the procedures for configuring ES-4810 network management modules. The procedures in this section explain how to initially configure the interfaces and the software that resides on the management modules.

Commands to control and modify the configuration of the other modules in the chassis through the management module are described in Chapter 3.

2.1 Overview

There are specific steps which must be performed initially to configure the management modules:

1. Assign an IP address and network mask to at least one of the interfaces on the management modules.
2. Assign each interface to a backplane segment and enable the interface.
3. Configure the frame acceptance mode for the desired interfaces if RMON is used.
4. Assign community names with access levels appropriate for the MIB variable you intend to use.
5. Define gateway addresses in the routing table.
6. Configure the sending of traps to network managers.

There are additional features including Content Addressable Memory and Spanning Tree Protocol that you can configure as well. These procedures are specific to the NMM-SEG-1 management module.

2.1.1 Online Help

You can get a listing of all of the commands, with a brief synopsis of each one, at the operator console. There are also other commands that provide various levels of on-line help. The following list describes the types of help commands available. Figure 2.2 through Figure 2.4 show examples of each one.

help [command]	With no options, this command lists all high-level commands. If the <i>command</i> parameter is specified, the options for the specified command are listed. This command can also be specified as: <i>command help</i> .
setup	Displays general instructions for set up of the management and user modules.
if setup	Provides a list of the commands that are used to configure the interfaces between the management module and the backplane segments.
community help	Lists the commands that are used to set up the community strings that provide up to five levels of access to MIB data via SNMP.

```
>> help
help      - list available commands
setup     - instructions for initial setup of the chassis
atmuplink - set/show ATM uplink information
boot      - set/show boot information
escam     - set/show ISCAM information
card      - set/show card information
chassis   - show chassis information
check     - analyze the chassis configuration
community - set/show community information
ebpseg    - set/show ethernet backplane segment information
egroup    - set/show ethernet group information
```

Figure 2.1 - The HELP Command on the Management Module (One of Two)

```
esport      - set/show switched ethernet port information
esgroup     - set/show switched ethernet group information
esbpbus     - set/show switched ethernet backplane bus information
espair      - set/show ethernet switched port pair information
group       - set/show ethernet/token-ring group information
if          - set/show interface information
ping        - ping an ip address
port        - set/show ethernet/token-ring port information
quit        - exit the command line interface
reboot      - reboot the network manager
rmon        - set/show rmon information
route       - set/show route information
slip        - enter SLIP mode on this port
snmp        - set/show snmp party information
stbridge    - set/show Spanning Tree Bridge information
stport      - set/show Spanning Tree Port level information
system      - set/show SNMP system information
trap        - set/show trap information
version     - show the network manager version
type 'help COMMAND' for more help
```

Figure 2.2 - The HELP Command on the Management Module (Two of Two)

Management Module Configuration

```
>> setup
>>The MINIMUM setup required for this chassis entails -
  Initializing all groups on all User port cards
  (ie. group/trgroup commands),
  Initializing all interfaces on the Network Management card
  (ie. if commands),
  Initializing the community names used to communicate
    via SNMP with the agent resident on the Network
    Management card
  (ie. community commands)

For each group on each User card, you need to initialize each
  port to either enabled or disabled, AND
  You need to assign each group to a backplane segment
For each interface on the Network Management card, you must
  enable it, and assign an IP address, a net mask,
  and a backplane segment
Common to all interfaces on the Network Management card,
  you need to assign a community name for read-only
  access and, if you wish to change configuration of the
  chassis via SNMP, you need to assign a community name
  for read-write access
>>
```

Figure 2.3 - The Setup Command on the Management Module

```
>> community help
community help
community show <entry>
community ip <entry> <ipaddress>
community priv <entry> <privilege(ie. ronly/rwrite/super/noacc/limronly/v2rw)>
community name <entry> <name>

The setup required for the community table entails -
  For each entry that you want to be an active entry
  in the community table, you want to do the following:
>> community ip <entry> <ipaddress>
>> community priv <entry> <priv>
>> community name <entry> <name>
    where - <entry> is the index into the community table
           - <ipaddress> is 0.0.0.0 by default, or a specific ip address if
             you only want one ip address to be able to use this community
           - <priv> is the privilege afforded this community
           - <name> is the ascii string that identifies this community
             (NOTE: set name to an empty string for any entry in the
                 table that you do not want to be used.)
>>
```

Figure 2.4 - The community help Command

For more details about a particular operator console command, refer to Chapter 5 of this manual.

2.1.2 Logging On and Off the Operator Console

The login feature prevents unauthorized access to the module management functions. The console login feature is enabled when at least one community string has been created with an access privilege of *super*. The login feature is disabled when there are no community strings with a *super* access privilege.

To log onto the console, enter the name of any community string which has a *super* access privilege at the login prompt, as shown in Figure 2.5.

```
login: ***** <Enter>
>>
```

Figure 2.5 - Logging On to the Console

During the installation of the management module, you were instructed to log onto the console with the login ID *super*. This is the default community string that has *super* access privilege as defined at the factory. It is recommended that you modify this community string as soon as possible. Refer to the procedure “Community String Configuration” for details on how to set up and modify community strings.

Enter **quit** at the console prompt “>>” to log out of the console.

2.1.3 Community String Configuration

Community strings are the “passwords” used by the management module to grant network management applications access to SNMPv1 MIB objects. The *super* access level community string is also used as the operator console login as describe in the section “Logging On and Off the Operator Console” on page 2-6.

Up to eight strings can be assigned, with each string corresponding to a specific level of access privilege for a network management application. To display the current community string configuration enter **community show**.

Note, if SNMPv1 is disabled, a warning message will be displayed when you issue this command.



The community strings for the management modules, ATM uplink modules, and ASX switch modules must be assigned separately and are not automatically shared among the components of the ES-4810.

```
>> community show
entry 1:  ip address 199.80.90.57, privilege: super name: vegas
entry 2:  ip address 0.0.0.0, privilege: rwrite  name: U2
entry 3:  ip address 0.0.0.0, privilege: ronly  name: private
entry 4:  ip address 0.0.0.0, privilege: limronly name: public
entry 5:  ip address 0.0.0.0, privilege: noacc  name:
entry 6:  ip address 0.0.0.0, privilege: noacc  name:
entry 7:  ip address 0.0.0.0, privilege: noacc  name:
entry 8:  ip address 0.0.0.0, privilege: noacc  name:
>>
```

Figure 2.6 - Displaying Community String Configuration

The command used to assign a community string name has the format:

```
community name entry name
```

where:

- entry** The entry (1-8) in the community table.
- name** A case sensitive string with a maximum of 15 characters. Enclose the string in double quotes when entering names with blanks in them. Enter an empty string for entries that you do not wish to use. Empty strings can be entered using "" or by omitting the name parameter entirely.

To enter an access privilege for a community string, use the command:

```
community priv entry privilege
```

where:

- entry** The entry (1-8) in the community table.
- privilege** One of the strings described in Table 2.1.

Table 2.1 - Community String Access Privilege Parameters

<i>privilege parameter</i>	Function
noacc	No privileges.
limronly	Permits read only access to the “system” group of MIB-II objects.
ronly	Permits read only access to all MIB objects designated as read-only or read/write (refer to MIB documentation).
rwrite	Permits read/write access to all MIB objects designated as read/write and read access to all MIB objects designated as read-only (refer to MIB documentation).
v2rwrite	Allows an agent running SNMPv1 to access the SNMPv2 MIB objects in RFCs 1447 and 1450. Permits read/write access to all MIB objects designated as read/write and read access to all MIB objects designated as read-only.
super	1)Enables console login feature; 2)Permits read/write access to all MIB objects designated as read/write and read access to all MIB objects designated as read-only; 3)Permits management module software updates via TFTP.



Remember to write down the community name of at least one community string which was assigned the super privilege.

If you forget the community string that you assigned to the super access level, refer to the following section 2-9

The access privilege associated with a particular community string can be further limited to a single IP address. By assigning an IP address to a community string, only accesses originating from that IP address will be allowed. To assign an IP address to a community string enter:

```
community ip entry ipaddress
```

where:

entry is the entry (1-8) in the community table.

ipaddress is the IP address from which access will be permitted, expressed in standard dot notation. The factory default address, “0.0.0.0”, allows access from any address (no restrictions)

2.1.4 Recovering from a Forgotten Login ID

If you have modified the *super* access level community string to something other than the default string *super* and you have forgotten the new string, it can be reset to the factory default string. Since the *super* level community string is used as the console login ID, if it is modified and lost, no one would be able to log onto the console.

The following procedure describes the steps for resetting the community string to the factory default. To perform the procedure, you will need a terminal with a serial port and an RS232 cable (included with the chassis).

To reset the login ID to the factory default:

1. Connect the terminal to the **TERMINAL** port.
2. Configure the terminal for 9600 baud, no parity, 8 data bits and 1 stop bit, start the terminal emulator program.
3. Press **<Enter>** and verify that the `login:` prompt is displayed.
4. Using a pointed object, press and hold the **LOAD** push-button on the front panel of the management module.
5. While still pressing the **LOAD** button, press and release the **RESET** button.
6. Release the **LOAD** button.
7. Wait for the Self-Test-Diagnostic to complete. Then, verify that the monitor program prompt `ES-4810 Bug>>` is displayed.
8. Type **restore** and press **<Enter>**.
9. Press and release the **RESET** button.
10. After the agent reboots, the terminal will display the login prompt.
11. Type the *super* community string and press **Enter** to log in.

2.2 Configuring Management Interfaces

The types of configuration functions that need to be performed on ES-4810 management modules include:

- Assigning the IP address and netmask to each interface
- Assigning each interface to a backplane segment/ring
- Configuring the frame acceptance mode

2.2.1 Displaying the Current Interface Configuration

You can use the `if show` command to display the current setup of any or all of the management module's interfaces to the packet bus. The format of the command is:

```
if show [if]
```

where:

if is the interface number. If the *if* parameter is not specified, all interfaces are displayed.

Figure 2.7 shows the output of the `if show` command on an ES-4810 management module. Each of the configurable parameters are labeled. This command will display both the current and new (NVRAM) values when they are different.

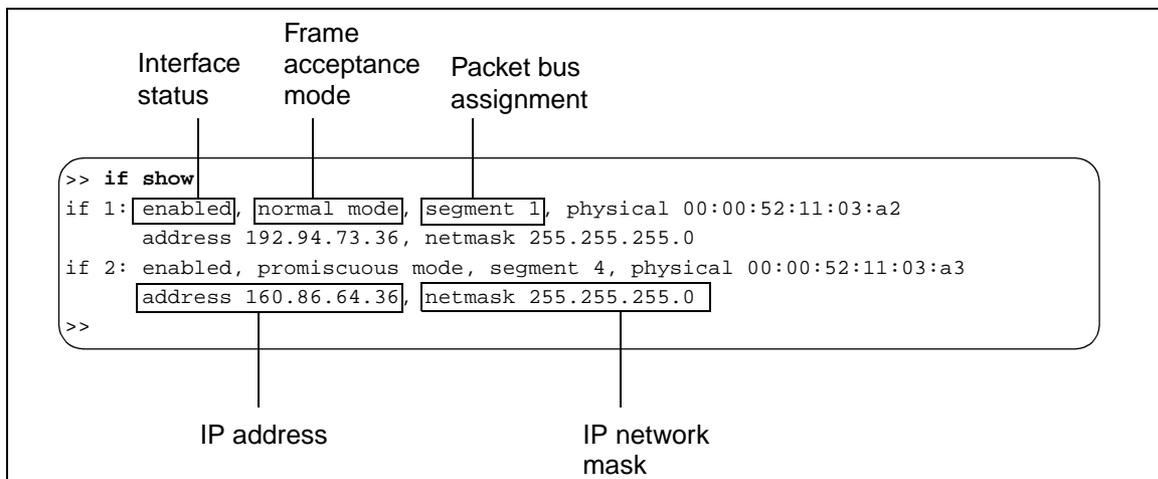


Figure 2.7 - Ethernet Interface Configuration Display

2.2.2 Assigning the IP Address and Network Mask

One of the first things that should be done after installing a management module in a chassis is to configure the IP address and subnet mask for each interface on the management module.

Each interface must have an IP address in order to communicate with other devices on the network. Without an address assigned, the interface can still monitor the backplane segments and collect statistics but no other devices can communicate with the interface.

Furthermore, for management modules with two or more interfaces, each interface must have IP addresses that are on completely different networks. When the management module is initially powered on, the IP addresses of the interfaces will be set to “0.0.0.0” or to an address that was used for testing at the factory. In either case, you should obtain a valid IP address for your site and configure each interface.

If you do not want to assign an IP address to the interface at this time, you can assign the null IP address, “0.0.0.0.” Using this address makes the hub unavailable over the network, though it can still be accessed using the operator console on a terminal connected to the hub’s `TERMINAL` port. The interface will continue to monitor the packet bus and collect statistics.

You can set the IP address for each interface using the following command:

```
if address if address
```

where:

if is the interface to be configured.

address specifies a valid IP address expressed in standard dot notation. The address can not be assigned to any other interface.

The subnet mask (also called IP net mask) is used to specify what part of an IP address designates the network and what part designates nodes or hosts on the network. Each octet in the netmask that is greater than zero indicates that the corresponding octet of the IP address specifies the network. Each octet set to zero indicates that the corresponding portion of the IP address designates a network node.

The octets are shown in Figure 2.8 in decimal, though they are actually used in binary. Each octet (eight bits) is a mask for that part of the address. Specifying “255” for the first three octets of the mask sets all 8 bits to 1 (on). This indicates that the first three octets of the IP address identify the network. However, the last eight bits of the final octet are zero, indicating that the last octet of the IP address identifies nodes on the network.

	Network			Node
IP address:	192	.86	.24	.10
Subnet mask:	255	.255	.255	.0

Figure 2.8 - Relationship of IP Address to Subnet Mask

If you wanted to have more networks and less nodes on each network, you could set some bits in the last octet of the netmask to 1. The bits set to 1 would designate networks, as shown in Figure 2.9.

	Network			Node
Subnet mask:	255	.255	.255	.240
Binary:	1111 1111	1111 1111	1111 1111	1111 0000

Figure 2.9 - Defining a Subnet Using Subnet Mask

In the example in Figure 2.9, only the last four bits of the last octet are used to number the nodes on the network. Using just the last four bits only allows 16 nodes per network. However, the first four bits of the octet allow the designation of 16 times more networks than were available before.

The subnet mask is used by the management to route packets. The network part of the IP address tells the device whether the destination for the packet is on its same network. Once the correct network for the packet is found, the node part is used to determine which device is the destination (or source).

The default value of the netmask is 255.255.255.0. To modify the netmask, you can issue the **if netmask** command from the operator console.

```
if netmask if netmask
```

where:

- if** is the interface to be configured.
- netmask** specifies a valid IP bit mask expressed as decimal values in standard dot notation.

The management module must be rebooted to make the IP Address or IP Net Mask change take effect. To reboot the management module, refer to “Rebooting or Resetting the Management Module” on page 2-33.

2.2.3 Assigning Interfaces to Packet Buses

The ES-4810 management modules vary in the number of available interfaces. Some interfaces are fixed to a particular packet bus and others can be switched between them.

For interfaces that can be assigned to any packet bus, it is not advisable to assign two interfaces to the same packet bus. This is unnecessary and may cause some management module functions to behave abnormally. ES-4810 Ethernet interfaces do not have to be disabled before being reassigned to a different packet bus.

Several management modules feature an auxiliary port. By connecting a network manager to the auxiliary port, you can communicate with the management module using SNMP or telnet (refer to “ES-4810 Description” on page 1-1).

To assign an interface, issue the following command at the operator console prompt:

```
if segment if segment
```

where:

- if** is the interface number.
- segment** is the packet bus number (1-2).

Figure 2.10 shows an example of displaying the interfaces for the NMM-SEG-1 management module.

```
>> if show
if 1: enabled, normal mode, segment 1, physical 00:00:52:30:c4:00
    address 160.86.8.178, netmask 255.255.255.0
    vlan 1
if 2: enabled, normal mode, segment 1,
    vlan 1
if 3: enabled, normal mode, segment AUX, physical 00:00:52:30:c4:02
    address 160.86.81.125, netmask 255.255.255.0
if 4: Term.: down, speed 9600, flow-ctl s/w, car-sense off, timeout none
    address 0.0.0.0, netmask 255.255.255.0, destaddr 0.0.0.0
if 5: Modem: down, speed 9600, flow-ctl s/w, car-sense off, timeout none
    address 0.0.0.0, netmask 255.255.255.0, destaddr 0.0.0.0
>>
```

Figure 2.10 - The if show Command on the NMM-SEG-1

2.2.4 Displaying and Setting Packet Bus Memos

You can define a memo for each packet bus. These memos can be used to provide information about the network segment on the user group associated with it.

To display the current setting of the memo for a particular Ethernet backplane segment, issue the following command:

```
ebpseg show segment
```

where *segment* is a packet bus number (1-2).

Figure 2.11 shows an example of defining a memo.

```
>> ebpseg show
Backplane Switch Bus 1 Engineering
Backplane Switch Bus 2 Marketing
>>
```

Figure 2.11 - Defining and Displaying a Backplane Memo

2.2.5 Configuring the Frame Acceptance Mode

To provide detailed statistics for the RMON functions of the management module, you need to configure the interfaces to accept all network traffic.



Frame acceptance mode for agent interfaces can only be set to normal mode. The information in this section applies to RMON interfaces.

RMON interfaces have two modes that determine which packets contribute to the statistics it collects: *normal* or *promiscuous*. These modes are described as follows:

Normal mode	The interface only accepts packets from the packet bus for all ports set to sniffed mode.
Promiscuous mode	The interface accepts every packet present on the packet bus.

In most cases, the RMON interface will be set in promiscuous mode. The normal mode provides a way to isolate and focus in on individual ports.

Use the following command to change the RMON interface mode:

```
if mode if mode
```

where:

if is the interface number.
mode is either normal or promiscuous.

2.3 Defining the Routing Table

Both of the packet buses in a chassis are separate networks. The ES-4810 management modules provide the capability to route packets between the packet bus and other networks. Through its interfaces, the agent looks at the destination IP address of packets that are addressed to the agent. If a packet is found with a destination address that is on a different network segment, the management module looks in its routing table for a gateway or router that can reroute the packet to the proper network.

The routing table can be viewed using the `route show` command as shown in Figure 2.12.

```
>> route show
Destination      Gateway
127.0.0.1        127.0.0.1
default          169.86.81.1
169.86.8.0       169.86.8.178
169.86.81.0      169.86.81.125
>>
```

Figure 2.12 - Displaying the Routing Table

The management module automatically defines a number of “static” routes in the table that can not be modified or deleted. One of these is a loopback interface that the management module uses for internal routing. There is also a route defined for each of the management module interfaces that has an IP address assigned. The management module shown in Figure 2.13 above has only two interfaces, whose addresses are 162.84.6.1 and 199.98.70.12.

When a packet is received that is destined for one of the networks listed in the Destination column, the packet is sent to the respective gateway for routing (in this case, one of the interfaces). If a packet is received that is destined for a network that is not listed in the table, it is dropped.

To avoid losing packets, you can add up to eight more routes to the table for other network segments that you know will be attached to the hub. These “user-defined” routes are stored in NVRAM and are maintained through resets and outages. Furthermore, an additional route can be specified as the “default” route. Any packets that can not be routed to any of the gateways in the table will be sent to the default gateway.

The `route show` command will also display the user-defined routes, if there are any defined. You can view just the user-defined routes stored in NVRAM by issuing the `route show nvram` command as shown in Figure 2.13.

```
>> route show nvram
Destination      Gateway          Net/Host
162.84.4.0       162.84.6.1      Net
162.84.2.2       199.98.70.12    Host
>>
```

Figure 2.13 - Displaying User-defined Routes in NVRAM

To add routes to the routing table, you can issue the following command from the operator console:

```
route add [net | host] destination_ip gateway_ip
route add [net | host] default gateway_ip
```

where:

- destination_ip*** specifies the destination network or host in standard dot notation. If you specify the key word `default` or the destination address “0.0.0.0,” the *gateway_ip* specified is used as the default gateway.
- gateway_ip*** specifies the IP address of a device on the local networks, expressed in standard dot notation.

Routes to a particular host must be distinguished from those to a network. The optional keywords `net` and `host` force the destination to be interpreted as a network or a specific host, respectively. If neither is specified, the management module software tries to determine whether the destination is a host or network.



Although `route add` commands can be processed without the optional `net` or `host` parameter, it should be included for best results.

Figure 2.14 shows an example of adding routes to the routing table.

```
>> route add default 199.98.70.177
>> route add 162.84.7.10 162.84.6.2
>> route show nvram
Destination      Gateway          Net/Host
default          199.98.70.177   Net
162.84.4.0       162.84.6.2      Net
162.84.2.2       199.98.70.12    Host
162.84.7.10      162.84.6.2      Host
>>
```

Figure 2.14 - Adding Routes to the Routing Table

You can not add more than eight routes, not including the default route. Also, an error message is issued if you try to add a gateway that is not on one of the networks known to the management module (i.e., one of the interface networks).

To delete any of the user-defined routes, use the **route delete** command, which has the following syntax:

```
route delete destination_ip gateway_ip

route delete default [gateway_ip]
```

where:

- destination_ip** specifies the destination address of the entry you want to delete. If you specify the key word `default` or the destination address `0.0.0.0`, the default gateway is deleted.
- gateway_ip** specifies the IP address of the entry you want to delete. When deleting the default gateway, this parameter is optional.

The example in Figure 2.15 shows the combined list of all routes using the **route show** command. It also shows a failed attempt to delete one of the “static” routes, which is not allowed.

```
>> route show
Destination      Gateway
127.0.0.1        127.0.0.1
162.84.6.0       162.84.6.1
199.98.70.0      199.98.70.12
default          199.98.70.177
162.84.2.2       199.98.70.12
162.84.7.10      162.84.6.2
>> route del 162.84.6.0 160.86.6.1
Deleting this route is not allowed
>> route del 162.84.7.0 162.84.6.2
>> route show nvram
Destination      Gateway      Net/Host
default          199.98.70.177 Net
162.84.2.2       199.98.70.12 Host
>> route del 0.0.0.0
>> route show nvram
Destination      Gateway      Net/Host
162.84.2.2       199.98.70.12 Host
>>
```

Figure 2.15 - Displaying All Routes

2.4 Configuring Traps

The ES-4810 management modules support the sending of the standard, “generic” SNMP traps as well as a number of FORE enterprise-specific traps to user-defined network management applications.

2.4.1 Configuring the Network Manager Trap Table

A table stored on the management module is used to determine which IP addresses of the network managers that will receive the trap information. Each IP address listed in the trap table will receive each trap generated by the agent on the management module. The trap table can hold a maximum of eight IP addresses. To display the current trap table use the `trap show` command as shown in Figure 2.16.

```
>> trap show
trap 1: address 199.80.90.5, community boingo
trap 2: address 175.78.25.4, community sade
trap 3: address 175.78.25.15, community inxs
>>
```

Figure 2.16 - Displaying the Trap Table

To add an IP address to the trap table, enter the command:

```
trap add address [community]
```

where:

- address** is an IP address of a network manager expressed in standard dot notation.
- community** is optional and is a string of up to 15 characters (defined by the receiving station) that is sent in the community field of the trap when a trap is sent. A string containing blanks must be enclosed in double quotes.

To change the community string of an existing trap, use the command:

```
trap community address community
```

where:

- address** is an existing IP address in the trap table expressed in standard dot notation.
- community** is a string of up to 15 characters (defined by the receiving station) that is sent in the community field of the trap when a trap is sent. A string containing blanks must be enclosed in double quotes.

To delete an IP address from the trap table, enter:

```
trap delete address
```

where *address* is an existing IP address in the trap table expressed in dot notation.

2.5 Enabling CAM

The NMM-SEG-1 management module offers a centralized CAM with the ability to maintain 8192 addresses. The CAM is resident on the management module and overrides the 4 address CAMs on the user modules. Once enabled, the management module will monitor the packet bus and make all forwarding decisions.

WARNING!



Only one NMM-SEG-1 management module interface should be attached and enabled on a packet bus at any given time. Attaching multiple NMM-SEG-1 interfaces to the same interface will likely result in damage to the management module.



Only one CAM enabled interface may be attached to any packet bus at a given time.

To enable the CAM use the following command:

```
escam enable
```

2.6 Setting Up Spanning Tree

The Spanning Tree Protocol is used to prevent loops in a network by intelligently blocking traffic and preventing it from being forwarded down duplicate paths. FORE Systems provides spanning tree capability on the Segmented Switch Management module (NMM-SEG-1).

2.6.1 Enabling Spanning Tree



Spanning Tree is only available on packet buses that have a NMM-SEG-1 management module operating with a central CAM already enabled.

To enable spanning tree use the following command:

```
stbridge enable
```

After being enabled, the management module (i.e bridge) will communicate with other bridges to become part of the network topology.

2.6.2 Setting Spanning Tree Parameters

Parameters such as priority, aging time, hello time, and forward delay time can be set for the spanning tree bridge. The values for most of these are only used when the bridge is the root bridge in the tree. Otherwise the values are obtained from the root bridge.

2.6.2.1 Priority

Priority is a value used as the first two hexadecimal numbers in generating a bridge ID. The default number is 8000H.

The bridge ID is used to determine which bridge will become the root. Bridges with lower IDs have a better chance of becoming the root bridge.

To set the priority enter the following command:

```
stbridge priority priority
```

where *priority* can be entered in decimal, octal or hexadecimal notation.

Decimal numbers can be any number but must not begin with a leading 0 (e.g. 1234). Octal numbers must begin with a leading 0 (e.g. 07). Hexadecimal numbers must begin with a leading 0x (e.g. 0xA9).

The valid range for priority is 0-65535 (0x0 - 0xFFFF).

2.6.2.2 Hello Time

Hello time is the number of seconds the bridge will pause between the transmission of Hello Messages (BPDUs). The default time is 2 seconds. The range of valid time is 1-10 seconds.

2.6.2.3 Max Age

The max age parameter determines how long a bridge should wait without receiving a BPDU before attempting reconfiguration. Normally, the bridge ports will receive BPDUs at regular intervals. The default time is 20 seconds. The range of valid time is 6-40 seconds.

To set the max age, enter the command:

```
stbridge age 15
```



The unit of time is in seconds. The value must meet the following requirements: $age_time \geq 2 \times (hello_time + 1)$

2.6.2.4 Forward Delay Time

The forward delay time is used to delay the changing of port states from listening to learning and learning to forwarding. This time is needed so every bridge on a network can receive information about the topology change before the port starts to forward packets. The default time is 15 seconds. The range of valid time is 4-30 seconds.

To set the forward delay time, enter the command:

```
stbridge forwarddelay 20
```



The value must meet the following requirements:

$delay_time \geq (age_time / 2) + 1$

By default, the traps issued by the spanning tree protocol are disabled. To enable new root traps enter the command:

```
stbridge newroot enable
```

To enable topology change traps, enter the command:

```
stbridge topchange enable
```

2.7 Configuring SNMPv2 Security Features

SNMPv2 is intended to provide improved security over SNMPv1. The parties, contexts, views and access control parameters must be configured in the agent to provide the network security feature. This configuration is only necessary, however, if you will be using a network management application that uses SNMPv2.

The Party database is initialized at the factory with no entries in nonvolatile storage. The reason for this is that the Initial Parties and Contexts, as defined in the Party MIB, use the agent's IP address as part of the `partyTable` and `contextTable` object ID (OID).

To list all of the commands available for configuring SNMPv2 security, issue the `snmp help` command as shown in Figure 2.17.

```
>> snmp help
snmp auth-key [hex|ascii] <party-index> <private-key>
snmp default
snmp help
snmp party-init <if>
snmp re-sync <party-index> <value>
snmp show

snmp v1 [enable|disable]
```

Follow these instructions to configure the SNMPv2 security:

- 1) Create the initial parties using any of the following:
 - A) Use 'snmp default' to clear the data base and create initial parties,
 - B) Use 'snmp party-init' to add a new set of initial parties.
- 2) Determine party index using the 'snmp show' command. Parties with a '*' are not installed because they were just created and/or need to have a private key assigned.
- 3) Assign a private key to each party using any of the following:
 - A) Use 'snmp auth-key h' to enter a hex key,
 - B) Use 'snmp auth-key a' to enter an ASCII key.
- 4) Type 'reboot' to activate the changes.

Figure 2.17 - Help for SNMP Commands

2.7.1 Displaying the Non-volatile Party Database

The `snmp show` command will display the current content of the non-volatile Party database. This command displays the configuration data that will be used to configure the agent's MIB on the next reboot. If changes have been made to the non-volatile database since the last reboot, then the `snmp show display` may not reflect the content of the agent's MIB.

The `snmp show` command presents the display shown in Figure 2.18. This example shows two of the four parties installed in the agent's MIB. Note that the `snmp show` command does not display contexts, ACLs, or views, even though these are present in the database.

```
>> snmp show
Party
Index Auth Priv L/R Party OID
-----
  1 none none Loc 1.3.6.1.6.3.3.1.3.160.86.5.161.1
  2 none none Rem 1.3.6.1.6.3.3.1.3.160.86.5.161.2
 *1 md5 none Loc 1.3.6.1.6.3.3.1.3.160.86.5.161.3
 *2 md5 none Rem 1.3.6.1.6.3.3.1.3.160.86.5.161.4
SNMPv1 is enabled. Agent authentication clock = 4776749
>>
```

Figure 2.18 - Displaying the Non-volatile Party Database

Observe that, in this example, four parties were created for an agent interface with an IP address of 160.86.5.161. This is in accordance with RFC1447.

Examine the entries under the `Party Index` column. The entries without a "*" character denote parties that were installed into the agent's MIB during the preceding reboot. Parties listed with the "*" character denote parties that have not been installed into the agent's MIB which can occur for one of two reasons:

1. The parties have just been created using the `snmp default` or `snmp party-init` commands, or
2. The party is an `md5Auth` party and has not had an authentication key assigned to it yet. An `md5Auth` party will not be installed into the agent's database until an authentication key has been assigned to it.

2.7.2 Configuring Initial Security Parameters

Configuration of the initial parties, contexts, ACLs, and views can be accomplished from a directly connected terminal or via telnet. This process is accomplished using FORE operator console commands. To configure the initial parameters, you must perform three basic steps:

1. Create initial parties, contexts, ACLs and views for each interface that has an IP address assigned.
2. Assign authorization keys for each *md5auth* party created.
3. Reboot the agent.

This procedure is described in detail below.

1. If you are accessing the agent through a terminal interface attached to the management module, use the **if show** command to ensure that an IP address is assigned to each of the agent's interfaces that you want to configure. If the addresses are all shown as 0.0.0.0, refer to the procedure to set the IP addresses on page 2-11.
2. To add a complete set of initial parties, contexts, ACLs, and views to an interface, issue the **snmp party-init if** command, where *if* is the interface you want to configure, as shown in Figure 2.19.

```
>> snmp party-init 1
>>
```

Figure 2.19 - Adding Initial Parties

The **snmp default** command can also be used to delete all entries already in the database and re-initialize the interfaces. This command creates a set of initial parties, contexts, ACLs, and views for each agent interface with a non-zero IP address. From the operator console, this is the only way to delete existing entries in the non-volatile database using the command line interface.

Once SNMPv2 access has been established via an NMS, the initial parties created here can be further modified or deleted using SNMPv2.

3. Determine the Party Index using the **snmp show** command. Parties with a "*" are not installed because they were just created and/or need to have a private key assigned.

- Issue the `snmp auth-key` command to assign authentication keys to each `md5Auth` party created in Step 2. Each interface which had a set of initial parties, contexts, ACLs, and views will have two `md5Auth` parties. There are two formats of the `snmp auth-key` command:

```
snmp auth-key hex party-index hex-key
```

```
snmp auth-key ascii party-index ascii-key
```

where:

- party-index** is the value displayed under the `Party Index` column of the `snmp show` command. Note that Party Indexes shown with an “*” in front of the number must be entered that way (*1, for instance).
- hex-key** is 16 bytes of data, entered in hexadecimal, with each byte separated by a space.
- ascii-key** is an ASCII string of printable characters. Strings containing spaces must be entered using double quotes. Strings with more than 16 characters are truncated. Strings with less than 16 characters are padded with nulls (0x00).

- Reboot the agent using the **RESET** button on the management module’s front panel or by issuing the `reboot` command at the operator console. The agent can only initialize its MIB from the non-volatile storage during the reboot process.

2.7.3 Enabling and Disabling SNMPv1

As mentioned in Chapter 1, the agent allows you to enable and disable the SNMPv1 protocol stack. The capability to accept and process SNMPv2 packets is always enabled on the agent since SNMPv2 access to the agent is controlled through the Party database.

To disable the SNMPv1 protocol stack, enter:

```
snmp v1 disable
```

With SNMPv1 disabled, the agent treats all SNMPv1 packets as if they had a bad community string (i.e. drop the packet and send an authentication failure trap). If you plan to use SNMPv2 network management stations exclusively to control your network and you desire to restrict access to the agent for security reasons, then you may wish to disable SNMPv1.

CAUTION



If SNMPv1 has been disabled you will no longer be able to access the agent using your SNMPv1 network management station.

To enable the SNMPv1 protocol stack, enter

```
snmp v1 enable
```

With SNMPv1 enabled, the agent functions like a normal SNMPv1 agent. The factory default setting has SNMPv1 enabled.

2.7.4 Resynchronizing SNMPv2 Agents and Management Stations

Experience has shown that one of the problems that will be encountered during these early days of SNMPv2 is that agents and managers will occasionally get their clocks out of sync with each other. For example, if a management station never updated its notion of the agent's *partyAuthClock* (perhaps due to a software bug), this would eventually prevent the two parties from communicating. Should this occur, the *partyAuthClock* can be reset through SNMP (provided the party's *partyAuthPrivate* authentication key is changed simultaneously).

The command line interface offers an alternative to the SNMP approach. The `snmp re-sync` command allows you to update the clock value for any *md5Auth* party. The format of the command is:

```
snmp re-sync party-index value
```

where

party-index	is the value displayed under the Party Index column of the <code>snmp show</code> command.
value	is a counter value from 0-4,294,967,295 and is usually set to zero to reset the clock.

The example in Figure 2.20 resets party three's clock to zero. Unlike the SNMP approach, this change only affects the non-volatile party database and not the agent's party MIB; for the change to take effect, the agent must be rebooted.

Management Module Configuration

```
>> snmp show
Party
Index Auth Priv L/R Party OID
-----
  1 none none Loc 1.3.6.1.6.3.3.1.3.192.94.73.12.1
  2 none none Rem 1.3.6.1.6.3.3.1.3.192.94.73.12.2
 *1 md5 none Loc 1.3.6.1.6.3.3.1.3.192.94.73.12.3
 *2 md5 none Rem 1.3.6.1.6.3.3.1.3.192.94.73.12.4
  3 none none Loc 1.3.6.1.6.3.3.1.3.160.86.6.1.1
  4 none none Rem 1.3.6.1.6.3.3.1.3.160.86.6.1.2
 *3 md5 none Loc 1.3.6.1.6.3.3.1.3.160.86.6.1.3
 *4 md5 none Rem 1.3.6.1.6.3.3.1.3.160.86.6.1.4
SNMPv1 is enabled. Agent authentication clock = 2461106
>> snmp re-sync 3 0
>> snmp show
Party
Index Auth Priv L/R Party OID
-----
  1 none none Loc 1.3.6.1.6.3.3.1.3.192.94.73.12.1
  2 none none Rem 1.3.6.1.6.3.3.1.3.192.94.73.12.2
 *1 md5 none Loc 1.3.6.1.6.3.3.1.3.192.94.73.12.3
 *2 md5 none Rem 1.3.6.1.6.3.3.1.3.192.94.73.12.4
  3 none none Loc 1.3.6.1.6.3.3.1.3.160.86.6.1.1
  4 none none Rem 1.3.6.1.6.3.3.1.3.160.86.6.1.2
 *3 md5 none Loc 1.3.6.1.6.3.3.1.3.160.86.6.1.3
 *4 md5 none Rem 1.3.6.1.6.3.3.1.3.160.86.6.1.4
SNMPv1 is enabled. Agent authentication clock = 2461108
>>
```

Figure 2.20 - Resetting a Party's Clock Value

2.8 Rebooting or Resetting the Management Module

The `reboot` command is used to re-initialize the management module. The `reboot` command must be used when changing the IP address or IP network mask. The `reboot` command can not be abbreviated.

CAUTION



Use of this command will erase all volatile data (such as statistics) on the module. Additionally, each management module interface will be re-initialized.

All ES-4810 management modules include compressed code. When the **RESET** button on the front panel is pressed and released, the message `Uncompressing Agent...` will be displayed on the operator console. A delay of approximately 5 to 15 seconds will occur after pressing and releasing the **RESET** button before the agent software comes up.

Management Module Configuration

CHAPTER 3

User Module Configuration

This chapter describes the procedures for configuring ES-4810 user modules using an ES-4810 management module.

3.1 Overview

You can configure and display port information for all ports in a group, for an individual port, or for an entire card. The ES-4810 user modules will operate with factory default settings, but you can modify the default configuration for your network.

The following switch configuration tasks can be performed using console commands:

- Assigning each group of ports to a backplane bus
- Initializing each user port in each group to be either enabled or disabled
- Entering memos for the individual user ports
- Assigning user port priority levels
- Modifying switch port characteristics, including:
 - Duplex mode
 - Speed (100BaseTX ports only)
- Configuring switch monitoring
- Configuring port security
- Configuring VLANs
- Configuring uplink ports
- Configuring redundant port pairs
- Configuring the ATM uplink
- Configuring Spanning Tree on user ports

3.2 Common Configuration Functions

There are common functions that can be performed on all ES-4810 modules, groups, and ports. For each element, you can show status, enable or disable them, reset them or set them to factory default settings stored in NVRAM.

In general, a **reset** command accomplishes three things:

1. It ensures that the hardware is configured as desired (i.e., the configuration matches the one stored in non-volatile memory). This is not really needed if everything is working properly but is done as a preventive measure.
2. If the item being reset supports the notion of a “reset action,” it is performed. The following reset actions apply to switch ports:
 - Learned addresses are freed.
 - Auto-negotiation on 100BaseTX ports is retried.
3. Resetting something resets all subordinate items. For example, resetting a card resets all groups on the card. Resetting a group resets all of the ports in that group.

The **default** commands have a similar effect, in that they are used to reset the component to a default configuration. You may want to use the **default** command as a corrective measure when troubleshooting problems with the chassis components or when reconfiguring a card that has been transferred from another chassis. The **default** command is described in more detail in the following sections.

3.3 Slot and Port Numbering in a Chassis

In the descriptions in this chapter, the slot (card) numbers for the chassis are figured lowest to highest, from left to right, as shown in Figure 3.1. The range of slot numbers is 1 to 12.

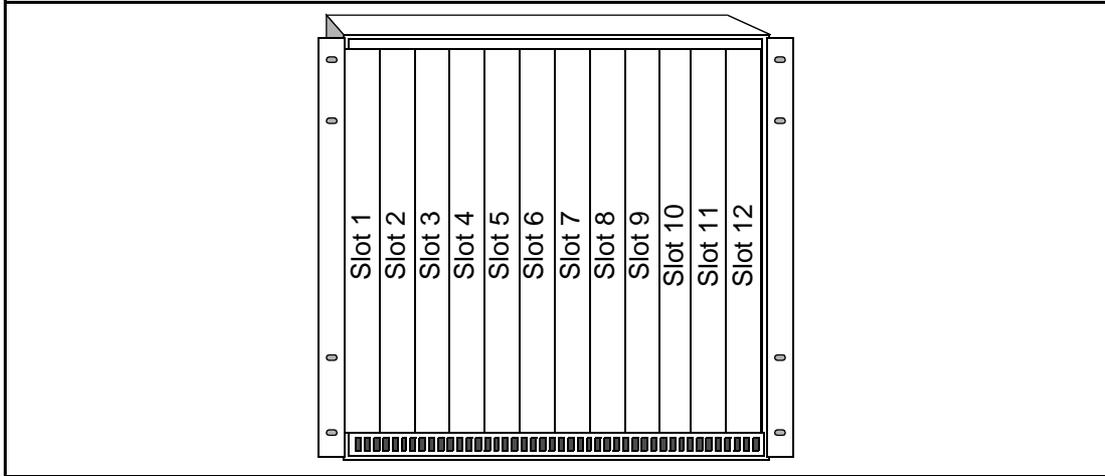


Figure 3.1 - Slot Numbering in a Chassis

The number used to identify a card in a chassis is the same as the number of the slot in which it is installed. The *card* parameter is entered with commands that display or alter information for all ports on the specified module.

The *group* parameter is entered with commands that display or alter information for all ports in the specified group. Group and port numbers, when entered as command parameters, use the format: *card-group* or *card-port*. For example, to specify Group 1 on the module in slot 4, enter: 4-1 as a parameter for a command.

3.4 Configuring Groups

Ports on ES-4810 modules are organized into work groups. You can assign groups to either of the two packet buses in the switching backplane, enable or disable all ports in a group, or reset the entire group to the factory default configuration.

3.5 Assigning Groups to a Packet Bus

Each ES-4810 module has one work group. The group can be assigned to either backplane bus or placed in standalone mode. Figure 3.2 illustrates group configuration on a typical module.

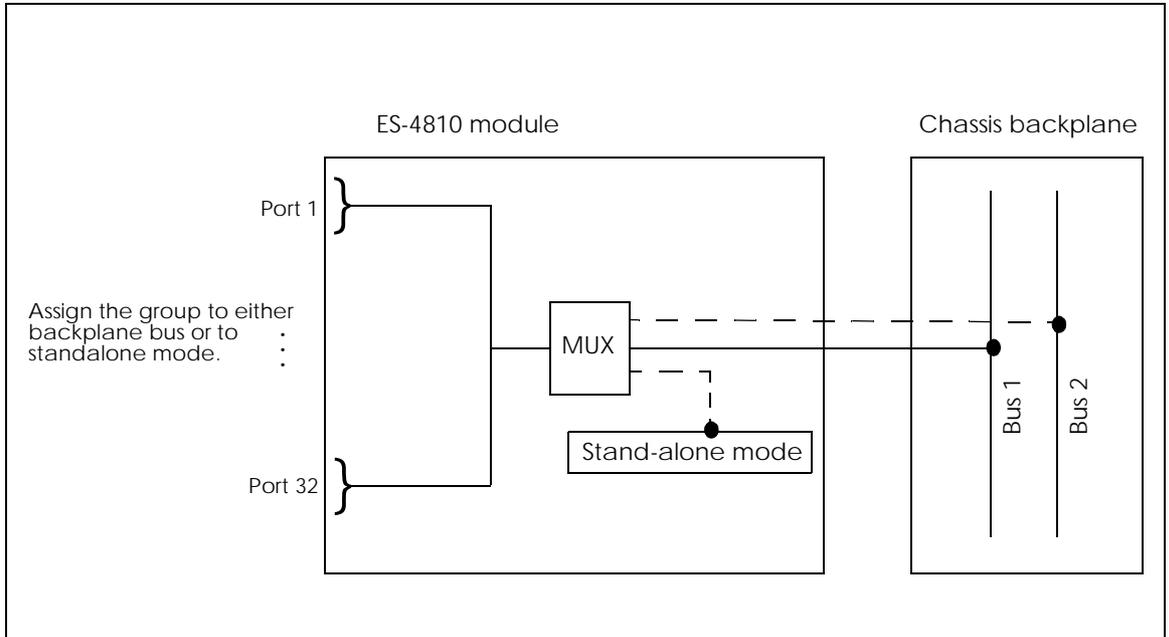


Figure 3.2 - Work Groups on ES-4810 Modules

When in stand-alone mode, the ports in the group can only communicate with other stand-alone ports on the same module.

To display the current group configuration, issue the command:

```
esgroup show [card-group]
```

where:

card is the slot number of the card that has the desired *group*.

group is a valid group number.

An asterisk (*) is a wildcard and may be substituted for *card-group* (all groups on all cards) or *group* (all groups on the specified card).

The example in Figure 3.3 shows group information for all Ethernet switch modules.

```
>> esgroup show
group 3-1: 32 ports, packetbus 1, aging time 0
group 9-1: 12 ports, packetbus 2, aging time 0
>>
```

Figure 3.3 - Displaying Current Group Configuration

The **esgroup show** command lists the following group information for each switch module:

- Number of ports in the group.
- Packet bus to which the group is assigned.
- The aging time for the group. The aging time is the amount of time before inactive learned addresses are deleted from the address database. See “Setting the Aging Time” on page 3-8 for information about setting this parameter.

3.6 Enabling/Disabling Groups of Ports

Groups of ports can be disabled and enabled using the commands:

```
esgroup disable card-group
```

where:

card is the slot number of the card that has the desired *group*.

group is a valid group number.

```
esgroup enable card-group
```

where:

card is the slot number of the card that has the desired *group*.

group is a valid group number.

3.7 Resetting Groups Of Ports

You can also reset groups or set them to the default configuration. To reset a group, issue the command:

```
esgroup reset card-group
```

where:

card is the number of the slot in which the card is installed.

group is a valid group number.

To set the group to the default configuration, issue the command:

```
esgroup default card-group
```

where:

card is the number of the slot in which the card is installed.

group is a valid group number.

3.8 Setting the Aging Time

The aging time specifies how long a learned address remains in a port's address database before being deleted, if that address has not been seen in packets received from the Ethernet segment.



When a NMM-SEG-1 management module is controlling the forwarding on a packet bus (CAM enabled), the aging time is determined by the `escam age`.

This parameter is specified in seconds for all ports in the group. The command format is

```
esgroup age card-group {0|10-1000000}
```

where

- card*** is the slot number of the card that has the desired *group*.
- group*** is a valid group number.
- {0|10-1000000}** is the aging time in seconds, which can be either 0 (aging is disabled) or a number from 10 to 1000000. The default is 300.

The example in Figure 3.4 sets the aging time for group 1 on card 3 to 600 seconds (10 minutes).

```
>> esgroup age 3-1 600
>> esgroup show 3-1
group 3-1: 32 ports, packetbus 1, aging time 600
>>
```

Figure 3.4 - Setting the Aging Time for a Group



To prevent an individual address in a port's address database from being deleted, use the `locked` parameter to the `esport addr` command as described in "Modifying the Address Database For Security" on page 3-25.

3.9 Configuring Ethernet Switch Ports

Ethernet switch user ports need to be enabled to allow users to communicate on the network. To view current port status and configuration, enter the command:

```
esport show [card-port]
```

where:

- card** is the number of the slot in which the card is installed.
- port** is a specific port number or an asterisk (*) for all ports on the card.

You can also specify ports on different cards in the same command line by repeating the *card-port* parameter, or use the wildcard “*” for all ports on all cards.

The command **esport showx** displays additional switch port configuration information.

Figure 3.5 gives an example of **esport show** and **esport showx** command output.

```
>> esport show 9-1
port 9-1: group 1, 100BaseTX, enabled, pri 0, A/100Mbps, dup A/hlf, no-link
>> esport showx 9-1
port 9-1: fwd limited, lrn norm, sniff norm, vlan 1, hist:S-Dis, L-Dis
>>
```

Figure 3.5 - Displaying the Current Ethernet Port Configuration

The **esport show** command displays the following information for each port:

- Memo, if one was assigned
- Group the port is assigned to
- Port type (10BaseT, 10BaseFL, 100BaseTX, or 100BaseFX)
- Enabled or disabled status
- Port priority (from 0 to 15)
- Port speed (speed at which the port is configured to operate)

10BaseT and 10BaseFL port speed is always 10Mbps. 100BaseTX ports can be configured to operate at 10Mbps, 100Mbps, or auto (auto-negotiated values are displayed with an “A/” prefix). 100BaseFX port speed is always 100Mbps.

User Module Configuration

- Duplex mode can be configured as `half`, `full`, `A/hlf`, or `A/ful` (the preceding `A/` indicates the value was auto-negotiate).
- Link status of the port (this field is shown only if there is no link)

The `esport showx` command displays the following information for each port:

- Forwarding mode (see `esport` in Chapter 5)
- Learning mode (`norm` or `sec`)
- Sniff mode (`norm` or `sniff`)
- Vlans this port belongs to (from 1 to 16, or `all`)
- Short history (`Dis` or `En`)
- Long history (`Dis` or `En`)

Modifying configurable Ethernet switch port options is described later in this chapter.

3.10 Enabling/Disabling ports

If any ports being used are not currently enabled, issue the `esport enable` command to enable them. Use the `esport disable` to disable them. The format of these commands is:

```
esport enable card-port
```

```
esport disable card-port
```

The example in Figure 3.6 enables all ports on card 3, then disables port 2.

```
>> esport enable 3-*  
>> esport disable 3-2  
>>
```

Figure 3.6 - Enabling and Disabling Ports

3.11 Enabling/Disabling Spanning Tree on Ports

Spanning tree is an available feature when you have a NMM-SEG-1 management module connected to the same packet bus as the user ports. If all these conditions are met, then spanning tree can be enable or disabled for individual ports.



To enable/disable spanning tree for individual ports, the CAM must be enabled using `escam enable` and spanning tree must be enabled using `stbridge enable`.

To enable spanning tree on a port, issue the command:

```
stport enable card-port
```

where:

- card*** is the number of the slot in which the card is installed or asterisk (*) for all.
- port*** is a valid port number for the card type or asterisk (*) for all.

To disable spanning tree on a port, issue the command:

```
stport disable card-port
```

3.12 Modifying Pathcost on Ports

Path cost is a value used by spanning tree to determine the best route for traffic. The lower path cost is considered the most desirable route. For example a 10 Mbps port has a default path cost of 100 and a 100 Mbps port has a default a cost of 10. If spanning tree needed to choose between these two ports, it would send the information out the 100 Mbps port.

In some cases, other factors will make it more desirable to route down a different path. This can be accomplished by changing the pathcost associated with a individual port.

To change the pathcost enter the command:

```
stport pathcost card-port pathcost
```

where:

<i>card</i>	is the slot in which the card is installed.
<i>port</i>	is the port number.
<i>pathcost</i>	is a value between 1 and 65535.

3.13 Modifying Spanning Tree Priority on Ports

To help ensure that a port becomes the primary port for the bridge, you can set a priority using the `stport priority` command. Using this command, you can specify a value that will be contained in the first two octets of the port's ID. The default is 128. This number will be combined with a port identifier which is used by spanning tree to uniquely identify each port. The range is 0-255 and can be entered in decimal, octal, or hexadecimal form.

To change the port priority enter:

```
esport priority card-port priority
```

where:

- card*** is the number of the slot in which the card is installed. or asterisk (*) for all.
- port*** is a valid port number for the card type or asterisk (*) for all.
- priority*** is a valid priority number.

3.14 Resetting Ports

You can also reset ports or set them to the default configuration. To reset ports, issue the command:

```
esport reset card-port
```

To set a port to the factory default configuration, use the command:

```
esport default card-port
```

3.15 Assigning Individual Ports To Groups

The `esport group` command allows you to assign a switch port to a group. (Currently, switch ports on ES-4810 modules are in one group, but this functionality is provided for future enhancements.) To switch a port from one group to another, enter the command:

```
esport group card-port group
```

where:

- card*** is the number of the slot in which the card is installed.
- port*** is a valid port number for the card type or asterisk (*) for all.
- group*** is a valid group number.

3.16 Specifying Port Memos

Ethernet switch ports also allow you to define a port memo to aid in network management. You can enter a string of up to 31 characters to describe what type of device is connected to the port or where it is located. To define the port memo, use the command:

```
esport memo card-port memo
```

where:

- card*** is the number of the slot in which the card is installed.
- port*** is a valid port number for the card type or asterisk (*) for all.
- memo*** is a string of up to 31 characters enclosed in double quotes.

3.17 Setting Port Priority

You can configure user ports to have a “priority” from 0 to 15. The priority of a port can be used by management applications to intelligently handle “important” ports; for example, require confirmation from the user before disabling the port. Ethernet port priority values can be set via the command line using the following command:

```
esport priority card-port priority
```

where:

- | | |
|------------------------|---|
| <i>card</i> | is the number of the slot in which the card is installed. |
| <i>port</i> | is a valid port number for the card type or asterisk (*) for all. |
| <i>priority</i> | is a priority value from 0 (the default) to 15. |

3.18 Setting Redundant Port Pairs

Most ES-4810 user modules provide redundant backup ports that can be assigned to take over in the event of failure, such as a bad cable. The primary ports and the redundant backup ports are referred to as “redundant port pairs”.

To show a user module’s redundant ports and the current state, issue the command:

```
espair show
```

To configure a redundant port pair, issue the command:

```
espair mode card-pair mode
```

where:

card is the number of the slot in which the card is installed.

pair is the number of the redundant pair.

mode is one of the following:

independent - Configures both ports in the pair to be used as individual user ports with no redundant backup.

redundant - Configures the ports to act as a redundant pair. Under normal circumstances, the primary port is active. If the primary port loses link, the link is switched to the backup port. The link will not switch back until the backup port loses link, or an **espair reset** command is executed on the pair.

autoprimary - Configures the ports to act as a redundant pair. However, when the primary port becomes active again, the link is automatically switched back.

User Module Configuration

```
>> espair show
pair 4-1: primary 1,secondary 2,mode redundant,active primary
pair 4-2: primary 25,secondary 26,mode independent
pair 7-1: primary 1,secondary 2,mode independent
pair 7-2: primary 25,secondary 26,mode independent
pair 9-1: primary 1,secondary 2,mode independent
pair 9-2: primary 25,secondary 26,mode independent
>> espair mode 7-1 redundant
pair 4-1: primary 1,secondary 2,mode redundant,active primary
pair 4-2: primary 25,secondary 26,mode independent
pair 7-1: primary 1,secondary 2,mode redundant,active primary
pair 7-2: primary 25,secondary 26,mode independent
pair 9-1: primary 1,secondary 2,mode independent
pair 9-2: primary 25,secondary 26,mode independent
>>
```

Figure 3.7 - Setting Redundant Port Pairs

3.19 Setting Port Monitoring Mode

As described in “Switch Monitoring” on page 1-9, to monitor network traffic in the Ethernet switch you must configure some ports as monitoring ports and others as monitored ports.

Monitored ports are set using the `esport sniff` command. Ports that can monitor other ports have their forwarding mode set to `sniffer`. A port can not simultaneously be an uplink port and monitor other ports, but an uplink port can be monitored.

To configure a port to be monitored, use the command:

```
esport sniff card-port monitoring_mode
```

where:

card	is the number of the slot in which the card is installed.
port	is a valid port number for the card type or asterisk (*) for all.
monitoring_mode	is the mode, which can be one of the following: <ul style="list-style-type: none"> <code>normal</code> port acts like a regular transparent learning bridge port. This is the default. <code>sniffed</code> In addition to normal packet forwarding, the port labels all transmitted and received packets so that they will also be forwarded by ports on the same packet bus that have their forwarding mode set to <code>sniffer</code>.

To configure a port to be a monitoring port, use the command:

```
esport forwarding card-port sniffer
```

where:

card	is the number of the slot in which the card is installed.
port	is a valid port number for the card type or asterisk (*) for all.
sniffer	is the forwarding mode for the monitoring port. The port will forward all packets that have been labeled by ports in <code>sniffed</code> mode.

User Module Configuration

The example in Figure 3.8 designates port 3-1 as a monitored port, then sets the forwarding mode for port 3-2 to *sniffer*. Port 3-2 will now forward packets labeled by port 3-1.

```
>> esport showx 3-1 3-2
port 3-1: fwd normal, lrn norm, sniff norm, vlan all, hist:S-Dis, L-Dis
port 3-2: fwd normal, lrn norm, sniff norm, vlan all, hist:S-Dis, L-Dis
>> esport sniff 3-1 sniffed
>> esport forwarding 3-2 sniffer
>> esport showx 3-1 3-2
port 3-1: fwd normal, lrn norm, sniff sniff, vlan all, hist:S-Dis, L-Dis
port 3-2: fwd sniffer, lrn norm, sniff norm, vlan all, hist:S-Dis, L-Dis
>>
```

Figure 3.8 - Setting Port Monitoring Mode

3.20 Setting Duplex Mode

ES-4810 10BaseT and 100BaseTX ports can operate in half or full duplex mode. 100BaseTX ports can also be configured to negotiate the duplex mode.

To set the duplex mode for a port, enter:

```
esport duplex card-port duplex_mode
```

where:

- card*** is the number of the slot in which the card is installed, or asterisk (*) for all ports on all cards.
- port*** is a valid port number for the card type or asterisk (*) for all.
- duplex_mode*** is the mode, which can be one of the following:
 - half*** port operates in half duplex mode. This is the default.
 - full*** port operates in full duplex mode.
 - auto*** on 100BaseTX ports only, duplex mode is negotiated by this port and the device attached to the port.



For auto-sensing to work properly in a 100BaseT port, set *speed* and *duplex_mode* to *auto*.

3.21 Setting 100BaseTX Port Speed

ES-4810 100BaseTX ports can operate at either 10Mbps or 100Mbps, or can be set to automatically use the same speed as the external device connected to the port. To see the configured port speed, use the `esport show` command described on 3-9. To see the speed at which the port is currently operating, use the `esport showx` command.

To set the speed for a 100BaseTX port, enter:

```
esport speed card-port speed
```

where:

- card*** is the number of the slot in which the card is installed.
- port*** is a valid port number for the card type or asterisk (*) for all.
- speed*** is the speed, which can be one of the following:
 - 100 100Mbps (default).
 - 10 10Mbps.
 - auto Speed is negotiated by this port and the device attached to the port.



For auto-sensing to work properly in a 100BaseT port, set *speed* and *duplex_mode* to auto.

3.22 Modifying the Address Database For Security

**NOTE**

The following section does not apply to a NMM-SEG-1 management module operating with CAM enabled. In this situation the addresses are stored in a central CAM on the management module and not with each individual port.

Each Ethernet switch port maintains an address database that it uses to make packet filtering and forwarding decisions. In normal operation, the port builds the address database automatically by learning the MAC addresses of the connected devices. You can also modify the address database and the learning process from the command line to make the address database more secure.

**NOTE**

You can not manually modify the address database for uplink ports or ports in sniffer mode because the entries are locked or reserved for use by the hardware. Also, secure learning is not applicable to ports in uplink or sniffer forwarding modes.

3.22.1 Displaying the Address Database



The following section does not apply to a NMM-SEG-1 management module operating with CAM enabled. In this situation the addresses are stored in a central CAM on the management module and not with each individual port.

To display the address database for a port, enter:

```
esport showaddr card-port
```

where:

card is the number of the slot in which the card is installed.

port is a valid port number for the card type or asterisk (*) for all.

Figure 3.9 shows an example of the address database.

```
>> esport showaddr 6-1
6-1-1:      free  00 00 00 00 00 00
6-1-2:      free  00 00 00 00 00 00
6-1-3:      free  00 00 00 00 00 00
6-1-4:      free  00 00 00 00 00 00
>>
```

Figure 3.9 - Displaying the Address Database for a Port

For each port specified, the **esport showaddr** command displays the four database entries, status of each entry (described below), and the associated MAC address.

3.22.2 Restricting Addresses in the Database



The following section does not apply to a NMM-SEG-1 management module operating in CAM mode. In this situation the addresses are stored in a central CAM on the management module and not with each individual port.

Each port's address database has four entries. Instead of allowing the port to learn the addresses in the database, you can specify the addresses to be used and "lock" them so they are not overwritten by new addresses. You can also "reserve" entries to limit the number of addresses in the database used to forward packets and prevent learning.

To lock or reserve an address, use the command:

```
esport addr card-port entry {locked|free|reserved} [address]
```

where:

<i>card</i>	is the number of the slot in which the card is installed, or asterisk (*) for all ports on all cards.
<i>port</i>	is a valid port number for the card type or asterisk (*) for all.
<i>entry</i>	is the number (1 to 4) of the entry in the address database or asterisk (*) for all.
{locked free reserved}	is the status for this entry: <i>locked</i> address is approved for use in the packet forwarding process, and will not be overwritten by another learned address. <i>free</i> the address at this location becomes invalid and the location is available for learning new addresses. <i>reserved</i> the address at this location is not used for forwarding, and this location is unavailable for learning new addresses.
[<i>address</i>]	is a 6-byte Ethernet MAC address in hexadecimal, with each pair of digits separated by a space.

Note that specifying an address is optional. If no address is given, the command applies to the address that is currently at the specified database location. To see the list of addresses that are currently in a port's database, use the **esport showaddr** command, as shown in Figure 3.10.

```
>> esport showaddr 6-1
6-1-1:          free  00 00 00 00 00 00
6-1-2:          free  00 00 00 00 00 00
6-1-3:          free  00 00 00 00 00 00
6-1-4:          free  00 00 00 00 00 00
>> esport addr 6-1 1 locked 00 01 02 03 04 05
>> esport showaddr 6-1
6-1-1:          locked 00 01 02 03 04 05
6-1-2:          free  00 00 00 00 00 00
6-1-3:          free  00 00 00 00 00 00
6-1-4:          free  00 00 00 00 00 00
>>
```

Figure 3.10 - Modifying the Address Database

3.22.3 Setting the Learning Mode



The following section does not apply to a NMM-SEG-1 management module operating with CAM enabled. In this situation the addresses are stored in a central CAM on the management module and not with each individual port.

To make a port's address database more secure, you can configure the port to learn addresses but not use them to forward packets until the addresses are validated. Refer to "Learning Process" on page 1-12 for a description of how addresses are added to the database.

To set the learning mode for a port, use the command:

```
esport learning card-port {normal|secure}
```

where:

- card** is the number of the slot in which the card is installed.
- port** is a valid port number for the card type or asterisk (*) for all.
- normal | secure** is the learning mode:
 - normal** addresses are used in the forwarding process as soon as they are learned. This is the default.
 - secure** the port learns addresses, but does not use the address to forward packets until the address is validated.

To validate an address that was added to the database while in secure learning mode, use the **locked** parameter to the **esport addr** command.

Figure 3.11 shows an example of setting a port in secure learning mode, then validating all the addresses in the database.

```
>> esport learning 9-1 secure
>> esport addr 9-1 * locked
>> esport showaddr 9-1
9-1-1:         locked  00 00 c0 1c 87 17
9-1-2:         locked  11 22 00 00 00 00
9-1-3:         locked  11 22 33 44 55 66
9-1-4:         locked  03 04 05 06 07 08
>>
```

Figure 3.11 - Setting Secure Learning for a Port

3.23 Modifying Forwarding of Unknown Unicasts on End-station Ports

For Ethernet switch ports that are not uplink ports, you can configure the port to filter or flood (forward to all ports) unknown unicast packets. A unicast packet is called “unknown” when its destination address is not recognized by any of the ports in the switch.

Uplink ports can not be configured to filter all unknown unicasts, since the purpose of an uplink port is to forward packets destined for a remote switch (see “Configuring Uplink Ports” on page 3-35). However, by configuring an end-station switch port to filter unknown unicasts, you can reduce unnecessary traffic on the local network. Broadcast packets are still forwarded to support learning on the port.

To modify the forwarding of unknown unicasts on end-station ports, use the command:

```
esport forwarding card-port {normal|limited}
```

where:

- | | |
|-------------------------|---|
| card | is the number of the slot in which the card is installed. |
| port | is a valid port number for the card type or asterisk (*) for all. |
| normal limited | is the forwarding mode:

normal - unknown unicasts are forwarded to all ports. This is the default.

limited - unknown unicasts are filtered. |

The example in Figure 3.12 sets the forwarding mode for all switch ports in the chassis to limited.

```
>> esport for * limited  
>>
```

Figure 3.12 - Setting the Forwarding Mode for End-station Ports

On the NMM-SEG-1 management module with CAM enabled, the forwarding modes operate differently than for other ES-4810 management modules. See Table 5.2 on page 5-23 for available modes.

3.24 Setting Long and Short History

By default, long and short histories are disabled for each port. Short history is configured with 50 buckets and a 30 second interval and long history is configured with 50 buckets and a 30 minute interval.

To modify the long and short histories on end-station ports, use the commands:

```
esport longhist card-port {enable|disable}
esport shorthist card-port {enable|disable}
```

where:

- card** is the number of the slot in which the card is installed or asterisk (*) for all.
- port** is a valid port number for the card type or asterisk (*) for all.
- enable|disable** is the mode:
 - enable turns history on.
 - disable turns history off.

The example in Figure 3.13 sets the long and short history for a single switch port in the chassis to enabled.



Enabling a history does not immediately create a history. A history will be created the next time the management module reboots or is inserted into a chassis.

```
>> esport longhist 9-1 enable
>> esport showx 9-1
port 9-1: fwd normal, lrn norm, sniff norm, vlan 1, hist:S-Dis, L-En

>> esport shorthist 9-1 enable
>> esport showx 9-1
port 9-1: fwd normal, lrn norm, sniff norm, vlan 1, hist:S-En , L-En
```

Figure 3.13 - Setting the History

3.25 Configuring VLANs

Each Ethernet switch port can belong to up to 16 VLANs. By default, ports belong to VLAN 1. By assigning ports to different VLANs, you can create logical groupings of ports to limit the number of broadcasts to all ports. See “Virtual LANs (VLANs)” on page 1-14 for more information about using VLANs.

To modify a port’s VLAN assignment, use the command:

```
esport vlan card-port vlan
```

where:

- card** is the number of the slot in which the card is installed.
- port** is a valid port number for the card type or asterisk (*) for all.
- vlan** is a list of numbers from 1 to 16, in any order and separated by spaces, that specifies the VLAN(s) this port belongs to, or `all` for all 16 VLANs.

The example in Figure 3.14 assigns port 1 on card 9 to VLAN 9.

```
>> esport vlan 9-1 9
>> esport showx
port 6-1: fwd normal, lrn normal, sniff normal, dup half, vlan 9, cur spd 10Mbps
```

Figure 3.14 - Setting the VLAN Assignment for a Port

3.26 Configuring ATM Uplinks



Do not enable an ATM uplink for a packet bus that already has Ethernet switch port-based uplinks. You must disable the port-based uplinks using the `esport forward` command prior to enabling the ATM uplink. You can continue to use port-based uplinks on other packet buses.

To configure the ATM uplink, insert the uplink module into the chassis and attach it to the appropriate packet bus by issuing the command:

```
esgroup bus card-group backplane
```

where:

- card*** is the slot number of the card that has the group.
- group*** is the group of ports to be moved.
- backplane*** is the number of the backplane to which the group will be moved.

The ATM uplink module must be attached to the same packet bus as the ES-4810 management module's management interface. For most management modules, this will be packet bus 1.

The ATM interface will need to be configured as described in the *ES-4810 ATM Uplink User's Manual (MANU0294)*.

Once the uplink module is connected on the packet bus and properly configured, it must be enabled to start collecting and forwarding mac addresses. To enable the uplink module, issue the command:

```
atmuplink enable backplane
```

where:

- backplane*** is the number of the backplane.

The ATM uplink module will now forward unrecognized packets to the uplink. Packets will also be received from the uplink and placed onto the packet bus.



The ATM uplink modules store a maximum of 8192 MAC address. While all addresses will be forwarded to the uplink, only the 8192 addresses will be received and placed on the packet bus.

A maximum of 1024 virtual connections (VCs) are supported on the ATM uplink.
Ports must be members of exactly one VLAN when using the ATM uplink.

3.27 Configuring Uplink Ports

Any ES-4810 port can be configured as an uplink port. An uplink port is used to connect the switch to:

- Another modules in the same chassis on a different packet bus.
- Modules in another ES-4810.
- A non-FORE switch.

100BaseTX and 100BaseFX uplink ports can be configured to forward only unknown unicast packets belonging to the same VLAN as the uplink port. Between ES-4810 modules, 100BaseTX and 100BaseFX ports can be configured to encapsulate VLAN information in the forwarded packets (this function is known as VLAN tagging).

Configuring an uplink port consists of setting the appropriate forwarding mode for the type of switch and port type. (See Table 1.2 on page 1-18 for a summary of switch port forwarding modes.)



The NMM-1 and NMM-2 management modules do not use the Spanning Tree Algorithm, therefore it is possible to create loops or multiple paths in the network that can degrade network performance. To avoid creating loops, refer to “Avoiding Loops in Switch Configuration” on page 1-20 before configuring uplink ports.

When configuring uplink ports for connecting two ES-4810 modules, the uplink ports at each end of the connection should have the same VLAN and forwarding mode configuration.

3.28 Connecting to Another ES-4810 or non-ES-4810

To configure an uplink port for connecting to either a FORE ES-4810 or a non-ES-4810, use the command:

```
esport forwarding card-port {uplink|vlanuplink}
```

where:

- | | |
|--------------------------|---|
| <i>card</i> | is the number of the slot in which the card is installed. |
| <i>port</i> | is a valid port number for the card type or asterisk (*) for all. |
| uplink vlanuplink | is the forwarding mode for the uplink port:
uplink forwards all unknown unicasts.
vlanuplink on 100BaseTX ports only, forwards only those unknown unicasts belonging to the same VLAN as this port. |

3.29 Using VLAN Tagging Between ES-4810 Modules

This option is available only for 100BaseTX ports.

To configure an uplink port to include VLAN information with forwarded packets, use the command:

```
esport forwarding card-port {uplinktag|vlanuplinktag}
```

where:

card	is the number of the slot in which the card is installed.
port	is a valid port number for the card type or asterisk (*) for all.
uplinktag vlanuplinktag	is the forwarding mode: <ul style="list-style-type: none"> uplinktag - Forward all unknown unicasts, and encapsulate VLAN information in the packet. vlanuplinktag - Forward only those unknown unicasts belonging to the same VLAN as this port, and encapsulate VLAN information in the packet.

On the NMM-SEG-1 management module, forwarding operates differently than for other ES-4810 management modules. See Table 5.1 on page 5-22 for available modes.

User Module Configuration

CHAPTER 4

Network Statistics Procedures

After the initial configuration is complete, the management module's agent software collects statistics and status about the other modules in the chassis. The information collected is stored in memory on the management module in variables that are defined by the MIBs. These variables are referred to as MIB objects and you can use console commands or an SNMP or RMON management application to view the values of several MIB objects. You can also configure other MIB objects that control the quantity and type of information collected.

In order to modify the contents of MIB objects or configuration parameters, the community string must be set to a *super* or *read-write* access level.

The agent software on the management module allows you to view and modify the following system parameters:

- System variables that identify the management module and chassis
- Configure RMON table size limits and automatic table creation
- Network statistics collected by the interfaces connected to the packet bus
- Diagnostic information for troubleshooting network problems

4.1 Viewing System Variables

The management module provides reference information useful for network identification and maintenance purposes.

To display general chassis backplane information and power supply status, issue the **chassis show** command. To determine the version of firmware running on the management module, issue the **version** command. These commands are shown in Figure 4.1.

```
>> chassis show
chassis   : FORE Systems ES-4810 chassis
           type 13, rev 1389/1.01, 12 slots
IS clocks: primary up, secondary up
power 1   : FORE Systems ES-4810-PS1 Power Supply
           type 97, status up
power 2   : FORE Systems ES-4810-PS1 Power Supply
           type 97, status down
> version
FORE Systems ES-4810
Version 4.7.3 (Mar 9 1998)
>>
```

Figure 4.1 - Chassis show and version Commands

To display the current settings of the MIB-II System Group variables, issue the **system show** command, as shown in Figure 4.2.

```
>> system show
Name:           stevet-uhub
Contact:        stevet
Location:       steve's office
Current sysUpTime: 7526933
>>
```

Figure 4.2 - System show Command

Commands are provided to allow you to define the System Group variables. To enter character strings that contain blanks, enclose them in double quotes.

To assign a name to this particular management module or chassis, use the following command:

```
system name name
```

where *name* is a character string up to 255 characters in length.

To assign an individual who is responsible for the operation of this module or chassis, issue the command:

```
system contact contact
```

where *contact* is a character string up to 255 characters in length.

Assign a location for this card or chassis by issuing the command:

```
system location location
```

where *location* is a character string up to 255 characters in length.

4.2 RMON Configuration

The agent on the management module provides the capability of collecting and storing a great deal of network information in various tables defined in the RMON MIB. Through the operator console, you can enable or disable the creation of certain RMON tables when the management module is rebooted or limit the size of some of the tables to save the amount of memory being used by the RMON statistics.

To view the current configuration, issue the `rmon show` command as shown in Figure 4.3.

```
>> rmon show
ECAM is Disabled.
      Stats -Short History--  --Long History--  Host Mat  Mac/
I/F  Sts  Int. Buckets Sts  Int. Buckets Sts  Sts  Sts  IP
---  ---  ---
  2   En   30    50 En  1800    50 En  Dis  Dis  Dis
  3   En   30    50 En  1800    50 En  Dis  Dis  Dis
>>
```

Figure 4.3 - Displaying the Current RMON Configuration

This command displays two tables. The first table displays the number of KBytes a group may use and the number in use for the following groups: Stats, History, Hosts, HostTopN, Matrix, MAC-IP, Filter, Capture, Alarm, Events, and Ring. The number of KBytes of memory that can be used by the RMON (not counting any overhead) is also displayed.

The second table displayed shows the current settings of the configurable parameters for each of the RMON tables. At the right side of the table, the column labeled `Mac/IP` shows the enabled or disabled state of the MAC to IP addressing function. This column is shown for all ENC and TRNC agents. However, the last two columns, labeled `RS Sts` and `SR Sts` indicate the state of the Ring Station Statistics and Source Routing Statistics respectively only for TRNC agents.

These parameters can be set from the operator console, as described in the following sections. After modifying any of the RMON configuration parameters, you can use the `rmon default` command to reset the values to the factory defaults.

The `rmon` commands described in this section do not allow you to analyze the statistics collected using RMON. You must have a remote RMON management application to view the statistics collected.

4.2.1 Enabling or Disabling RMON Table Creation

Upon initialization, the ES-4810 management modules allow the creation of an RMON Statistics, long and short History, Hosts, Matrix table and MAC-to-IP mapping table for each of the interfaces. The management modules are shipped from the factory with the Statistics and History table creation enabled. The other tables are disabled by default.

Table 4.1 describes each of the RMON tables.

Table 4.1 - RMON tables

Table	Description
Hosts table	Collects statistics about all hosts on the network, including host address, number of packets and octets sent and received (including multicast and broadcast packets), and other information.
Matrix table	Contains statistics about the source and destination of packets and octets seen on the network.
Statistics table	Collects cumulative statistics about network traffic, especially concerning error conditions such as collisions, dropped packets, CRC alignment errors, burst errors, line errors, and many more.
Long History table	Collects the same information that is collected in the Statistics table, but it is collected in sets called “buckets” over designated time periods. The factory default collection interval for the long table is 30 minutes (1800 seconds).
Short History table	Collects the same information that is collected in the Statistics table, but it is collected in sets called “buckets” over designated time periods. The factory interval for the short history table is 30 seconds.
MAC-IP table	Stores the MAC addresses mapped to the IP addresses of each device on the ring or segment. This table automatically provides this mapping to provide a more familiar identifier for the user.

The collection interval and number of buckets determines how much information can be stored in the History table. These parameters can be configured as explained in “Configuring the Size of the History Tables” on page 4-6.

The following commands are provided to allow you to enable or disable the creation of these RMON tables for each interface (specified by the *if* parameter):

```
rmon host if enable | disable
rmon matrix if enable | disable
```

```
rmon stats if enable | disable
```

```
rmon long if enable | disable
```

```
rmon short if enable | disable
```

```
rmon macip if enable | disable
```

These commands will change the configuration displayed by the `rmon show` command, but will not take effect until the management module is rebooted.

4.2.2 Configuring the Size of the History Tables

The Short and Long History tables collect data over specific time intervals and store it in the table buffer in collection sets called “buckets.” The table buffer is circular and when all of the buckets have been filled, the buffer wraps back to the first bucket and begins overwriting it with new data.

The collection interval and number of buckets in the buffer are configurable. The factory default number of buckets is 50. As an example, if the interval is set to 30 minutes per bucket, this allows for the collection of 25 hours of continuous statistics before it is overwritten. You can modify the intervals and numbers of buckets for each table and each interface using the following commands:

```
rmon long if interval interval
```

```
rmon long if buckets buckets
```

```
rmon short if interval interval
```

```
rmon short if buckets buckets
```

where:

if specifies the interface for which the table is to be configured.

interval specifies the amount of time, in seconds, that statistics collection will occur. The valid range is 1 – 3600 seconds. The factory default value is 30 seconds.

buckets specifies the number of collection sets per *interval* that make up the History table buffer. The valid range for this parameter is 1–65535.

The example in Figure 4.4 configures the size of the Short History table.

```
>> rmon short 2 interval 10
>> rmon short 1 buckets 30
>> rmon show
Stats -Short History--  --Long History--  Host Mat  Mac/ RS  SR
I/F  Sts  Int.  Buckets  Sts  Int.  Buckets  Sts  Sts  Sts  IP  Sts  Sts
---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---
 1  En   30   30  En  1800   50  En  En  En  En  En  En
 2  Dis  10   50  Dis  1800   50  Dis  En  Dis  Dis  Dis  Dis
 3  Dis  30   50  Dis  1800   50  Dis  Dis  Dis  Dis  Dis  Dis
 4  Dis  30   50  Dis  1800   50  Dis  Dis  Dis  Dis  Dis  Dis
>>
```

Figure 4.4 - Configuring the Size of the History Tables

4.2.3 Enabling RMON Statistics Collection

Since the Hosts and Matrix RMON tables are meant to provide information on all hosts and data on the network, these statistics will only be meaningful if all data traffic across an interface is counted. The management (RMON) interfaces must be set to “promiscuous” frame capture mode to collect statistics on all network traffic.



The promiscuous mode can only be set on RMON interfaces and does not apply to agent interfaces.

This mode can be set using the `if mode promiscuous` command described in Chapter 5. You will also need a remote RMON management application to view the statistics collected in the Hosts and Matrix tables.

4.3 Viewing Ethernet Network Statistics

The ES-4810 management module only collects statistics on those packet buses that are assigned to one of the interfaces on the management module.

These network statistics can be viewed using operator console commands or a network management application such as ForeView.

For each group of statistics described in the following sections, an example is shown of the statistics that are displayed. Though the command output is not described in detail in this manual, the RFC or other reference document is cited where you can find detailed descriptions of each of the fields.

4.3.1 Displaying Interface Statistics

The Interface group of statistics consist of data from the *ifTable* from MIB-II (RFC 1213). To display the Interface Group for the management module, enter the command:

```
if stats if
```

where *if* is an interface number or an asterisk (*) for all interfaces.

```
>>if stats 1
if 1:  Rx Octets:          1151296   Tx Octets:          0
      Rx Unicasts:         0       Tx Unicasts:         0
      Rx NonUnicasts:      0       Tx NonUnicasts:     0
      Rx Discards:         0       Tx Discards:         0
      Rx Unknowns:         17989   Tx Errors:          0
      Rx Errors:           0       Tx Errors:          0
>>
```

Figure 4.5 - Displaying Ethernet Interface Statistics

4.4 Diagnostics

The ES-4810 management module provides self-test diagnostics called the Self-Test-Diagnostic (STD) test. In addition, you can check the integrity of all of the components and the compatibility of all configurable options using the **check** command.

4.4.1 Checking the Configuration

To test the current configuration of all of the elements in the chassis, issue the command **check config**. This command checks the configuration of the management module interfaces against the external network and the user modules in the chassis. Any problems are reported to the screen for correction. If no problems are detected, the following message is displayed:

```
No configuration problems found.
```

4.4.2 Self-Test-Diagnostic Test

The Self-Test-Diagnostic (STD) test is only executed on operator request through the use of the **RESET** and **LOAD** buttons on the front panel of the module. The STD test consists of a thorough functionality test of the management module hardware.

To execute the STD test, follow this procedure:

1. Using a pen or other pointed object, press and hold the **LOAD** button while momentarily pressing the **RESET** button.
2. Release the **LOAD** button.

The 7-segment LED displays are used to monitor the test while it is running. A lowercase **b** indicates that a test has begun; the number of the test running is displayed on the other LEDs.

Momentarily press the **RESET** button by itself to return the management module to operational use.



Any card that fails this test should be returned to FORE Systems for repair.

Network Statistics Procedures

CHAPTER 5

Operator Console Command Reference

This chapter is a complete alphabetical reference for all commands used to configure and manage ES-4810 management and user modules.

Commands are available through a terminal connected to the terminal port of the management module. For more information on the uses of the operator console commands, refer to Chapter 2.

Some operator console commands require slot (card) numbers, group numbers or port numbers. Chassis slots are numbered sequentially starting at 1. Group numbers consist of the card and group. Port numbers consist of the card and port. The group or port designator is entered in the format: *card-group* or *card-port*.

All of the following commands have a `help` version that lists all of the options on the command. For brevity, the `help` commands have been listed, but not described.

Most operator console commands can be abbreviated using only enough characters to make it unique (usually 2-4 characters).

5.1 atmuplink

Enables and disables the passing of addresses to the ATM-1/155 or ATM-2/155 uplink modules when installed in the chassis.

5.1.1 Format

```
atm[uplink] h[elp]
atm[uplink] sh[ow] bus
atm[uplink] en[able] bus
atm[uplink] dis[able] bus
```

5.1.2 Description

The `atmuplink` command is used to enable or disable the passing of addresses to the ATM uplink. When disabled, no packets will be passed through the ATM uplink.

5.1.3 Options

bus An Ethernet switch packet bus (1 or 2).

5.1.4 Example

```
>>atmuplink show
Backplane #1:
A ATM uplink card is not attached to this backplane
>>atmuplink enable 1
>>atmuplink show
Backplane #1:
MAC addresses registered with ATM uplink = 105
MAC addresses not registered due to exceeding maximum limit = 0
MAC addresses found with multiple vlans = 0
>>atmuplink disable 1
>>atmuplink show
Backplane #1:
Address updating for the ATM uplink is disabled on this backplane
>>
```

Figure 5.1 - Example of atmuplink show Command

5.2 card

Provides information and control of user modules.

5.2.1 Format

```
ca[rd] h[elp]

ca[rd] sh[ow] [card]

ca[rd] en[able] card

ca[rd] dis[able] card

ca[rd] r[eset] card

ca[rd] de[fault] card

ca[rd] se[rial] card serial

ca[rd] reboot card
```

5.2.2 Description

The **card show** command displays the current state of the card and other identification information. You can enable and disable all ports on a module using the **card enable** and **card disable** commands.

To reset the card's hardware configuration to the configuration stored in non-volatile memory, use the **card reset** command. This command also causes ports to switch to the primary port of a redundant port pair, if it has link. The **card reset** command also causes all groups on the specified card to be reset.

The **card default** command resets the card's configuration to the factory defaults.

The **card reboot** command reboots the user card as if it had been removed and re-inserted into the chassis.

5.2.3 Options

- card** Specifies the card in the chassis. If no *card* parameter is specified, a summary of all cards is listed. The valid range of card numbers depends on the type of chassis and the number of slots in the chassis.
- serial** Specifies the serial number (printed on the card) for a module.

5.2.4 Example

```
>> card show 3
card 3: FORE Systems ESM 10BaseFL Ethernet Card
      type 4, rev 1.16, width 1, sn
Ethernet: 4 groups, 18 ports, 2 pairs
group 3-1: 4 ports, segment 1
group 3-2: 4 ports, segment 1
group 3-3: 5 ports, segment 1
group 3-4: 5 ports, segment 4
pair 3-1: primary 12, secondary 13, mode independent
pair 3-2: primary 17, secondary 18, mode independent
>>
```

Figure 5.2 - Example of card show Command

5.3 chassis

Shows the current configuration of the chassis.

5.3.1 Format

```
ch[chassis] h[elp]
```

```
ch[chassis] s[how]
```

5.3.2 Description

The **chassis show** command provides management information about the chassis in which the management module is installed. The chassis type, the hardware revision, and number of slots in the chassis are listed. Each power supply part number is listed followed by the status (up or down). The *type* parameter represents the chassis type.

5.3.3 Example

```
>> chassis show
chassis   : FORE Systems ES-4810 chassis
           type 13, rev 1389/1.01, 12 slots
IS clocks: primary up, secondary up
power 1   : FORE Systems ES-4810-PS1 Power Supply
           type 97, status up
power 2   : FORE Systems ES-4810-PS1 Power Supply
           type 97, status down
>>
```

Figure 5.3 - Example of chassis show Command

5.3.4 See Also

version, system

5.4 check

Analyzes the current configuration of the chassis.

5.4.1 Format

```
che[ck] h[elp]
```

```
che[ck] c[onfig]
```

5.4.2 Description

The **check** command performs a diagnostic test of the configuration of the management modules and user modules in the chassis. Any configuration anomalies will be listed, or a message stating that there were no problems.

5.4.3 Examples

If you have not assigned a *super*-level community string:

```
>> check config
warning - It is recommended that you have at least one
community name with super user privileges.
USE THE community COMMAND.
>>
```

Figure 5.4 - Example of check config Command

If more than one community string has the same name:

```
>> check config
error: You should NOT have more than one entry
      in the community table with the SAME NAME.
      USE THE community COMMAND.
entry 1:  ip address 0.0.0.0, privilege: rwrite  name: super
entry 2:  ip address 0.0.0.0, privilege: ronly  name: public
entry 3:  ip address 0.0.0.0, privilege: ronly  name: public
entry 4:  ip address 0.0.0.0, privilege: noacc  name:
entry 5:  ip address 0.0.0.0, privilege: noacc  name:
entry 6:  ip address 0.0.0.0, privilege: noacc  name:
entry 7:  ip address 0.0.0.0, privilege: noacc  name:
entry 8:  ip address 0.0.0.0, privilege: noacc  name:
>>
```

Figure 5.5 - Example of check config Command

5.5 community

Allows modification of community strings that provide access privileges to MIB data.

5.5.1 Format

```

co[mmunity] h[elp]

co[mmunity] s[how] entry

co[mmunity] n[ame] entry name

co[mmunity] p[riv] entry privilege

co[mmunity] i[p] entry ipaddress

```

5.5.2 Description

The **community** command defines five different levels of access to the SNMP MIB data on the management module agent. Up to eight community strings can be entered in the community table stored in non-volatile memory.

The community strings are defined using the **community name** command. The **community priv** command lets you set the string specified by the *community* parameter to one of five access privileges.

The access privilege associated with a particular community string can be further limited to a single IP address using the **community ip** command. By assigning an IP address to a community string, only accesses originating from that IP address will be allowed.

5.5.3 Options

- entry** Specifies the table entry (1-8) that is being updated.
- name** Specifies the community string to be assigned to a privilege level. The *name* is case-sensitive, and may consist of up to 15 characters, though strings with spaces in them must be enclosed in double quotes. To enter an empty string, use "" or omit the *name* parameter.

privilege Specifies the privilege level to be assigned to the specified table entry. The five levels that can be specified are:

super Allows read and write access to all MIB information and software updates via TFTP. Also, any community string defined with the super access level can be used as a login ID on the operator console.

ronly Provides read only access to most MIB data.

rwrite Provides read and write access to most MIB data.

v2rwrite Allows an agent running SNMPv1 to access the SNMPv2 MIB objects in RFCs 1447 and 1450.

noacc Specifically prevents access to MIB data for community strings with this level access.

limonly Provides a limited read only access to the MIB. Only the MIB-II System group objects can be read.

ipaddress Specifies the IP address, in standard dot notation, that the access privilege is restricted to.

5.5.4 Example

```
>> community show
entry 1:  ip address 0.0.0.0, privilege: super    name: super
entry 2:  ip address 0.0.0.0, privilege: ronly   name: public
entry 3:  ip address 47.2.3.4, privilege: ronly  name: anystring
entry 4:  ip address 0.0.0.0, privilege: noacc   name:
entry 5:  ip address 0.0.0.0, privilege: noacc   name:
entry 6:  ip address 0.0.0.0, privilege: noacc   name:
entry 7:  ip address 0.0.0.0, privilege: noacc   name:
entry 8:  ip address 0.0.0.0, privilege: noacc   name:
>>
```

Figure 5.6 - Example of community show Command

5.5.5 See Also

`trap`

5.6 esbpbus

Assign and display a memo to an Ethernet switch backplane bus.

5.6.1 Format

```
esb[pbus] h[elp]
```

```
esb[pbus] s[how]
```

```
esb[pbus] mem[o] bus memo
```

5.6.2 Description

The `esbpbus show` command displays the memo, if any, for each backplane bus. To assign a 31-character memo, use the `esbpbus memo` command.

5.6.3 Options

- bus** An switch backplane bus (1 or 2).
- memo** A 31 ASCII-character memo enclosed in double quotes.

5.6.4 Example

```
>>esbpbus show
Backplane Switch Bus 1 marketing
Backplane Switch Bus 2 support
>> esbpbus memo 1 "engineering"
>> esbpbus show
Backplane Switch Bus 1 engineering
Backplane Switch Bus 2 support
```

Figure 5.7 - Example of esbpbus Command

5.7 **escam**

Provides information and control of the Content Addressable Memory (CAM) database on NMM-SEG-1 management modules.

5.7.1 **Format**

```
esc[am] h[elp]

esc[am] s[how]

esc[am] en[able]

esc[am] dis[able]

esc[am] re[set]

esc[am] d[efault]

esc[am] age {10-1000000}

esc[am] co[unters]
```

5.7.2 **Description**

The **escam show** command displays the current state of CAM database. The **escam counters** command displays just the events portion of the **escam show** command which reports the number of learned addresses and the number of address that were discarded because the CAM was completely full.

You can enable and disable the CAM using the **escam enable** and **escam disable** commands. By default, the CAM will be disabled and will need to be enabled to begin the learning of addresses and forwarding of packets.

To reset the CAM database to the configuration stored in non-volatile memory, use the **escam reset** command. Resetting the CAM will purge all addresses from the CAM database and restore the CAM settings from NVRAM.

To set the CAM database back to its default setting, issue the command **escam default**. Setting the CAM database to default will purge all addresses from the CAM database and disable the CAM. The aging time will be set back to 300 seconds.

The aging time specifies how long a learned address remains in the CAM database before being deleted. The address will only be deleted if that address has not been seen by the CAM during the aging time period.

5.7.3 Options

time Aging time must be between 10-1000000 seconds and cannot be disabled.

5.7.4 Example

```
>> esc show
CAM STATUS
-----
CAM Admin: enabled
CAM Oper: enabled
CAM Bus: 2
CAM Size: 8192 entries
CAM Learned Events: 0
CAM Learned Invalid Events: 0
CAM Learned Discard Events: 0
CAM Age: 1000000
>> esc counters
CAM COUNTERS
-----
CAM Learned Events: 1230
CAM Learned Invalid Events: 0
CAM Learned Discard Events: 0
>>
```

Figure 5.8 - Example of esc Command

5.8 esgroup

Configure switched Ethernet port work groups.

5.8.1 Format

```

    esg[roup] h[elp]

    esg[roup] s[how] [card-group]

    esg[roup] bus card-group bus

    esg[roup] en[able] card-group

    esg[roup] dis[able] card-group

    esg[roup] r[eset] card-group

    esg[roup] d[efault] card-group

    esg[roup] age card-group {0|10-1000000}
  
```

5.8.2 Description

The **esgroup show** command displays the current configuration of the specified group of switched Ethernet ports. To assign a group of ports to one of the packet buses, issue the **esgroup bus** command. You can enable and disable a group of ports using the **esgroup enable** and **esgroup disable** commands.

To reset the configuration of all ports in a group to the configuration stored in non-volatile memory, use the **esgroup reset** command. The **esgroup default** command resets the configuration of the ports in the group to the factory defaults.

The **esgroup age** command sets the aging time. The aging time specifies how long a learned address remains in a port's address database before being deleted, if that address has not been seen in packets received from the Ethernet segment.

5.8.3 Options

card is the slot number of the card that has the desired *group*.

- group** is a valid group number.
- bus** is the backplane bus number (1 or 2), or stand-alone for stand-alone mode.
- {0|10-1000000}** is the aging time in seconds, which can be either 0 (aging is disabled) or a number from 10 to 1000000.

5.8.4 Example

```
>> esgroup show
group 3-1: 32 ports, packetbus 1, aging time 0
group 9-1: 12 ports, packetbus 2, aging time 100, CAM capable
>>
```

Figure 5.9 - Example of esgroup Command

In this example card 3 is not capable of operating on a packet bus with the NMM-SEG-1 when CAM is enabled. Card 9 is CAM capable and would operate properly.

5.9 espair

Configure switched redundant port pairs.

5.9.1 Format

```

    espa[ir] h[elp]

    espa[ir] s[how] [card-pair]

    espa[ir] r[eset] card-pair

    espa[ir] d[efault] card-pair

    espa[ir] m[ode] card-pair mode
  
```

5.9.2 Description

Most ES-4810 cards provide redundant backup ports that can be assigned to take over in the event of failure of another port. The primary ports and redundant backup ports are referred to as “redundant port pairs”. On the ESM-24/FEM-2, ports 1 and 2 make up one redundant port pair and ports 25 and 26 make up the second.

The **espair mode** command allows the backup ports to operate in two different redundant modes: *redundant* or *autoprimary*. Each port in a pair can also be configured to operate individually by setting the mode to *independent*.

5.9.3 Options

- card** specifies the slot number in which the card is installed.
- pair** specifies the pair number for the ports.
- mode** is the redundancy mode for the pair:
 - independent* Configures both ports in the pair to be used as individual user ports with no redundant backup.

`redundant` Configures the ports to act as a redundant pair. Under normal circumstances, the primary port is active. If the primary port loses link, the link is switched to the backup port. Link is not switched back unless link is lost to the backup port.

`autoprimary` Configures the ports to act as a redundant pair. However, if the backup port has had the link switched to it, when the primary port becomes active again, the link is automatically switched back.

5.9.4 Example

```
>>espair show
pair 9-1: primary 1,secondary 2,mode independent
pair 9-2: primary 25,secondary 26,mode independent
>>espair mode 9-2 autoprimary
>>espair show
pair 9-1: primary 1,secondary 2,mode independent
pair 9-2: primary 25,secondary 26,mode autoprimary,active neither
>>
```

Figure 5.10 - Example of espair Command

5.9.5 See Also

`esport`, `esgroup`

5.10 esport

Configure switched Ethernet ports.

5.10.1 Format

```

    esp[ort] h[elp]

    esp[ort] s[how] card-port
    esp[ort] showx card-port
    esp[ort] camshow card-port
    esp[ort] en[able] card-port
    esp[ort] dis[able] card-port
    esp[ort] r[eset] card-port
    esp[ort] d[efault] card-port
    esp[ort] gr[oup] card-port group
    esp[ort] mem[o] card-port memo
    esp[ort] pri[ority] card-port priority
    esp[ort] sniff card-port monitoring_mode
    esp[ort] duplex card-port duplex_mode
    esp[ort] speed card-port speed
    esp[ort] showaddr card-port
    esp[ort] addr card-port entry status [address]
    esp[ort] learning card-port learning_mode
    esp[ort] vlan card-port vlan
  
```

```
esp[ort] forwarding card-port forwarding_mode
```

```
esp[ort] shorthist card-port {enable|disable}
```

```
esp[ort] longhist card-port {enable|disable}
```

5.10.2 Description

The **esport** command is the Ethernet switch-specific port configuration command. You can enable and disable ports, show the current setting, reset them or set them to the factory defaults, and specify a port memo and port priority. You can also set extended switch port attributes such as monitoring mode, duplex mode, security, forwarding mode, and VLAN information. On 100BaseTX ports, you can specify the port speed.

You cannot manually modify the address database for uplink ports or ports in sniffer mode because the entries are locked or reserved for use by the hardware. Also, secure learning is not applicable to ports in these forwarding modes.

The **esport show** command displays basic switch port information; the **esport showx** command displays extended switch port information. When using a NMM-SEG-1 management module, you can also display the addresses learned by the CAM using the **esport camshow** command.

5.10.3 Options

card	is the slot number of the card that has the desired <i>port</i> or asterisk (*) for all cards.
port	is a valid port number or asterisk (*) for all ports.
group	is a valid group number or asterisk (*) for all groups.
memo	is a string of up to 31 characters enclosed in double quotes.
priority	is a priority value from 0 (the default) to 15.
monitoring_mode	is the port monitoring mode, which can be one of the following: normal port acts like a regular transparent learning bridge port. This is the default.

- sniffed* In addition to normal packet forwarding, the port labels all transmitted and received packets so that they will also be forwarded by ports on the same packet bus that are configured in *sniffer* mode.
- duplex_mode*** is the duplex mode, which can be one of the following:
- half* port operates in half duplex mode. This is the default.
 - full* port operates in full duplex mode.
 - auto* on 100BaseTX ports only, duplex mode is negotiated by this port and the device attached to this port.
- speed*** is the speed for 100BaseTX ports, which can be one of the following:
- 100* 100Mbps (default).
 - 10* 10Mbps.
 - auto* Speed is negotiated by this port and the device attached to the port.
- entry*** is the number (1 to 4) of the entry in the address database or asterisk (*) for all.
- status*** is the status for this entry:
- locked* address is approved for use in the packet forwarding process, and will not be overwritten by another learned address.
 - free* the address at this location becomes invalid and the location is available for learning new addresses.
 - reserved* the address at this location is not used for forwarding, and this location is unavailable for learning new addresses.
- [*address*]** is a 6-byte Ethernet MAC address in hexadecimal, with each pair of digits separated by a space.

learning_mode is the learning mode:

normal addresses are used in the forwarding process as soon as they are learned. This is the default.

secure the port learns addresses, but does not use the address to forward packets until the address is validated.

vlan is a list of numbers from 1 to 16, in any order and separated by spaces, that specifies the VLAN(s) this port belongs to, or *all* for all 16 VLANs. A value of *none* can be used if the port does not belong to a VLAN.

forwarding_mode for all management modules except a CAM enabled NMM-SEG-1 the *forwarding_mode* is one the forwarding modes listed in Table 5.1. When the CAM is enabled the forwarding modes are listed in Table 5-2.

Table 5.1 - Forwarding Modes for Ethernet Switch Ports

Port type	Forwarding mode	Description ¹	Notes
non-uplink	<i>normal</i>	Forward all unknown unicasts.	Default forwarding mode.
	<i>limited</i>	Filter all unknown unicasts.	Uplink ports cannot filter all unknown unicasts.
	<i>sniffer</i>	Forward packets that have been labeled by ports in <i>sniffed</i> mode.	Uplink ports cannot monitor other ports, but they can be monitored.
uplink	<i>uplink</i>	Forward all unknown unicasts.	Secure learning and modifying the address database does not apply to uplink ports.

Table 5.1 - Forwarding Modes for Ethernet Switch Ports

Port type	Forwarding mode	Description ¹	Notes
uplink, 100BaseTX and 100BaseFX only	uplinktag	Forward all unknown unicasts, and encapsulate VLAN information in the packet.	For use between ES-4810s only. Secure learning and modifying the address database does not apply to uplink ports.
	vlanuplink	Forward only those unknown unicasts belonging to the same VLAN as this port.	Secure learning and modifying the address database does not apply to uplink ports.
	vlanuplinktag	Forward only those unknown unicasts belonging to the same VLAN as this port, and encapsulate VLAN information in the packet.	For use between ES-4810s only. Secure learning and modifying the address database does not apply to uplink ports.

¹. Multicasts and broadcasts are forwarded according to the VLAN assignment of the port except in sniffer mode.

Table 5.2 - Forwarding Modes for Ethernet Switch Ports (NMM-SEG-1 in CAM Mode)

Forwarding modes	Description
normal, limited, uplink, vlanuplink	Forward all unknown unicast, broadcast, and multicast packets to all ports in the same VLAN as the source port.
uplinktag, vlanuplinktag	Forward all unknown unicast, broadcast, and multicast packets to all ports in the same VLAN as the source port and encapsulate VLAN information in the packet.
sniffer	Forward packets that have been labeled by ports in sniffed mode.

5.10.4 Example

```
>> esport show 9-1  
port 9-1: group 1, 100BaseTX, enabled, pri 0, A/100Mbs, A/ful  
>> esport showx 9-1  
port 9-1: fwd limited, lrn norm, sniff norm, vlan all, hist:S-Dis, L-Dis  
>>
```

Figure 5.11 - Example of esport Command

5.11 group

Displays and modifies the configuration of groups of ports.

5.11.1 Format

```

    gr[oup] h[elp]

    gr[oup] sh[ow] [card-group]

    gr[oup] en[able] card-group

    gr[oup] dis[able] card-group

    gr[oup] r[eset] card-group

    gr[oup] d[efault] card-group

    gr[oup] seg[ment] card-group segment
  
```

5.11.2 Description

The **group show** command displays the current configuration of the specified group of ports. To assign a group of ports to the Ethernet backplane segment, issue the **group segment** command. You can enable and disable a group using the **group enable** and **group disable** commands.

To reset the configuration of all ports in a group to the configuration stored in non-volatile memory, use the **group reset** command. The **group default** command resets the configuration of the ports in the group to the factory defaults.

5.11.3 Options

card	The slot in the chassis in which the card is installed.
group	The work group on the card or an asterisk (*) for all groups on the card.
segment	The backplane segment to which you want the group assigned (1-4). You can also specify <i>stand-alone</i> or the on-board segments 1-4 (<i>on1-on4</i>).

5.11.4 See Also

`port`

5.12 if

Allows configuration of a management module's interfaces to the backplane or packet bus.

5.12.1 Format

```
if h[elp]

if s[how] [if]

if addr[ess] if address

if dis[able] if

if en[able] if

if mo[de] if mode

if net[mask] if netmask

if seg[ment] if segment

if st[ats] if

if vlan if vlan
```

5.12.2 Description

The **if** (interface) command provides several configuration functions. The interfaces on the management module allow the agent software to monitor and control the packet buses they are assigned to and to communicate with devices on the network.

After initial installation of the management module in a chassis, a network IP address and subnetwork mask must be defined for each interface using the **if address** and **if net-mask** commands. The IP address allows communication with other devices on the network. However, without an IP address assigned, an interface can still monitor packet buses and collect statistics. The netmask is used to specify which part of the IP address designates the network and which part specifies nodes on the network. After modifying these addresses, the management module must be rebooted for the addresses to be applied.

Next, each interface must be assigned to a packet bus using the `if segment` command. The interfaces to the packet bus can be enabled or disabled using the `if enable` and `if disable` commands.

In order to provide detailed statistics for the RMON functions of the management module, you would need to configure the interfaces to accept all network traffic. Each interface has two modes that determine which packets contribute to the statistics it collects: *normal* or *promiscuous*. The `if mode` command configures the interface to operate in “normal” or “promiscuous” mode. In promiscuous mode, all frames are copied from the packet bus.

5.12.3 Options

<i>if</i>	Specifies the interface (1-5) to be configured or displayed.
<i>address</i>	Specifies the IP address for an interface. The <i>address</i> parameter must be entered in standard dot notation and the address cannot be assigned to any other interface. To disable communication with the agent software over the network, use the IP address 0.0.0.0. This will prohibit SNMP management access to the hub, but will not interfere with statistics collection or access using the operator console through the terminal port.
<i>mode</i>	Specifies the frame acceptance mode for the interface: <i>normal</i> The interface only accepts traffic specifically addressed to it, including broadcast frames. <i>promiscuous</i> The interface accepts every packet of data present on the network. This does not apply to agent interfaces.
<i>netmask</i>	Specifies the subnet network mask for an interface. The <i>netmask</i> parameter must be entered in standard dot notation.
<i>state</i>	Specifies the participation state for active monitor selection: <i>true</i> or <i>false</i> .
<i>segment</i>	Specifies the backplane ring or segment to assign the interface to (1-4).
<i>vlan</i>	A comma delimited list of VLANs to which the interface belongs and will receives traffic.

5.13 port

Sets and displays information for the ES-4810 Ethernet ports.

5.13.1 Format

```
p[ort] h[elp]

p[ort] s[how] [card-port]

p[ort] r[eset] card-port

p[ort] d[efault] card-port

p[ort] gr[oup] card-port group

p[ort] en[able] card-port

p[ort] dis[able] card-port

p[ort] mem[o] card-port memo

p[ort] pri[ority] card-port priority
```

5.13.2 Description

The **port** command allows you to enable and disable ports, show the current setting, reset them, set them to the factory defaults or define a priority that can be used by network management applications. You can also label each port with a 31-character port memo to aid network administration. The memos you assign are displayed with all port level commands.

The **port group** command applies to all Ethernet modules that have 10BaseT ports that are individually switchable to the groups and on-board segments on the cards. Any combination of up to 13 ports on a card can be assigned to a single group. Ethernet modules with 12 10BaseT ports have four on-board segments. Ethernet modules with 24 10BaseT ports have eight on-board segments. Note that the group-to-segment assignments permitted on these cards are restricted.

5.13.3 Options

<i>card</i>	Specifies the number of the slot in which the card is installed.
<i>port</i>	Specifies a valid port number for the card type or asterisk (*) for all.
<i>memo</i>	Specifies a character string of up to 31 characters. If spaces are included in the string, it must be enclosed in double quotes.
<i>priority</i>	A number (0-15) that specifies a port's relative importance within the network. This priority does not affect insertion onto the backplane segment or ring (or any operator console command processing).
<i>group</i>	The group (1-4) to which the port is to be assigned.

5.13.4 See Also

`group`, `card`, `esport`

5.14 rmon

Displays the configuration, resets defaults and sets the size limits for RMON tables.

5.14.1 Format

```
rm[on] h[elp]

rm[on] s[how]

rm[on] d[efault]
```

5.14.2 Description

The **rmon** commands are used to configure the collection of statistics using the Remote Network Monitoring (RMON) function of the management module. This command displays all RMON command formats (**rmon help**), shows the current configuration (**rmon show**), and lets you reset the configuration to the factory default (**rmon default**).

The **rmon show** command displays two tables. The first table lists the number of kilobytes a group may use and the number in use for the following groups: stats, History, Hosts, Host-Topn, Matrix, MAC-IP, Filter, Capture, Alarm, Events, and Ring. Also displayed is the number of kilobytes of memory the RMON code may use, not counting any overhead.

The second table displayed is the same as in the prior release with the following additions: 1) a Mac/IP column for all ENC and TRNC agents and 2) Ring Station Statistics and Source Routing Statistics columns only for TRNC agents.

Note that, to view the RMON statistics collected on the management module, you must use a remote RMON management application.

Changes to the RMON configuration parameters take effect after you reboot the management module. Issuing the **rmon default** command resets the values of each command parameter to the default value (defaults are listed in the “Options” section for each command parameter).

5.14.3 Example

```
>> rmon show
Category      Limit      Used      Category      Limit      Used
-----      -
rmon          507k      212k      macip          none        5k
stats         none       1k        filter         none        8k
history       none       36k      capture        none       124k
hosts         none       10k      alarm          none         0k
htopn         none         0k      event          none         0k
matrix        none       12k      ring           none        18k

      Stats -Short History-- --Long History-- Host Mat Mac/ RS  SR
I/F  Sts  Int. Buckets Sts  Int. Buckets Sts  Sts  Sts  IP  Sts  Sts
---  ---  -
1    En   30    50  En  1800    50  En  En  En  En  En  En
2    Dis  30    50  Dis 1800    50  Dis Dis Dis Dis Dis Dis
3    Dis  30    50  Dis 1800    50  Dis Dis Dis Dis Dis Dis
4    Dis  30    50  Dis 1800    50  Dis Dis Dis Dis Dis Dis
>>
```

Figure 5.12 - Example of rmon Command

5.14.4 See Also

rmon host, rmon long, rmon matrix, rmon short, rmon stats, rmon limit

5.15 rmon host

Enables or disables the creation of the RMON Hosts table at management module initialization.

5.15.1 Format

```
rm[on] host if en[able]
```

```
rm[on] host if dis[able]
```

5.15.2 Description

The `rmon host` command can be used to disable or enable the creation of the RMON Hosts table for a specified interface when the management module is rebooted. By default from the factory, the creation of the Hosts table is disabled.

When enabled, one Hosts table is created for each interface specified by the `if` parameter. The Hosts table contains statistics about all of the hosts on the network, including the host address, number of packets and octets sent and received (including multicast and broadcast packets), and other information. Refer to RFC 1271 for a more detailed description of the RMON Hosts table.

Hosts table statistics will only be collected for the interfaces that are configured in “promiscuous” frame capture mode (see `if mode`). You will also need a remote RMON management application to view the statistics collected in the Hosts table.

After enabling or disabling Hosts table creation, the change will appear to have taken effect as displayed by the `rmon show` command. However, changes to the RMON configuration take effect only after you reboot the management module.

In general, the RMON functions provided by the agent place an additional burden on the agent (in terms of memory usage and CPU cycles). To maximize the utilization of the memory and CPU, RMON tables should be created sparingly (i.e., only created when needed).

5.15.3 Example

```

>> rmon show
Category    Limit    Used    Category    Limit    Used
-----    -
rmon        507k    212k    macip        none     5k
stats       none     1k      filter       none     8k
history     none     36k    capture      none    124k
hosts       none     10k    alarm        none     0k
htopn      none     0k     event        none     0k
matrix      none     12k    ring         none    18k

      Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1 En 30 50 En 1800 50 En En En En En En
2 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
3 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
4 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis

>> rmon host 2 enable
Reboot the agent for changes to take affect.
(after the reboot...)
>> rmon show
Category    Limit    Used    Category    Limit    Used
-----    -
rmon        507k    212k    macip        none     5k
stats       none     1k      filter       none     8k
history     none     36k    capture      none    124k
hosts       none     10k    alarm        none     0k
htopn      none     0k     event        none     0k
matrix      none     12k    ring         none    18k

      Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1 En 30 50 En 1800 50 En En En En En En
2 Dis 30 50 Dis 1800 50 Dis En Dis Dis Dis Dis
3 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
4 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis

>>

```

Figure 5.13 - Example of rmon host Command

5.15.4 See Also

`rmon`, `rmon long`, `rmon matrix`, `rmon stats`, `rmon short`, `rmon limit`

5.16 rmon long

Enables or disables the creation of the long RMON History table at management module initialization.

5.16.1 Synopsis

```
rm[on] long if en[able]

rm[on] long if dis[able]

rm[on] long if interval interval

rm[on] long if buckets buckets
```

5.16.2 Description

The **rmon long** command can be used to disable or enable the creation of a “long” version of the RMON History table for a specified interface when the management module is rebooted. By factory default, the creation of the History tables is enabled. The long History table differs from the “short” History table in that the default statistics collection interval for the long table is 30 minutes (1800 seconds), whereas it is 30 seconds for the short table.

The History table contains detailed statistics about network traffic over designated time periods, especially concerning error conditions such as collisions, dropped packets, CRC alignment errors, burst errors, line errors, and many more. Refer to RFC 1271 and RFC 1513 for a more detailed description of the RMON History tables.

Using the **interval** and **buckets** forms of the command, you can configure the quantity of data that will be collected in the table before it is overwritten. The **interval** parameter specifies the collection time period and the **buckets** parameter specifies the number of collection sets in the History table buffer. The buffer is circular and when all buckets have been filled, the buffer wraps back to the first bucket and begins overwriting it with new data.

On the ES-4810 management module, the interfaces must also be configured in “promiscuous” frame copy mode to collect data for the History tables. This can be configured on the management module using the **if mode promiscuous** command.

On the ES-4810 management modules, you must have a remote RMON management application to view the statistics collected. After enabling or disabling table creation, the change will appear to have taken effect as displayed by the **rmon show** command. However, changes to the RMON configuration take effect only after you reboot the management module.

5.16.3 Options

- if** Specifies the interface for which RMON table generation is to be enabled, disabled, or configured.
- interval** Specifies the amount of time, in seconds, that statistics collection will occur. The valid range is 1 - 3600 seconds. The factory default value is 30 minutes (1800 seconds).
- buckets** Specifies the number of collection sets that make up the History table buffer. The buffer is circular and when all buckets have been filled, the buffer wraps back to the first bucket and begins overwriting it with new data. The factory default number of buckets is 50. At 30 minutes per bucket, this allows for the collection of 25 hours of continuous statistics. The valid range for this parameter is 1-65535.

5.16.4 Example

```
>> rmon long 1 buckets 30
>> rmon show
```

Category	Limit	Used	Category	Limit	Used
rmon	507k	212k	macip	none	5k
stats	none	1k	filter	none	8k
history	none	36k	capture	none	124k
hosts	none	10k	alarm	none	0k
htopn	none	0k	event	none	0k
matrix	none	12k	ring	none	18k

Stats		-Short History--			--Long History--			Host	Mat	Mac/	RS	SR
I/F	Sts	Int.	Buckets	Sts	Int.	Buckets	Sts	Sts	Sts	IP	Sts	Sts
1	En	30	50	En	1800	50	En	En	En	En	En	En
2	Dis	30	50	Dis	1800	30	Dis	En	Dis	Dis	Dis	Dis
3	Dis	30	50	Dis	1800	50	Dis	Dis	Dis	Dis	Dis	Dis
4	Dis	30	50	Dis	1800	50	Dis	Dis	Dis	Dis	Dis	Dis

```
>>
```

Figure 5.14 - Example of rmon long Command

5.16.5 See Also

`rmon`, `rmon host`, `rmon matrix`, `rmon stats`, `rmon short`, `if`, `if mode`, `if cafmode`, `rmon limit`

5.17 rmon macip

Enables or disables the creation of the RMON MAC-IP table at management module initialization.

5.17.1 Format

```
rm[on] macip if en[able]
```

```
rm[on] matrix if dis[able]
```

5.17.2 Description

The `rmon macip` command can be used to disable or enable the creation of the RMON MAC to IP address mapping table for a specified interface when the management module is rebooted. By default from the factory, the creation of the MAC-IP table is disabled.

When enabled, one Matrix table is created for each interface specified by the `if` parameter. The MAC-IP table stores the MAC addresses mapped to the IP addresses of each device on the ring. This table automatically provides this mapping to provide a more familiar identifier for the user.

After enabling or disabling the table creation, the change will appear to have taken effect as displayed by the `rmon show` command. However, changes to the RMON configuration take effect only after you reboot the management module.

In general, the RMON functions provided by the agent place an additional burden on the agent (in terms of memory usage and CPU cycles). To maximize the utilization of the memory and CPU, RMON tables should be created sparingly (i.e., only created when needed).

5.17.3 Example

```

>> rmon show
Category      Limit      Used      Category      Limit      Used
-----      -
rmon          507k      212k      macip          none        5k
stats         none       1k        filter         none        8k
history       none       36k      capture        none       124k
hosts         none       10k      alarm          none        0k
htopn         none       0k        event          none        0k
matrix        none       12k      ring           none       18k

      Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1 En 30 50 En 1800 50 En En En En En En
2 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
3 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
4 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis

>> rmon macip 2 enable
Reboot the agent for changes to take affect.
(after the reboot...)
>> rmon show
Category      Limit      Used      Category      Limit      Used
-----      -
rmon          507k      212k      macip          none        5k
stats         none       1k        filter         none        8k
history       none       36k      capture        none       124k
hosts         none       10k      alarm          none        0k
htopn         none       0k        event          none        0k
matrix        none       12k      ring           none       18k

      Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1 En 30 50 En 1800 50 En En En En En En
2 Dis 30 50 Dis 1800 50 Dis Dis Dis En Dis Dis
3 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
4 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis

>>

```

Figure 5.15 - Example of rmon macip Command

5.17.4 See Also

`rmon`, `rmon short`, `rmon long`, `rmon long`, `rmon stats`, `rmon host`, `rmon limit`

5.18 rmon matrix

Enables or disables the creation of the RMON Matrix table at management module initialization.

5.18.1 Format

```
rm[on] matrix if en[able]
```

```
rm[on] matrix if dis[able]
```

5.18.2 Description

The `rmon matrix` command can be used to disable or enable the creation of the RMON Matrix table for a specified interface when the management module is rebooted. By default from the factory, the creation of the Matrix table is disabled.

When enabled, one Matrix table is created for each interface specified by the `if` parameter. The Matrix table contains statistics about the source and destination of packets and octets seen on the network. Refer to RFC 1271 for a more detailed description of the RMON Hosts table.

Matrix table statistics will only be collected for the interfaces that are configured in “promiscuous” frame capture mode (see `if mode promiscuous`). You will also need a remote RMON management application to view the statistics collected in the Matrix table.

After enabling or disabling Matrix table creation, the change will appear to have taken effect as displayed by the `rmon show` command. However, changes to the RMON configuration take effect only after you reboot the management module.

In general, the RMON functions provided by the agent place an additional burden on the agent (in terms of memory usage and CPU cycles). To maximize the utilization of the memory and CPU, RMON tables should be created sparingly (i.e., only created when needed).

5.18.3 Example

```
>> rmon show
Category    Limit    Used    Category    Limit    Used
-----    -
rmon        507k    212k    macip       none     5k
stats       none     1k      filter      none     8k
history     none     36k    capture     none     124k
hosts       none     10k    alarm       none     0k
htopn       none     0k     event       none     0k
matrix      none     12k    ring        none     18k

    Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1  En   30   50  En  1800   50  En  En  En  En  En  En
2  Dis  30   50  Dis 1800   50  Dis Dis Dis Dis Dis Dis
3  Dis  30   50  Dis 1800   50  Dis Dis Dis Dis Dis Dis
4  Dis  30   50  Dis 1800   50  Dis Dis Dis Dis Dis Dis

>> rmon matrix 2 enable
Reboot the agent for changes to take affect.
(after the reboot...)
>> rmon show
Category    Limit    Used    Category    Limit    Used
-----    -
rmon        507k    212k    macip       none     5k
stats       none     1k      filter      none     8k
history     none     36k    capture     none     124k
hosts       none     10k    alarm       none     0k
htopn       none     0k     event       none     0k
matrix      none     12k    ring        none     18k

    Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1  En   30   50  En  1800   50  En  En  En  En  En  En
2  Dis  30   50  Dis 1800   50  Dis En  En  Dis Dis Dis
3  Dis  30   50  Dis 1800   50  Dis Dis Dis Dis Dis Dis
4  Dis  30   50  Dis 1800   50  Dis Dis Dis Dis Dis Dis

>>
```

Figure 5.16 - Example of rmon matrix Command

5.18.4 See Also

`rmon`, `rmon long`, `rmon host`, `rmon stats`, `rmon short`, `rmon limit`

5.19 rmon short

Enables or disables the creation of the short RMON History table at management module initialization.

5.19.1 Format

```
rm[on] short if en[able]
rm[on] short if dis[able]
rm[on] short if interval interval
rm[on] short if buckets buckets
```

5.19.2 Description

The `rmon short` command can be used to disable or enable the creation of a “short” version of the RMON History table for a specified interface when the management module is rebooted. By factory default, the creation of the History tables is enabled. The short History table differs from the “long” History table in that the default statistics collection interval for the short table is 30 seconds, whereas it is 30 minutes (1800 seconds) for the long table.

The History table contains detailed statistics about network traffic over designated time periods, especially concerning error conditions such as collisions, dropped packets, CRC alignment errors, burst errors, line errors, and many more. Refer to RFC 1271 and RFC 1513 for a more detailed description of the RMON History tables.

Using the `interval` and `buckets` forms of the command, you can configure the quantity of data that will be collected in the table before it is overwritten. The `interval` parameter specifies the collection time period and the `buckets` parameter specifies the number of collection sets in the History table buffer. The buffer is circular and when all buckets have been filled, the buffer wraps back to the first bucket and begins overwriting it with new data.

The interfaces must also be configured in “promiscuous” frame copy mode to collect data for the History tables. This can be configured using the `if mode promiscuous` command.

After enabling or disabling table creation, the change will appear to have taken effect as displayed by the `rmon show` command. However, changes to the RMON configuration take effect only after you reboot the management module.

5.19.3 Options

- if** Specifies the interface for which RMON table generation is to be enabled, disabled, or configured.
- interval** Specifies the amount of time, in seconds, that statistics collection will occur. The valid range is 1 - 3600 seconds. The factory default value is 30 seconds.
- buckets** Specifies the number of collection sets per *interval* that make up the History table buffer. The buffer is circular and when all buckets have been filled, the buffer wraps back to the first bucket and begins overwriting it with new data. The factory default number of buckets is 50. At 30 minutes per bucket, this allows for the collection of 25 hours of continuous statistics. The valid range for this parameter is 1-65535.

5.19.4 Example

```
>> rmon short 2 interval 10
>> rmon show
Category      Limit      Used      Category      Limit      Used
-----
rmon          507k      212k      macip          none        5k
stats         none       1k        filter         none        8k
history       none      36k      capture        none      124k
hosts         none      10k      alarm          none         0k
htopn         none       0k        event          none         0k
matrix        none      12k      ring           none        18k

      Stats -Short History-- --Long History--  Host Mat  Mac/ RS   SR
I/F  Sts  Int. Buckets Sts  Int. Buckets Sts  Sts  Sts  IP  Sts  Sts
----
  1  En   30    50  En  1800    50  En  En  En  En  En  En
  2  Dis  10    50  Dis 1800    50  Dis En  Dis Dis Dis Dis
  3  Dis  30    50  Dis 1800    50  Dis Dis Dis Dis Dis Dis
  4  Dis  30    50  Dis 1800    50  Dis Dis Dis Dis Dis Dis
>>
```

Figure 5.17 - Example of rmon short Command

5.19.5 See Also

`rmon, rmon host, rmon long, rmon matrix, rmon stats, rmon limit, if, if mode, if cafmode`

5.20 rmon stats

Enables or disables the creation of the RMON Statistics table at management module initialization.

5.20.1 Format

```
rm[on] stats if en[able]
```

```
rm[on] stats if dis[able]
```

5.20.2 Description

The **rmon stats** command can be used to disable or enable the creation of the RMON Statistics table for a specified interface when the management module is rebooted. By default from the factory, the creation of the Statistics table is enabled.

If enabled from the command line, one Statistics table is created for each interface specified by the *if* parameter upon reboot of the management module. The Statistics table contains cumulative statistics about network traffic, especially concerning error conditions such as collisions, dropped packets, CRC alignment errors, burst errors, line errors, and many more. Refer to RFC 1271 and RFC 1513 for a more detailed description of the RMON Statistics table.

The interfaces must be configured in “promiscuous” frame copy mode to collect data for the Statistics tables. This can be configured using the **if mode promiscuous** command.

You must have a remote RMON management application to view the statistics collected.

After enabling or disabling Statistics table creation, the change will appear to have taken effect as displayed by the **rmon show** command. However, changes to the RMON configuration take effect only after you reboot the management module.

5.20.3 Example

```
>> rmon show
```

Category	Limit	Used	Category	Limit	Used
rmon	507k	79k	macip	none	4k
stats	none	1k	filter	none	0k
history	none	37k	capture	none	0k
hosts	none	10k	alarm	none	0k
htopn	none	0k	event	none	0k
matrix	none	11k	ring	none	18k

```

      Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1 En 30 50 En 1800 50 En En En En En En
2 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
3 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
4 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
>> rmon stats 1 disable
Reboot the agent for changes to take affect.
>> rmon show
```

Category	Limit	Used	Category	Limit	Used
rmon	507k	79k	macip	none	4k
stats	none	1k	filter	none	0k
history	none	37k	capture	none	0k
hosts	none	10k	alarm	none	0k
htopn	none	0k	event	none	0k
matrix	none	11k	ring	none	18k

```

      Stats -Short History-- --Long History-- Host Mat Mac/ RS SR
I/F Sts Int. Buckets Sts Int. Buckets Sts Sts Sts IP Sts Sts
--- ---
1 Dis 30 50 En 1800 50 En En En En En En
2 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
3 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
4 Dis 30 50 Dis 1800 50 Dis Dis Dis Dis Dis Dis
>>
```

Figure 5.18 - Example of rmon stats Command

5.20.4 See Also

`rmon`, `rmon host`, `rmon long`, `rmon matrix`, `rmon short`, `if`, `if mode`, `if cafmode`, `rmon limit`

5.21 route

Displays and alters the addresses used to route packets to other networks.

5.21.1 Format

```

    ro[ute] h[elp]

    ro[ute] s[how] [nvram]

    ro[ute] add [net | host] destination_ip gateway_ip

    ro[ute] add [net | host] d[efault] gateway_ip

    ro[ute] del[ete] destination_ip gateway_ip

    ro[ute] del[ete] d[efault] [gateway_ip]
  
```

5.21.2 Description

The ES-4810 management modules provide the capability to route packets between the packet buses and other networks. The agent software monitors the destination IP address of packets addressed to its interfaces. If a packet is received with a destination address for a different network segment, the management module looks in its routing table for a gateway that can reroute the packet to the proper network. The routing table can be viewed and modified using the **route** command.

The management module automatically defines a number of “static” routes in the table that cannot be modified or deleted. One of these is a loopback interface that the management module uses for internal routing. There is also a route defined for each of the interfaces on the management module that has an IP address assigned. The **route show** command will also display the user-defined routes, if there are any defined. You can view just the user-defined routes stored in NVRAM by issuing the **route show nvram** command.

Using the **route add** command, up to eight more routes can be added to the table for other network segments that will be attached to the hub. These “user-defined” routes are stored in NVRAM and are maintained through resets and outages. Furthermore, an additional route can be specified as the “default” route using the **route add default** command. Any packets that cannot be routed to any of the gateways in the table will be sent to the default gateway. You cannot add more than eight routes, not including the default route. Also, you cannot add a gateway that is not on one of the networks that are known to the management module (i.e., one of the interface networks). To delete any of the user-defined routes, use the **route**

delete command.

Routes to a particular host must be distinguished from those to a network. The optional keywords `net` and `host` force the destination to be interpreted as a network or a specific host, respectively. If neither is specified, the management module software tries to determine whether the destination is a host or network. Although `route add` commands can be processed without the optional `net` or `host` parameter, it should be included for best results.

5.21.3 Options

- | | |
|------------------------------|--|
| <i>destination_ip</i> | Specifies the destination network or host you want to add or delete in standard dot notation. If you specify the key word <code>default</code> or the destination address <code>0.0.0.0</code> on the <code>add</code> command, the <i>gateway_ip</i> specified is used as the default gateway. The key word <code>default</code> or the destination address “0.0.0.0” on the <code>delete</code> command deletes the default gateway. |
| <i>gateway_ip</i> | Specifies the IP address in standard dot notation of a device on a local network to be used for routing packets. When deleting the default gateway, this parameter is optional. |

5.21.4 Example

```
>> route add default 199.98.70.177
>> route add net 162.84.7.0 162.84.6.2
>> route show nvram
Destination      Gateway          Net/Host
default          199.98.70.177  Net
162.84.2.2       199.98.70.12   Host
162.84.7.0       162.84.6.2     Net
>> route show
Destination      Gateway          Net/Host
127.0.0.1        127.0.0.1
162.84.6.0        162.84.6.1
199.98.70.0       199.98.70.12
default          199.98.70.177
162.84.2.2       199.98.70.12
162.84.7.0       162.84.6.2
>> route del 162.84.6.0 162.84.6.1
Deleting this route is not allowed
>> route del 162.84.7.0 162.84.6.2
>> route show nvram
Destination      Gateway          Net/Host
default          199.98.70.177  Net
162.84.2.2       199.98.70.12   Host
>> route del 0.0.0.0
>> route show nvram
Destination      Gateway          Net/Host
162.84.2.2       199.98.70.12   Host
>>
```

Figure 5.19 - Example of route add Command

5.22 setup

Presents a brief description of the required setup procedure for the chassis.

5.22.1 Format

`se[tup]`

5.22.2 Example

```
>> setup
The MINIMUM setup required for this chassis entails -
  Initializing all groups on all User port cards
  (ie. group commands),
  Initializing all interfaces on the Network Management card
  (ie. if commands),
  Initializing the community names used to communicate
    via SNMP with the agent resident on the Network
    Management card
  (ie. community commands)

For each group on each User card, you need to initialize each
  port to either enabled or disabled, AND
  You need to assign each group to a backplane segment
For each interface on the Network Management card, you must
  enable it, and assign an IP address, a net mask,
  and a backplane segment
Common to all interfaces on the Network Management card,
  you need to assign a community name for read-only
  access and, if you wish to change configuration of the
  chassis via SNMP, you need to assign a community name
  for read-write access
>>
```

Figure 5.20 - Example of setup Command

5.22.3 See Also

`check`

5.23 snmp

Allows configuration of SNMPv2 Security parameters for each interface.

5.23.1 Format

```
sn[mp] h[elp]

sn[mp] s[how]

sn[mp] default

sn[mp] au[th-key] hex party-index hex-key

sn[mp] au[th-key] ascii party-index ascii-key

sn[mp] par[ty-init] if

sn[mp] re-[sync] party-index value

sn[mp] v1 en[able] | dis[able]
```

5.23.2 Description

The **snmp** commands allow you to configure the SNMPv2 security features on the management module agents. SNMPv2 provides more secure manager/agent communications by providing a way to “tag” the messages being sent between them. The messages being sent from a manager to an agent can be given authorization keys that specify who sent the message. Also, the messages sent back to the managers from the agent can have authorization keys specifying which agent sent the message.

The **snmp help** command provides a list of the SNMP commands and a synopsis of the configuration procedure. The **snmp show** command lists the current configuration of the SNMP security parameters for each interface. The entities that are communicating (i.e., managers and agents) are referred to as “parties.” The **snmp v1** command can be used to enable or disable SNMP version 1.

The **show** command lists the initial parties created automatically for each interface that is assigned an IP address. The entries in the column labeled `Party Index` without a “*” character denote parties that were installed into the agent’s MIB during the preceding reboot. Parties listed with the “*” character denote parties that have not been installed into the agent’s MIB which can occur for one of two reasons:

- The parties have just been created using the `snmp default` or `snmp party-init` commands, or
- The party is an `md5Auth` party and has not had an authentication key assigned to it yet. An `md5Auth` party will not be installed into the agent's database until an authentication key has been assigned.

The `snmp default` command can also be used to delete any entries already in the database and re-initialize the interfaces. From the operator console, this is the only way to delete existing entries in the tables. This command creates a set of initial parties, contexts, ACLs, and views for each agent interface with a non-zero IP address.

The SNMPv2 security feature also allows PDUs to be given a time stamp to prevent authorized PDUs from being captured and replayed at a later time. This is referred to as “replay protection”. In order for the time stamp to be validated by the receiving party, the sending and receiving parties must have synchronized clocks. These clocks are loosely synchronized and there is a window of time in which the time stamp must fall into. This window of time accounts for delays that may occur in transmission of the PDU and for drift that may occur in the clocks in the sending and receiving devices. The `snmp re-sync` command is provided to allow the re-synchronization of the party clocks. The `snmp re-sync` command allows you to update the clock value for any `md5Auth` party.

5.23.3 Options

<i>if</i>	is the interface (1-5) that you want to initialize parties on.
<i>party-index</i>	is the value displayed under the Party Index column of the <code>snmp show</code> command.
<i>hex-key</i>	is 16 bytes of data, entered in hexadecimal, with each byte separated by a space.
<i>ascii-key</i>	is an ASCII string of printable characters. Strings containing spaces must be entered using double quotes. Strings with more than 16 characters are truncated. Strings with less than 16 characters are padded with nulls (0x00).
<i>value</i>	is a counter value from 0-32767 and is usually set to zero to reset the clock.

5.23.4 Example

```
>> snmp show
Party
Index Auth Priv L/R Party OID
-----
  1 none none Loc 1.3.6.1.6.3.3.1.3.160.86.5.35.1
  2 none none Rem 1.3.6.1.6.3.3.1.3.160.86.5.35.2
 *1 md5 none Loc 1.3.6.1.6.3.3.1.3.160.86.5.35.3
 *2 md5 none Rem 1.3.6.1.6.3.3.1.3.160.86.5.35.4
SNMPv1 is enabled. Agent authentication clock = 7239009
>>
```

Figure 5.21 - Example of snmp Command

5.24 stbridge

Allows configuration and control of the Spanning Tree Protocol (STP) on NMM-SEG-1 management modules.

5.24.1 Format

```
stb[ridge] h[elp]

stb[ridge] s[how]

stb[ridge] re[set]

stb[ridge] de[fault]

stb[ridge] en[able]

stb[ridge] di[sable]

stb[ridge] pr[iority] priority

stb[ridge] age age_time

stb[ridge] hello hello_time

stb[ridge] f[orwarddelay] delay_time

stb[ridge] n[ewroot] {enable|disable}

stb[ridge] t[opchange] {enable|disable}
```

5.24.2 Description

The **stbridge** command is used to enable, disable, and configure the Spanning Tree Protocol on NMM-SEG-1 network management modules. To see the current spanning tree configuration, use the command **stbridge show**.

Enabling and disabling of spanning tree is accomplished with the **stbridge enable** and **stbridge disable** commands. By default, spanning tree will be disabled and will need to be enabled to start spanning tree operations. Likewise, spanning tree can only be enabled when the NMM-SEG-1 CAM is enabled.

To help ensure that the device becomes the root bridge, you can set a priority using the `stbridge priority` command. Using this command, you can specify a value that will be contained in the first two octets of the bridge ID.

During spanning tree calculation, the bridge transmits and receives “Hello Messages” that contain spanning tree information. The `stbridge hello` command is used to set the duration of time between “Hello Message” transmissions (*hello_time*). This applies to transmissions by this bridge on any port when it is acting as the root bridge.

The age time and forward delay time is set using `stbridge forwarddelay` and `age` commands. When acting as the root bridge these parameters are propagated to other non-root bridges. The age parameter determines the length of time a bridge will wait without receiving a BPDU before re-configuring. The forward delay time controls the length of time a bridge waits while new topology information is being propagated throughout the network.

If you want a trap to be sent to a management station each time a port changes topology or a new root is established, enable this feature using the `stbridge newroot` and `topchange` commands. By default both are disabled.

5.24.3 Options

priority Specifies the 4-digit hexadecimal priority of the bridge and becomes the first and second octets of the bridge identifier. Range 0-65535. The default is 8000 hexadecimal. The value can be entered in decimal, octal, or hexadecimal as described in “Setting Spanning Tree Parameters” on page 2-23.

hello_time Specifies the number of seconds the bridge will pause between the transmission of Hello Messages (BPDU's). The default time is 2 seconds. The valid range is 1-10 seconds.

age_time - the *age_time* parameter takes effect when a bridge is the root bridge. Any bridge that is not the root bridge uses the root bridge's *age_time*. The value specified (in seconds) determines how long a bridge waits without receiving a BPDU before attempting a re-configuration. Under normal conditions, the bridge ports should receive BPDUs at regular intervals.

Set the value of *age_time* to the following:

$$age_time \geq 2 \times (hello_time + 1)$$

The default value is 20 seconds and the valid range is 6-40 seconds

delay_time - specifies the number of seconds that ports will wait before it changes from listening to learning state and from learning to forwarding state. This delay is needed so every bridge on the network can receive information about the topology change before the port starts to forward packets. The default time is 15 seconds and the range of valid inputs is 4-30 seconds.

Set the value of *delay_time* to the following condition:

$$delay_time \geq (age_time/2) + 1$$

5.24.4 Example

```
>> stbridge show
Spanning Tree Protocol is enabled and has 148 ports on packetbus 1
bridge ID      : 80:00:00:00:52:32:90:00
designated root: 80:00:00:00:52:32:90:00
root path cost : 0
root port      : root bridge
bridge         : max age 20 secs, hello time 2 secs, forward delay 15 secs
root bridge    : max age 20 secs, hello time 2 secs, forward delay 15 secs
new root traps are disabled
topology change traps are disabled
>> stbridge newroot enable
>> stbridge topchange enable
>> stbridge show
Spanning Tree Protocol is enabled and has 148 ports on packetbus 1
bridge ID      : 80:00:00:00:52:32:90:00
designated root: 80:00:00:00:52:32:90:00
root path cost : 0
root port      : root bridge
bridge         : max age 20 secs, hello time 2 secs, forward delay 15 secs
root bridge    : max age 20 secs, hello time 2 secs, forward delay 15 secs
new root traps are enabled
topology change traps are enabled
>> stbridge age 25
>> stbridge forwarddelay 20
>> stbridge show
Spanning Tree Protocol is enabled and has 148 ports on packetbus 1
bridge ID      : 80:00:00:00:52:32:90:00
designated root: 80:00:00:00:52:32:90:00
root path cost : 0
root port      : root bridge
bridge         : max age 25 secs, hello time 2 secs, forward delay 20 secs
root bridge    : max age 25 secs, hello time 2 secs, forward delay 20 secs
new root traps are enabled
topology change traps are enabled
```

Figure 5.22 - Example of stbridge Command (one of two)

```
>> stbridge default
>> stbridge show
Spanning Tree Protocol is disabled and has 148 ports on packetbus 1
bridge ID      : 80:00:00:00:52:32:90:00
designated root: 80:00:00:00:52:32:90:00
root path cost : 0
root port      : root bridge
bridge         : max age 20 secs, hello time 2 secs, forward delay 15 secs
root bridge    : max age 20 secs, hello time 2 secs, forward delay 15 secs
new root traps are disabled
topology change traps are disabled
>>
```

Figure 5.23 - Example of stbridge Command (two of two)

5.25 stport

Displays status and provides control of the Spanning Tree Protocol (STP) on individual ports. This command applies only to management modules that support STP.

5.25.1 Format

```

    stp[ort] h[elp]

    stp[ort] sh[ow]

    stp[ort] st[atus] card-port

    stp[ort] re[set] card-port

    stp[ort] de[fault] card-port

    stp[ort] en[able] card-port

    stp[ort] di[sable] card-port

    stp[ort] pa[thcost] card-port pathcost

    stp[ort] pri[ority] card-port priority
  
```

5.25.2 Description

The **stport** command provides the ability to enable and disable spanning tree on individual ports. By default all ports will have spanning tree enabled.

NOTE: By default, spanning tree is disabled for the bridge but enabled for individual ports. By enabling spanning tree on the bridge, the current state specified for each port attached to the packet bus will become active.

To enable and disable individual ports use the command **stport enable** and **disable**.

If you've made changes to the configuration of a port and want to set it back to the default settings (including path cost, priority, and spanning tree enabled), enter the **stport default** command.

The path cost parameter provides a mechanism by which route decisions can be made. The path cost associated with each type of port is determined by the type of port and the speed of the network interface. The associated cost is inversely proportional to the speed. By default 10

Mbps ports have a cost of 100 and 100 Mbps ports have a cost of 10. The path cost for ports can be changed to encourage or discourage the routing of network traffic across specific connections.

5.25.3 Options

card	Specifies the number of the slot in which the card is installed.
port	Specifies a valid port number for the card type or asterisk (*) for all.
pathcost	Pathcost sets the path cost of each port. The range of valid inputs is 1-65535 when entered in decimal form.
priority	To help ensure that a port becomes the primary port for the bridge, you can set a priority using the stport priority command. Using this command, you can specify a value that will be contained in the first two octets of the port's ID. The default is 128. This number will be combined with a port identifier which is used by spanning tree to uniquely identify each port. The range is 0-255 and can be entered in decimal, octal, or hexadecimal form as described in "Enabling Spanning Tree" on page 2-23.

5.25.4 Example

```
>> stport show 4-*
port 4-1: port ID 0x8022, port enabled, stp enabled, forwarding, cost 100
port 4-2: port ID 0x8023, port enabled, stp enabled, forwarding, cost 100
port 4-3: port ID 0x8024, port enabled, stp enabled, forwarding, cost 100
port 4-4: port ID 0x8025, port enabled, stp enabled, forwarding, cost 100
port 4-5: port ID 0x8026, port enabled, stp enabled, forwarding, cost 100
port 4-6: port ID 0x8027, port enabled, stp enabled, forwarding, cost 100
port 4-7: port ID 0x8028, port enabled, stp enabled, forwarding, cost 100
port 4-8: port ID 0x8029, port enabled, stp enabled, forwarding, cost 100
port 4-9: port ID 0x802a, port enabled, stp enabled, forwarding, cost 100
port 4-10: port ID 0x802b, port enabled, stp enabled, forwarding, cost 100
port 4-11: port ID 0x802c, port enabled, stp enabled, forwarding, cost 100
port 4-12: port ID 0x802d, port enabled, stp enabled, forwarding, cost 100
port 4-13: port ID 0x802e, port enabled, stp enabled, forwarding, cost 100
port 4-14: port ID 0x802f, port enabled, stp enabled, forwarding, cost 100
port 4-15: port ID 0x8030, port enabled, stp enabled, forwarding, cost 100
port 4-16: port ID 0x8031, port enabled, stp enabled, forwarding, cost 100
port 4-17: port ID 0x8032, port enabled, stp enabled, forwarding, cost 100
port 4-18: port ID 0x8033, port enabled, stp enabled, forwarding, cost 100
port 4-19: port ID 0x8034, port enabled, stp enabled, forwarding, cost 100
port 4-20: port ID 0x8035, port enabled, stp enabled, forwarding, cost 100
port 4-21: port ID 0x8036, port enabled, stp enabled, forwarding, cost 100
port 4-22: port ID 0x8037, port enabled, stp enabled, forwarding, cost 100
port 4-23: port ID 0x8038, port enabled, stp enabled, forwarding, cost 100
port 4-24: port ID 0x8039, port enabled, stp enabled, forwarding, cost 100
port 4-25: port ID 0x803a, port enabled, stp enabled, forwarding, cost 100
port 4-26: port ID 0x803b, port enabled, stp enabled, forwarding, cost 100
port 4-27: port ID 0x803c, port enabled, stp enabled, forwarding, cost 100
port 4-28: port ID 0x803d, port enabled, stp enabled, forwarding, cost 100
port 4-29: port ID 0x803e, port enabled, stp enabled, forwarding, cost 100
port 4-30: port ID 0x803f, port enabled, stp enabled, forwarding, cost 100
port 4-31: port ID 0x8040, port enabled, stp enabled, forwarding, cost 100
port 4-32: port ID 0x8041, port enabled, stp enabled, forwarding, cost 100
```

Figure 5.24 - Example of the stport Command (one of two)

```
>> stport pathcost 4-4 200
>> stport show 4-4
port 4-4: port ID 0x8025, port enabled, stp enabled, forwarding, cost 200
>> stport stat 4-4
port 4-4: desg root 80:00:00:00:52:32:90:00, desg port 0x8025
           desg bridge 80:00:00:00:52:32:90:00, desg cost 0
>> stport disable 4-11
>> stport show 4-11
port 4-11: port ID 0x802c, port enabled, stp disabled, ignore, cost 100
>>
```

Figure 5.25 - Example of Command (two of two)

5.26 system

Presents the general management module and chassis information.

5.26.1 Format

```
s[ystem] h[elp]
```

```
s[ystem] s[how]
```

```
s[ystem] n[ame] name
```

```
s[ystem] c[ontact] contact
```

```
s[ystem] l[ocation] location
```

5.26.2 Description

The **system** command provides management module identification and information. The fields in the **system show** output message can be updated using the *name*, *contact*, and *location* forms of the command. Note that input strings containing blanks must be enclosed in double quotes.

5.26.3 Options

- | | |
|------------------------|--|
| <i>name</i> | Specifies a name for the chassis or module and is a character string up to 255 characters long. |
| <i>contact</i> | Specifies a person to contact in the event of problems or changes and is a character string up to 255 characters long. |
| <i>location</i> | Specifies the location of the chassis or module and is a character string up to 255 characters long. |

5.26.4 Example

```
>> system show  
Name.                cnt12.fore.com  
Contact:             Bob Roberts x2240  
Location:            Network Closet #7  
Current sysUpTime: 7526933  
>>
```

Figure 5.26 - Example of the system Command

5.27 trap

Configures the sending of traps for switched Ethernet agents.

5.27.1 Format

```
trap h[elp]
trap s[how] [address]
trap a[dd] address [community]
trap del[ete] address
trap co[mmunity] address community
```

5.27.2 Description

The **trap** command configures the sending of SNMP traps from the agent on the ES-4810 management module to network managers on the net. The management modules support the following SNMP generic and FORE Systems' enterprise-specific traps:

- Cold start of the management module
- Authentication failure
- Power supply insertion, removal, or failure
- Card insertion or removal from the chassis

A trap table stored on the module is used to determine which IP addresses of the network managers that will receive the trap information. Each IP address listed in the trap table will receive each trap generated by the agent on the module. The trap table can hold a maximum of eight IP addresses.

Use the **trap show** command to display the current trap table configuration. You can add new manager addresses to the table using the **trap add** command. To delete an entry from the table, use the **trap delete** command.

The **trap community** command allows you to specify the community string to be passed to the manager with each trap that is sent. The community string should correspond to one set up on the remote manager to allow access to its agent.

5.27.3 Options

address	Specifies a valid IP address expressed in standard dot notation.
community	(Optional on the <code>trap add</code> command) Specifies a quoted string of up to 15 characters that are sent in the community field of the trap when a trap is sent. The content of the community string is defined by the receiving station.

5.27.4 Example

```
>> trap show
trap 1: address 192.94.73.212, community public
>>
```

Figure 5.27 - Example of trap Command

5.28 version

Displays the version of firmware that is running on the management module.

5.28.1 Format

`v[ersion]`

5.28.2 Example

```
>> version
FORE Systems ES-4810
Version 4.7.3 (Mar 9 1998)
>>
```

Figure 5.28 - Example of version Command

Operator Console Command Reference

CHAPTER 6

Upgrading the Management Module Firmware

When you receive your management module from FORE Systems, the latest version of firmware is installed on the module. However, at some point you might have to upgrade the firmware on the module.

**NOTE**

The firmware upgrade must be performed separately for each network management module (NMM-1, NMM-2, or NMM-SEG-1), each ATM uplink, and each ASX switch installed in the ES-4810. Also, the firmware procedures are different for the different types of modules:

- For an NMM-1, NMM-2, or NMM-SEG-1, use the procedure given below.
- For an ATM uplink, refer to the procedures in the ATM Uplink User's Manual (MANU0294).
- For an ASX-200BX switch, refer to the ASX-200BX documentation.

6.1 Determining the Current Firmware Version

To check what version of firmware is installed on the management module, use one of the following methods. If you have management modules that are not using the current version of firmware then follow the instructions in the following sections.

6.1.1 Using the Version Command

At the Operator console, issue the `version` command. A string is displayed, indicating the current software version on the management module:

```
FORE Systems ES-4810  
Version 4.7.3 (Mar 9 1998)  
>>
```

The version also appears above the login prompt when you first access the console via telnet or terminal.

6.1.2 Using SNMP

Retrieve the MIB-II `sysDescr.0` object.

If you have management modules that are not using the current version of firmware then follow the instructions provided below.

6.2 Requirements for the Upgrade Process

To do this, you need:

- Upgrade files from FORE Systems. To get the proper upgrade file, contact the FORE Systems Technical Assistance Center (TAC) as described in “Technical Support” on page ii of the Preface.
- A TFTP client on the same IP subnet as the management module that you are upgrading. The management module acts as a TFTP server. The TFTP client and the management module must be assigned to the same IP subnet, with the appropriate IP address and subnet mask.



The IP address and subnet of each management module is independent of the IP address and configuration of other management modules, ATM modules, or ASX switches installed in the ES-4810.

- A community string configured on the management module to allow the TFTP client access with `super` privileges. Refer to “community” on page 5-9 for information on defining community strings.

6.3 Upgrading the Firmware

To upgrade the firmware:

1. Make sure you have the correct file.
2. Make sure the transfer mode is set to `binary` or `octet`.
3. Use your TFTP client to upload the upgrade file to the management module. Use the superuser password on the management module as the destination filename.

For example, if the upgrade file is named `mmm_4810_473` and the superuser password is `super`, on a typical TFTP client you would issue a command similar to the following:

```
put mmm_4810_473 super
```

The upgrade file will cause the management module to reprogram the Flash PROMs. There will be a 2.5 minute delay before you can communicate with the management module.

CAUTION



Do **not** reset or remove power from a management module while the upgrade is in progress.



When a download is initiated, the agent deletes all RMON tables (including those created automatically by the agent) in order to free up memory. If the download is cancelled before it completes, the RMON tables are not restored. Reboot the agent to restore the automatic RMON tables.

Glossary

802.1d Spanning Tree Bridging - the IEEE standard for bridging; a MAC layer standard for transparently connecting two or more LANs (often called subnetworks) that are running the same protocols and cabling. This arrangement creates an extended network, in which any two workstations on the linked LANs can share data.

802.3 Ethernet - the IEEE standard for Ethernet; a physical-layer standard that uses the CSMA/CD access method on a bus-topology LAN.

802.5 Token Ring - the IEEE physical-layer standard that uses the token-passing access method on a ring-topology LAN.

AAL (ATM Adaptation Layer) - the AAL divides the user information into segments suitable for packaging into a series of ATM cells. There are several types of AALs in use. FORE Systems currently supports AAL 5 and AAL 3/4. AAL 3/4 supports connection-oriented VBR data transfer and connectionless VBR data transfer, respectively. AAL 5 is defined as Simple and Efficient Adaptation Layer (SEAL).

AAL Connection - an association established by the AAL between two or more next higher layer entities.

ABR (Available Bit Rate) - a type of traffic for which the ATM network attempts to meet that traffic's bandwidth requirements. It does not guarantee a specific amount of bandwidth and the end station must retransmit any information that did not reach the far end.

ACR (Allowable Cell Rate) - parameter defined by the ATM Forum for ATM traffic management. ACR varies between the MCR and the PCR, and is dynamically controlled using congestion control mechanisms.

Address Mask - a bit mask used to identify which bits in an address (usually an IP address) are network significant, subnet significant, and host significant portions of the complete address. This mask is also known as the subnet mask because the subnetwork portion of the address can be determined by comparing the binary version of the mask to an IP address in that subnet. The mask holds the same number of bits as the protocol address it references.

Agent (SNMP) - a component of network- and desktop-management software, such as SNMP, that gathers information from MIBs.

AIS (Alarm Indication Signal) - a line AIS is asserted when a 111 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line AIS is removed when any pattern other than 111 is detected in these bits for five consecutive frames.

alarm - an unsolicited message from a device, typically indicating a problem with the system that requires attention.

AMI (ATM Management Interface) - the user interface to FORE Systems' *ForeThought* switch control software (SCS). AMI lets users monitor and change various operating configurations of FORE Systems switches and network module hardware and software, IP connectivity, and SNMP network management.

ANSI (American National Standards Institute) - a private organization that coordinates the setting and approval of some U.S. standards. It also represents the United States to the International Standards Organization.

API (Application Program Interface) - a language format that defines how a program can be made to interact with another program, service, or other software; it allows users to develop custom interfaces with FORE products.

APP (application program) - a complete, self-contained program that performs a specific function directly for the user.

AppleTalk - a networking protocol developed by Apple Computer for communication between Apple's products and other computers. Independent of the network layer, AppleTalk runs on LocalTalk, EtherTalk and TokenTalk.

ARP (Address Resolution Protocol) - a method used to resolve higher level protocol addressing (such as IP) into the appropriate header data required for ATM; i.e., port, VPI, and VCI; also defines the AAL type to be used.

ASCII (American Standard Code for Information Interchange) - a standard character set that (typically) assigns a 7-bit sequence to each letter, number, and selected control characters.

Assigned Cell - a cell that provides a service to an upper layer entity or ATM Layer Management entity (ATMM-entity).

asxmon - a FORE program that repeatedly displays the state of the switch and of all its active ports.

ATDM (Asynchronous Time Division Multiplexing) - a method of sending information that resembles normal TDM, except that time slots are allocated as needed rather than preassigned to specific transmitters.

ATM (Asynchronous Transfer Mode) - a transfer mode in which the information is organized into cells. It is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic.

ATM Forum - an international non-profit organization formed with the objective of accelerating the use of ATM products and services through a rapid convergence of interoperability specifications. In addition, the Forum promotes industry cooperation and awareness.

ATM Layer link - a section of an ATM Layer connection between two adjacent active ATM Layer entities (ATM-entities).

ATM Link - a virtual path link (VPL) or a virtual channel link (VCL).

ATM Peer-to-Peer Connection - a virtual channel connection (VCC) or a virtual path connection (VPC) directly established, such as workstation-to-workstation. This setup is not commonly used in networks.

ATM Traffic Descriptor - a generic list of parameters that can be used to capture the intrinsic traffic characteristics of a requested ATM connection.

ATM User-to-User Connection - an association established by the ATM Layer to support communication between two or more ATM service users (i.e., between two or more next higher layer entities or between two or more ATM entities). The communication over an ATM Layer connection may be either bidirectional or unidirectional. The same Virtual Channel Identifier (VCI) is used for both directions of a connection at an interface.

atmarp - a FORE program that shows and manipulates ATM ARP entries maintained by the given device driver. This is also used to establish PVC connections.

atmconfig - a FORE program used to enable or disable SPANS signalling.

atmstat - a FORE program that shows statistics gathered about a given adapter card by the device driver. These statistics include ATM layer and ATM adaptation layer cell and error counts. This can also be used to query other hosts via SNMP.

AUI (Attachment User Interface) - IEEE 802.3 interface between a media attachment unit (MAU) and a network interface card (NIC). The term AUI can also refer to the rear panel port to which an AUI cable might attach.

Auto-logout - a feature that automatically logs out a user if there has been no user interface activity for a specified length of time.

B8ZS (Bipolar 8 Zero Substitution) - a line coding technique used to accommodate the ones density requirements of T1 facilities.

Backbone - the main connectivity device of a distributed system. All systems that have connectivity to the backbone connect to each other. This does not stop systems from setting up private arrangements with each other to bypass the backbone for cost, performance, or security.

Bandwidth - usually identifies the capacity or amount of data that can be sent through a given circuit; may be user-specified in a PVC.

baud - unit of signalling speed. The speed in baud is the number of discrete conditions or signal events per second. If each signal event represents only one bit, the baud rate is the same as bps; if each signal event represents more than one bit (such as a dibit), the baud rate is smaller than bps.

BECN (Backward Explicit Congestion Notification) - bit set by a Frame Relay network in frames traveling in the opposite direction of frames encountering a congested path. Data terminal equipment (DTE) receiving frames with the BECN bit set can request that higher-level protocols take flow control action as appropriate. Compare with *FECN*.

BES (Bursty Errored Seconds) - a BES contains more than 1 and fewer than 320 path coding violation error events, and no severely errored frame or AIS defects. Controlled slips are not included in determining BESs.

BGP (Border Gateway Protocol) - used by gateways in an internet connecting autonomous networks. It is derived from experiences learned using the EGP.

BIP (Bit Interleaved Parity) - an error-detection technique in which character bit patterns are forced into parity, so that the total number of one bits is always odd or always even. This is accomplished by the addition of a one or zero bit to each byte, as the byte is transmitted; at the other end of the transmission, the receiving device verifies the parity (odd or even) and the accuracy of the transmission.

B-ISDN (Broadband Integrated Services Digital Network) - a common digital network suitable for voice, video, and high-speed data services running at rates beginning at 155 Mbps.

BNC (Bayonet-Neill-Concelman) - a bayonet-locking connector for miniature coax.

BPDU (Bridged Protocol Data Unit) - Spanning-tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network.

bps (bits per second) - a measure of speed or data rate. Often combined with metric prefixes in kbps for thousands of bits per second (k for kilo-) and in Mbps for millions of bits per second (M for mega-).

BPV (Bipolar Violation) - an error event on a line in which the normal pattern of alternating high (one) and low (zero) signals is disrupted. A bipolar violation is noted when two high signals occur without an intervening low signal, or vice versa.

Bridge - a device that expands a Local Area Network by forwarding frames between data link layers associated with two separate cables, usually carrying a common protocol. Bridges can usually be made to filter certain packets (to forward only certain traffic).

Broadband - a service or system requiring transmission channels capable of supporting rates greater than the Integrated Services Digital Network (ISDN) primary rate.

Broadband Access - an ISDN access capable of supporting one or more broadband services.

Router (bridging/router) - a device that routes some protocols and bridges others based on configuration information.

Bursty Second - a second during which there were at least the set number of BES threshold event errors but fewer than the set number of SES threshold event errors.

BUS (Broadcast and Unknown Server) - in an emulated LAN, the BUS is responsible for accepting broadcast, multicast, and unknown unicast packets from the LECs to the broadcast MAC address (FFFFFFFFFFFF) via dedicated point-to-point connections, and forwarding the packets to all of the members of the ELAN using a single point-to-multipoint connection.

CAC (Connection Admission Control) - the procedure used to decide if a request for an ATM connection can be accepted based on the attributes of both the requested connection and the existing connections.

Call - an association between two or more users or between a user and a network entity that is established by the use of network capabilities. This association may have zero or more connections.

Carrier - a company, such as any of the “baby Bell” companies, that provide network communications services, either within a local area or between local areas.

CBR (Constant Bit Rate) - a type of traffic that requires a continuous, specific amount of bandwidth over the ATM network (e.g., digital information such as video and digitized voice).

CBR port - a port on the *CellPath 300* for transmitting and receiving CBR traffic.

cchan - a FORE program used to manage virtual channels on a FORE Systems ATM switch running asxd.

CCITT (Consultative Committee for International Telephone and Telegraph) - an international consultative committee that sets international communications recommendations, which are frequently adopted as standards; develops interface, modem, and data network recommendations. Membership includes PTTs, scientific and trade associations, and private companies. CCITT is part of the International Communications Union (a United Nations treaty organization in Geneva).

CDV (Cell Delay Variation) - a quantification of cell clumping for a connection. The cell clumping CDV (y_k) is defined as the difference between a cell's expected reference arrival time (ck) and its actual arrival time (ak). The expected reference arrival time (ck) of cell k of a specific connection is $\max [c_{\{k-1\}} + T, a_k]$. T is the reciprocal of the negotiated peak cell rate.

CE (Connection Endpoint) - a terminator at one end of a layer connection within a SAP.

CEI (Connection Endpoint Identifier) - an identifier of a CE that can be used to identify the connection at a SAP.

Cell - an ATM Layer protocol data unit (PDU). The basic unit of information transported in ATM technology, each 53-byte cell contains a 5-byte header and a 48-byte payload.

Cell Delineation - the protocol for recognizing the beginning and end of ATM cells within the raw serial bit stream.

Cell Header - ATM Layer protocol control information.

Cell Port - a port on the *CellPath 300* that transmits and receives traffic in cell format.

Cell Rate Adaptation - a function performed by a protocol module in which empty cells (known as unassigned cells) are added to the output stream. This is because there always must be a fixed number of cells in the output direction; when there are not enough cells to transmit, unassigned cells are added to the output data stream.

Cell Transfer Delay - the transit delay of an ATM cell successfully passed between two designated boundaries.

CES (Circuit emulation Services) - The *CellPath 90* supports Circuit Emulation Services (CES) for applications requiring a fixed delay, lossless end-to-end connection through the network. In essence, CES provides a virtual private line service to the connecting application.

Channelization - capability of transmitting independent signals together over a cable while still maintaining their separate identity for later separation.

CLP (Cell Loss Priority) - the last bit of byte four in an ATM cell header; indicates the eligibility of the cell for discard by the network under congested conditions. If the bit is set to 1, the cell may be discarded by the network depending on traffic conditions.

Cold Start Trap - a *CellPath 300* SNMP trap which is sent when the unit has been power-cycled (see trap).

Comm Port - the front panel DCE port that allows access to the *CellPath 300* user interface via a connected terminal.

Community String - the password that allows an SNMP manager to access the agent information. Each request from a manager is accompanied by a community string.

Concentrator - a communications device that offers the ability to concentrate many lower-speed channels into and out of one or more high-speed channels.

Congestion Management - a *CellPath 300* feature that helps ensure reasonable service for VBR connections in an ATM network. For each connection, the *CellPath 300* maintains a priority, sustained cell rate (SCR), and peak cell rate (PCR). During times of congestion, the *CellPath 300* reduces the bandwidth to the SCR, based on the priority of the connection.

Connection - the concatenation of ATM Layer links in order to provide an end-to-end information transfer capability to access points.

Connectionless Service - a type of service in which no pre-determined path or link has been established for transfer of information, supported by AAL 4.

Connection-Oriented Service - a type of service in which information always traverses the same pre-established path or link between two points, supported by AAL 3.

Controlled Slip - a situation in which one frame's worth of data is either lost or replicated. A controlled slip typically occurs when the sending device and receiving device are not using the same clock.

Corresponding Entities - peer entities with a lower layer connection among them.

cpath - a FORE program used to manage virtual paths on a FORE Systems ATM switch running asxd.

CPE (Customer Premise Equipment) - equipment that is on the customer side of the point of demarcation, as opposed to equipment that is on a carrier side. *See also* point of demarcation.

cport - a FORE program used to monitor and change the state of ports on a FORE Systems ATM switch running asxd.

CRC (Cyclic Redundancy Check) - an error detection scheme in which a number is derived from the data that will be transmitted. By recalculating the CRC at the remote end and comparing it to the value originally transmitted, the receiving node can detect errors.

Cross Connection - a mapping between two channels or paths at a network device such as the *CellPath 300*.

CD (Controlled Slip) - a situation in which one frame's worth of data is either lost or replicated. A controlled slip typically occurs when the sending device and receiving device are not using the same clock.

CS (Convergence Sublayer) - a portion of the AAL. Data is passed first to the CS where it is divided into rational, fixed-length packets or PDUs (Protocol Data Units). For example, AAL 4 processes user data into blocks that are a maximum of 64 kbytes long.

CTS (Clear To Send) - and RS-232 modem interface control signal (sent from the modem to the DTE on pin 5) which indicates that the attached DTE may begin transmitting; issuance in response to the DTE's RTS.

D4 framing - See SF)

DARPA (Defense Advanced Research Projects Agency) - the US government agency that funded the ARPANET.

Datagram - a packet of information used in a connectionless network service that is routed to its destination using an address included in the datagram's header.

DCE (Data Communications Equipment) - a definition in the RS232C standard that describes the functions of the signals and the physical characteristics of an interface for a communication device such as a modem.

DCS (Digital Cross-connect System) - an electronic patch panel used to route digital signals in a central office.

Demultiplexing - a function performed by a layer entity that identifies and separates SDUs from a single connection to more than one connection (*see* multiplexing).

DFA (DXI Frame Address) - a connection identifier associated with ATM DXI packets that serves the same functions as, and translates directly to, the VPI/VCI on an ATM cell.

DIP Switch (Dual In-line Package) - a device that has two parallel rows of contacts that let the user switch electrical current through a pair of those contacts to on or off. They are used to reconfigure components and peripherals.

DLCI (Data Link Connection Identifier) - a connection identifier associated with frame relay packets that serves the same functions as, and translates directly to, the VPI/VCI on an ATM cell.

Domain Name Server - a computer that converts names to their corresponding Internet numbers. It allows users to telnet or FTP to the name instead of the number.

DNS (Domain Name System) - the distributed name and address mechanism used in the Internet.

DSn (Digital Standard n (0, 1, 1C, 2, and 3)) - a method defining the rate and format of digital hierarchy, with asynchronous data rates defined as follows:

DS0	64kbps	1 voice channel
DS1	1.544Mbps	24 DS0s
DS1C	3.152 Mbps	2 DS1s
DS2	6.312 Mbps	4 DS1s
DS3	44.736 Mbps	28 DS1s

Synchronous data rates (SONET) are defined as:

STS-1/OC-1	51.84 Mbps	28 DS1s or 1 DS3
STS-3/OC-3	155.52 Mbps	3 STS-1s byte interleaved
STS-3c/OC-3c	155.52 Mbps	Concatenated, indivisible payload
STS-12/OC-12	622.08 Mbps	12 STS-1s, 4 STS-3cs, or any mixture
STS-12c/OC-12c	622.08 Mbps	Concatenated, indivisible payload
STS-48/OC-48	2488.32 Mbps	48 STS-1s, 16 STS-3cs, or any mixture

DSR (Data Set Ready) - an RS-232 modem interface control signal (sent from the modem to the DTE on pin 6) which indicates that the modem is connected to the telephone circuit. Usually a prerequisite to the DTE issuing RTS.

DTE (Data Terminal Equipment) - generally user devices, such as terminals and computers, that connect to data circuit-terminating equipment. They either generate or receive the data carried by the network.

DTR (Data Terminal Ready) - an RS232 modem interface control signal (sent from the DTE to the modem on pin 20) which indicates that the DTE is ready for data transmission and which requests that the modem be connected to the telephone circuit.

DXI - a generic phrase used in the full names of several protocols, all commonly used to allow a pair of DCE and DTE devices to share the implementation of a particular WAN protocol. The protocols all define the packet formats used to transport data packets between DCE and DTE devices.

E1 - Wide-area digital transmission scheme used predominantly in Europe that carries data at a rate of 2.048 Mbps. E1 lines can be leased for private use from common carriers.

E3 - Wide-area digital transmission scheme used predominantly in Europe that carries data at a rate of 34.368 Mbps. E3 lines can be leased for private use from common carriers.

EEPROM (Electrically Erasable Programmable Read Only Memory) - an EPROM that can be cleared with electrical signals rather than the traditional ultraviolet light.

EFCI (Explicit Forward Congestion Indication) - the second bit of the payload type field in the header of an ATM cell, the EFCI bit indicates network congestion to receiving hosts. On a congested switch, the EFCI bit is set to "1" by the transmitting network module when a certain number of cells have accumulated in the network module's shared memory buffer. When a cell is received that has its EFCI bit set to "1," the receiving host notifies the sending host, which should then reduce its transmission rate.

EGP (Exterior Gateway) Protocol - used by gateways in an internet, connecting autonomous networks.

EIA (Electronics Industries Association) - a USA trade organization that issues its own standards and contributes to ANSI; developed RS-232. Membership includes USA manufacturers.

EISA (Extended Industry Standard Architecture) - a bus architecture for desktop computers that provides a 32-bit data passage while maintaining compatibility with the ISA or AT architecture.

elarp - a FORE program that shows and manipulates MAC and ATM address mappings for LAN Emulation Clients (LECs).

elconfig - a FORE program that shows and modifies LEC configuration. Allows the user to set the NSAP address of the LAN Emulation Configuration Server (LECS), display the list of Emulated LANs (ELANs) configured in the LECS for this host, display the list of ELANs locally configured along with the membership state of each, and locally administer ELAN membership.

EM - the *CellPath* 300 extension module; paired with the system controller and supporting an optional PCMCIA card.

Embedded SNMP Agent - an SNMP agent can come in two forms: embedded or proxy. An embedded SNMP agent is integrated into the physical hardware and software of the unit. The *CellPath* 300 has an internal, integrated SNMP agent.

EMI (Electromagnetic Interference) - signals generated and radiated by an electronic device that cause interference with radio communications, among other effects.

End-to-End Connection - when used in reference to an ATM network, a connection that travels through an ATM network, passing through various ATM devices and with endpoints at the termination of the ATM network.

EPROM - Erasable Programmable Read Only Memory (see PROM).

EQL (Equalization) - the process of compensating for line distortions.

ES (End System) - a system in which an ATM connection is terminated or initiated. An originating end system initiates the ATM connection, and a terminating end system terminates the ATM connection. OAM cells may be generated and received.

ES (Errored Seconds) - a second during which at least one code violation occurred.

ESF (Extended Superframe) - T1 framing standard that provides frame synchronization, cyclic redundancy, and data link bits.

Ethernet - a 10-Mbps, coaxial standard for LANs in which all nodes connect to the cable where they contend for access.

Fairness - as related to Generic Flow Control (GFC), fairness is defined as meeting all of the agreed quality of service (QoS) requirements by controlling the order of service for all active connections.

Far-End - in a relationship between two devices in a circuit, the far-end device is the one that is remote.

FCC - a board of commissioners appointed by the President under the Communications Act of 1934, with the authority to regulate all interstate telecommunications originating in the United States, including transmission over phone lines.

FDDI (Fiber Distributed Data Interface) - high-speed data network that uses fiber-optic as the physical medium. Operates in similar manner to Ethernet or Token Ring, only faster.

FDM (Frequency Division Multiplexing) - a method of dividing an available frequency range into parts with each having enough bandwidth to carry one channel.

FEBE (Far End Block Error) - an error detected by extracting the 4-bit FEBE field from the path status byte (G1). The legal range for the 4-bit field is between 0000 and 1000, representing zero to eight errors. Any other value is interpreted as zero errors.

FECN (Forward Explicit Congestion Notification) - bit set by a Frame Relay network to inform data terminal equipment (DTE) receiving the frame that congestion was experienced in the path from source to destination. DTE receiving frames with the FECN bit set can request that higher-level protocols take flow control action as appropriate. Compare with *BECN*.

FERF (Far End Receive Failure) - a line error asserted when a 110 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line FERF is removed when any pattern other than 110 is detected in these bits for five consecutive frames.

FIFO (First-In, First-Out) - a method of coordinating the sequential flow of data through a buffer.

Flag - a specific bit pattern used to identify the beginning or end of a frame.

Frame - a variable length group of data bits with a specific format containing flags at the beginning and end to provide demarcation.

Frame Relay - a fast packet switching protocol based on the LAPD protocol of ISDN that performs routing and transfer with less overhead processing than X.25.

Frame Synchronization Error - an error in which one or more time slot framing bits are in error.

Framing - a protocol that separates incoming bits into identifiable groups so that the receiving multiplexer recognizes the grouping.

FT-PNNI (ForeThought PNNI) - a FORE Systems routing and signalling protocol that uses private ATM (NSAP) addresses; a precursor to ATM Forum PNNI (*see* PNNI).

FTP (File Transfer Protocol) - a TCP/IP protocol that lets a user on one computer access, and transfer data to and from, another computer over a network. ftp is usually the name of the program the user invokes to accomplish this task.

GCRA (Generic Cell Rate Algorithm) - an algorithm which is employed in traffic policing and is part of the user/network service contract. The GCRA is a scheduling algorithm which ensures that cells are marked as *conforming* when they arrive when expected or later than expected and *non-conforming* when they arrive sooner than expected.

GFC (Generic Flow Control) - the first four bits of the first byte in an ATM cell header. Used to control the flow of traffic across the User-to-Network Interface (UNI), and thus into the network. Exact mechanisms for flow control are still under investigation and no explicit definition for this field exists at this time. (This field is used only at the UNI; for NNI-NNI use (between network nodes), these four bits provide additional network address capacity, and are appended to the VPI field.)

GIO - a proprietary bus architecture used in certain Silicon Graphics, Inc. workstations.

Header - protocol control information located at the beginning of a protocol data unit.

HDB3 (High Density Bipolar) - line-code type standard for T1 where each block of three zeros is replaced by 00V or B0V, where B represents an inserted pulse conforming to the AMI rule (ITU-T G.701, item 9004) and V represents an AMI violation (ITU-T G.701, item 9007). The choice of 00V or B0V is made so that the number of B pulses between consecutive V pulses is odd (successive V pulses are of alternate polarity so that no d.c. component is introduced). Compare with *AMI*.

HDLC (High-Level Data Link Control) - Bit-oriented synchronous data link layer protocol developed by the ISO. Derived from SDLC, HDLC specifies a data encapsulation method on synchronous serial links using frame characters and checksums. See also *SDLC*.

HEC (Header Error Control) - a CRC code located in the last byte of an ATM cell header that is used for checking cell integrity only.

HIPPI (High Performance Parallel Interface) - ANSI standard that extends the computer bus over fairly short distances at speeds of 800 and 1600 Mbps.

HPUX - the Hewlett-Packard version of UNIX.

HSSI (High-Speed Serial Interface) - a serial communications connection that operates at speeds of up to 1.544 Mbps.

Hub - a device that connects several other devices, usually in a star topology.

I/O Module - FORE's interface cards for the LAX-20 LAN Access Switch, designed to connect Ethernet, Token Ring, and FDDI LANs to *ForeRunner* ATM networks.

ICMP (Internet Control Message Protocol) - the protocol that handles errors and control messages at the IP layer. ICMP is actually a part of the IP protocol layer. It can generate error messages, test packets, and informational messages related to IP.

IEEE (Institute of Electrical and Electronics Engineers) - the world's largest technical professional society. Based in the U.S., the IEEE sponsors technical conferences, symposia & local meetings worldwide, publishes nearly 25% of the world's technical papers in electrical, electronics & computer engineering, provides educational programs for members, and promotes standardization.

IETF (Internet Engineering Task Force) - a large, open, international community of network designers, operators, vendors and researchers whose purpose is to coordinate the operation, management and evolution of the Internet to resolve short- and mid-range protocol and architectural issues.

ILMI (Interim Local Management Interface) - the standard that specifies the use of the Simple Network Management Protocol (SNMP) and an ATM management information base (MIB) to provide network status and configuration information.

Interface Data - the unit of information transferred to/from the upper layer in a single interaction across a SAP. Each Interface Data Unit (IDU) controls interface information and may also contain the whole or part of the SDU.

internet - while an internet is a network, the term "internet" is usually used to refer to a collection of networks interconnected with routers.

Internet - (note the capital "I") the largest internet in the world including large national backbone nets and many regional and local networks worldwide. The Internet uses the TCP/IP suite. Networks with only e-mail connectivity are not considered on the Internet.

Internet Addresses - the numbers used to identify hosts on an internet network. Internet host numbers are divided into two parts; the first is the network number and the second, or local, part is a host number on that particular network. There are also three classes of networks in the Internet, based on the number of hosts on a given network. Large networks are classified as Class A, having addresses in the range 1-126 and having a maximum of 16,387,064 hosts. Medium networks are classified as Class B, with addresses in the range 128-191 and with a maximum of 64,516 hosts. Small networks are classified as Class C, having addresses in the range 192-254 with a maximum of 254 hosts. Addresses are given as dotted decimal numbers in the following format:

nnn.nnn.nnn.nnn

In a Class A network, the first of the numbers is the network number, the last three numbers are the local host address.

In a Class B network, the first two numbers are the network, the last two are the local host address.

In a Class C network, the first three numbers are the network address, the last number is the local host address.

The following table summarizes the classes and sizes:

<u>Class</u>	<u>First #</u>	<u>Max# Hosts</u>
A	1-126	16,387,064
B	129-191	64,516
C	192-223	254

Network mask values are used to identify the network portion and the host portion of the address. For:

Class A - the default mask is 255.0.0.0

Class B - the default mask is 255.255.0.0

Class C - the default mask is 255.255.255.0

Subnet masking is used when a portion of the host ID is used to identify a subnetwork. For example, if a portion of a Class B network address is used for a subnetwork, the mask could be set as 255.255.255.0. This would allow the third byte to be used as a subnetwork address. All hosts on the network would still use the IP address to get on the Internet.

IP (Internet Protocol) - a connectionless, best-effort packet switching protocol that offers a common layer over dissimilar networks.

IP Address - a unique 32-bit integer used to identify a device in an IP network. You will most commonly see IP addresses written in “dot” notation; for instance, 192.228.32.14 (see IP net-mask).

IP Netmask - a pattern of 32 bits that is combined with an IP address to determine which bits of an IP address denote the network number and which denote the host number. Netmasks are useful for sub-dividing IP networks. IP netmasks are written in “dot” notation; for instance, 255.255.255.0 (see IP address).

IPX Protocol (Internetwork Packet Exchange) - a NetWare protocol similar to the Xerox Network Systems (XNS) protocol that provides datagram delivery of messages.

IS (Intermediate system) - a system that provides forwarding functions or relaying functions or both for a specific ATM connection. OAM cells may be generated and received.

ISA Bus - a bus standard developed by IBM for expansion cards in the first IBM PC. The original bus supported a data path only 8 bits wide. IBM subsequently developed a 16-bit version for its AT class computers. The 16-bit AT ISA bus supports both 8- and 16-bit cards. The 8-bit bus is commonly called the PC/XT bus, and the 16-bit bus is called the AT bus.

ISDN (Integrated Services Digital Network) - an emerging technology that is beginning to be offered by the telephone carriers of the world. ISDN combines voice and digital network services into a single medium or wire.

ISO (International Standards Organization) - a voluntary, non treaty organization founded in 1946 that is responsible for creating international standards in many areas, including computers and communications.

Isochronous - signals carrying embedded timing information or signals that are dependent on uniform timing; usually associated with voice and/or video transmission.

ITU (International Telecommunications Union) - the telecommunications agency of the United Nations, established to provide standardized communications procedures and practices, including frequency allocation and radio regulations, on a worldwide basis.

J2 - Wide-area digital transmission scheme used predominantly in Japan that carries data at a rate of 6.312 Mbps.

Jitter - analog communication line distortion caused by variations of a signal from its reference timing position.

Jumper - a patch cable or wire used to establish a circuit, often temporarily, for testing or diagnostics; also, the devices, shorting blocks, used to connect adjacent exposed pins on a printed circuit board that control the functionality of the card.

LAN (Local Area Network) - a data network intended to serve an area of only a few square kilometers or less. Because the network is known to cover only a small area, optimizations can be made in the network signal protocols that permit higher data rates.

lane - a program that provides control over the execution of the LAN Emulation Server (LES), Broadcast/Unknown Server (BUS), and LAN Emulation Configuration Server (LECS) on the local host.

LAN Access Concentrator - a LAN access device that allows a shared transmission medium to accommodate more data sources than there are channels currently available within the transmission medium.

LAPB (Link Access Procedure, Balanced) - Data link protocol in the X.25 protocol stack. LAPB is a bit-oriented protocol derived from HDLC. See also HDLC and X.25.

LAX-20 - a FORE Systems LAN Access Switch, designed to connect Ethernet, Token Ring, and FDDI LANs to *ForeRunner* ATM networks. The LAX-20 is a multiport, multiprotocol internet-working switch that combines the advantages of a high-performance LAN switch and a full-featured ATM interface capable of carrying LAN traffic.

Layer Entity - an active layer within an element.

Layer Function - a part of the activity of the layer entities.

Layer Service - a capability of a layer and the layers beneath it that is provided to the upper layer entities at the boundary between that layer and the next higher layer.

Layer User Data - the information transferred between corresponding entities on behalf of the upper layer or layer management entities for which they are providing services.

le - a FORE program that implements both the LAN Emulation Server (LES) and the Broadcast/Unknown Server (BUS).

LEC (LAN Emulation Client) - the component in an end system that performs data forwarding, address resolution, and other control functions when communicating with other components within an ELAN.

lecs - a FORE program that implements the assignment of individual LECs to different emulated LANs.

LECS (LAN Emulation Configuration Server) - the LECS is responsible for the initial configuration of LECs. It provides information about available ELANs that a LEC may join, together with the addresses of the LES and BUS associated with each ELAN.

leq - a FORE program that provides information about an ELAN. This information is obtained from the LES, and includes MAC addresses registered on the ELAN together with their corresponding ATM addresses.

LES (LAN Emulation Server) - the LES implements the control coordination function for an ELAN. The LES provides the service of registering and resolving MAC addresses to ATM addresses.

Link Down Trap - a *CellPath* 300 SNMP trap that signifies that the Ethernet interface has transitioned from a normal state to an error state, or has been disconnected.

Link Up Trap - a *CellPath* 300 SNMP trap that signifies that the Ethernet interface has transitioned from an error condition to a normal state.

LLC (Logical Link Control) - a protocol developed by the IEEE 802 committee for data-link-layer transmission control; the upper sublayer of the IEEE Layer 2 (OSI) protocol that complements the MAC protocol; IEEE standard 802.2; includes end-system addressing and error checking.

LOF (Loss Of Frame) - a type of transmission error that may occur in wide-area carrier lines.

Loopback - a troubleshooting technique that returns a transmitted signal to its source so that the signal can be analyzed for errors. Typically, a loopback is set at various points in a line until the section of the line that is causing the problem is discovered.

looptest - a program that tests the interface for basic cell reception and transmission functionality. It is usually used for diagnostic purposes to determine if an interface is functioning properly.

LOP (Loss Of Pointer) - a type of transmission error that may occur in wide-area carrier lines.

LOS (Loss Of Signal) - a type of transmission error that may occur in wide-area carrier lines.

MAC (Media Access Control) - a media-specific access control protocol within IEEE 802 specifications; currently includes variations for Token Ring, token bus, and CSMA/CD; the lower sublayer of the IEEE's link layer (OSI), which complements the Logical Link Control (LLC).

MAU (Media Attachment Unit) - device used in Ethernet and IEEE 802.3 networks that provides the interface between the AUI port of a station and the common medium of the Ethernet. The MAU, which can be built into a station or can be a separate device, performs physical layer functions including conversion of the digital data from the Ethernet interface, collision detection, and injection of bits onto the network.

Maximum Burst Tolerance - the largest burst of data that a network device is guaranteed to handle without discarding cells or packets. Bursts of data larger than the maximum burst size may be subject to discard.

MCR (Minimum Cell Rate) - parameter defined by the ATM Forum for ATM traffic management. MCR is defined only for ABR transmissions, and specifies the minimum value for the ACR.

Metasignalling - an ATM Layer Management (LM) process that manages different types of signalling and possibly semipermanent virtual channels (VCs), including the assignment, removal, and checking of VCs.

Metasignalling VCs - the standardized VCs that convey metasignalling information across a User-to-Network Interface (UNI).

MIB (Management Information Base) - the set of parameters that an SNMP management station can query or set in the SNMP agent of a networked device (e.g., router).

MIC (Media Interface Connector) - the optical fiber connector that joins the fiber to the FDDI controller.

MicroChannel - a proprietary 16- or 32-bit bus developed by IBM for its PS/2 computers' internal expansion cards; also offered by others.

MTU (Maximum Transmission Unit) - the largest unit of data that can be sent over a type of physical medium.

Multi-homed - a device that has both an ATM and another network connection, typically Ethernet.

Multiplexing - a function within a layer that interleaves the information from multiple connections into one connection (*see demultiplexing*).

Multipoint Access - user access in which more than one terminal equipment (TE) is supported by a single network termination.

Multipoint-to-Point Connection - a Point-to-Multipoint Connection may have zero bandwidth from the Root Node to the Leaf Nodes, and non-zero return bandwidth from the Leaf Nodes to the Root Node. Such a connection is also known as a Multipoint-to-Point Connection.

Multipoint-to-Multipoint Connection - a collection of associated ATM VC or VP links, and their associated endpoint nodes, with the following properties:

1. All N nodes in the connection, called Endpoints, serve as a Root Node in a Point-to-Multipoint connection to all of the (N-1) remaining endpoints.
2. Each of the endpoints can send information directly to any other endpoint, but the receiving endpoint cannot distinguish which of the endpoints is sending information without additional (e.g., higher layer) information.

Near-End - in a relationship between two devices in a circuit, the near-end device is the one that is local.

Network Module - ATM port interface cards which may be individually added or removed from any *ForeRunner* ATM switch to provide a diverse choice of connection alternatives. Each network module provides between one and six full-duplex ATM physical connections to the *ForeRunner* switch.

NMS (Network Management Station) - the system responsible for managing a network or a portion of a network. The NMS talks to network management agents, which reside in the managed nodes.

NNI (Network-to-Network Interface or Network Node Interface) - the interface between two public network pieces of equipment.

nonvolatile - a term used to describe a data storage device (memory) that retains its contents when power is lost.

NuBus - a high-speed bus used in the Macintosh family of computers, structured so that users can put a card into any slot on the board without creating conflict over the priority between those cards

OAM (Operation and Maintenance) Cell - a cell that contains ATM LM information. It does not form part of the upper layer information transfer.

octet - a grouping of 8 bits; similar, but not identical, to a byte.

OID (Object Identifier) - the address of a MIB variable.

OOF (Out-of-Frame) - a signal condition and alarm in which some or all framing bits are lost.

OpenView - Hewlett-Packard's network management software.

OSI (Open Systems Interconnection) - the 7-layer suite of protocols designed by ISO committees to be the international standard computer network architecture.

OSPF (Open Shortest Path First) Protocol - a routing algorithm for IP that incorporates least-cost, equal-cost, and load balancing.

Out-of-Band Management - refers to switch configuration via the serial port or over Ethernet, not ATM.

packet - a group of bits - including information bits and overhead bits - transmitted as a complete package on a network. Usually smaller than a transmission block.

Packet Port - a port on the *CellPath 300* that transmits and receives packet traffic.

Packet Switching - a communications paradigm in which packets (messages) are individually routed between hosts with no previously established communications path.

Payload Scrambling - a technique that eliminates certain bit patterns that may occur within an ATM cell payload that could be misinterpreted by certain sensitive transmission equipment as an alarm condition.

PBX (Private Branch Exchange) - a private phone system (switch) that connects to the public telephone network and offers in-house connectivity. To reach an outside line, the user must dial a digit like 8 or 9.

PCI (Peripheral Component Interconnect) - a local-bus standard created by Intel.

PCM (Pulse Code Modulation) - a modulation scheme that samples the information signals and transmits a series of coded pulses to represent the data.

PCR (Peak Cell Rate) - parameter defined by the ATM Forum for ATM traffic management. In CBR transmissions, PCR determines how often data samples are sent. In ABR transmissions, PCR determines the maximum value of the ACR.

PDN (Public Data Network) - a network designed primarily for data transmission and intended for sharing by many users from many organizations.

PDU (Protocol Data Unit) - a unit of data specified in a layer protocol and consisting of protocol control information and layer user data.

Peak Cell Rate - at the PHY Layer SAP of a point-to-point VCC, the Peak Cell Rate R_{pis} is the inverse of the minimum inter-arrival time T_0 of the request to send an ATM-SDU.

Peer Entities - entities within the same layer.

PHY (Physical Layer) - the actual cards, wires, and/or fiber-optic cabling used to connect computers, routers, and switches.

Physical Layer (PHY) Connection - an association established by the PHY between two or more ATM-entities. A PHY connection consists of the concatenation of PHY links in order to provide an end-to-end transfer capability to PHY SAPs.

PLCP (Physical Layer Convergence Protocol) - a framing protocol that runs on top of the T1 or E1 framing protocol.

PLM (Physical Layer Module) - interface card in the *CellPath 300* that provides the logic to support the physical layer of the network link. A PLM has the actual physical port mounted on it. Various PLMs support various physical layers, such as OC-3c/STM1 or DS3.

PLP (Packet Level Protocol) - Network layer protocol in the X.25 protocol stack. Sometimes called X.25 Level 3 or X.25 Protocol. See also X.25.

PM (Protocol Module) - interface card in the *CellPath 300* that provides the logic supporting the protocol layer of the network link. Various PMs support various protocols, such as ATM cell, Frame Relay, or CBR traffic.

PMD (Physical Medium Dependent) - a sublayer concerned with the bit transfer between two network nodes. It deals with wave shapes, timing recovery, line coding, and electro-optic conversions for fiber based links.

PNNI (Private Network Node Interface or Private Network-to-Network Interface) - a protocol that defines the interaction of private ATM switches or groups of private ATM switches

ping (Packet Internet Groper) - a program used to test reachability of destinations by sending them an ICMP echo request and waiting for a reply.

Point-to-Multipoint Connection - a collection of associated ATM VC or VP links, with associated endpoint nodes, with the following properties:

1. One ATM link, called the Root Link, serves as the root in a simple tree topology. When the Root node sends information, all of the remaining nodes on the connection, called Leaf nodes, receive copies of the information.
2. Each of the Leaf Nodes on the connection can send information directly to the Root Node. The Root Node cannot distinguish which Leaf is sending information without additional (higher layer) information. (See the following note for Phase 1.)
3. The Leaf Nodes cannot communicate directly to each other with this connection type.

Note: Phase 1 signalling does not support traffic sent from a Leaf to the Root.

Point-to-Point Connection - a connection with only two endpoints.

Point of Demarcation - the dividing line between a carrier and the customer premise that is governed by strict standards that define the characteristics of the equipment on each side of the demarcation. Equipment on one side of the point of demarcation is the responsibility of the customer. Equipment on the other side of the point of demarcation is the responsibility of the carrier.

Policing - the function that ensures that a network device does not accept traffic that exceeds the configured bandwidth of a connection.

Primitive - an abstract, implementation-independent interaction between a layer service user and a layer service provider.

Priority - the parameter of ATM connections that determines the order in which they are reduced from the peak cell rate to the sustained cell rate in times of congestion. Connections with lower priority (4 is low, 1 is high) are reduced first.

PROM (Programmable Read-Only Memory) - a chip-based information storage area that can be recorded by an operator but erased only through a physical process.

Protocol - a set of rules and formats (semantic and syntactic) that determines the communication behavior of layer entities in the performance of the layer functions.

Protocol Control Information - the information exchanged between corresponding entities using a lower layer connection to coordinate their joint operation.

Proxy - the process in which one system acts for another system to answer protocol requests.

Proxy Agent - an agent that queries on behalf of the manager, used to monitor objects that are not directly manageable.

PSN (Packet Switched Network) - a network designed to carry data in the form of packets. The packet and its format is internal to that network.

PT (Payload Type) - bits 2...4 in the fourth byte of an ATM cell header. The PT indicates the type of information carried by the cell. At this time, values 0...3 are used to identify various types of user data, values 4 and 5 indicate management information, and values 6 and 7 are reserved for future use.

PVC (Permanent Virtual Circuit (or Channel)) - a circuit or channel through an ATM network provisioned by a carrier between two endpoints; used for dedicated long-term information transport between locations.

Q.2931 - Derived from Q.93B, the narrowband ISDN signalling protocol, an ITU standard describing the signalling protocol to be used by switched virtual circuits on ATM LANs.

Real-Time Clock - a clock that maintains the time of day, in contrast to a clock that is used to time the electrical pulses on a circuit.

Relaying - a function of a layer by means of which a layer entity receives data from a corresponding entity and transmits it to another corresponding entity.

RFCs (Requests For Comment) - IETF documents suggesting protocols and policies of the Internet, inviting comments as to the quality and validity of those policies. These comments are collected and analyzed by the IETF in order to finalize Internet standards.

RFI (Radio Frequency Interference) - the unintentional transmission of radio signals. Computer equipment and wiring can both generate and receive RFI.

RIP (Routing Information Protocol) - a distance vector-based protocol that provides a measure of distance, or hops, from a transmitting workstation to a receiving workstation.

RISC (Reduced Instruction Set Computer) - a generic name for CPUs that use a simpler instruction set than more traditional designs.

Router - a device that forwards traffic between networks or subnetworks based on network layer information.

RTS (Request To Send) - an RS-232 modem interface signal (sent from the DTE to the modem on pin 4) which indicates that the DTE has data to transmit.

SBus - hardware interface for add-in boards in later-version Sun 3 workstations.

SAP (Service Access Point) - the point at which an entity of a layer provides services to its LM entity or to an entity of the next higher layer.

SAR (Segmentation And Reassembly) - the SAR accepts PDUs from the CS and divides them into very small segments (44 bytes long). If the CS-PDU is less than 44 bytes, it is padded to 44 with zeroes. A two-byte header and trailer are added to this basic segment. The header identifies the message type (beginning, end, continuation, or single) and contains sequence numbering and message identification. The trailer gives the SAR-PDU payload length, exclusive of pad, and contains a CRC check to ensure the SAR-PDU integrity. The result is a 48-byte PDU that fits into the payload field of an ATM cell.

SC - *CellPath* 300 System Controller; paired with the Extension Module (EM).

SCR (sustainable cell rate) - parameter defined by the ATM Forum for ATM traffic management. For VBR connections, SCR determines the long-term average cell rate that can be transmitted.

SCSI (Small Computer Systems Interface) - a standard for a controller bus that connects disk drives and other devices to their controllers on a computer bus. It is typically used in small systems.

SDLC (Synchronous Data Link Control) - IBM's data link protocol used in SNA networks.

SDU (Service Data Unit) - a unit of interface information whose identity is preserved from one end of a layer connection to the other.

SEAL (Simple and Efficient Adaptation Layer) - also called AAL 5, this ATM adaptation layer assumes that higher layer processes will provide error recovery, thereby simplifying the SAR portion of the adaptation layer. Using this AAL type packs all 48 bytes of an ATM cell information field with data. It also assumes that only one message is crossing the UNI at a time. That is, multiple end-users at one location cannot interleave messages on the same VC, but must queue them for sequential transmission.

Segment - a single ATM link or group of interconnected ATM links of an ATM connection.

Semipermanent Connection - a connection established via a service order or via network management.

SES (Severely Errored Seconds) - a second during which more event errors have occurred than the SES threshold.

SF (Superframe) - Common framing type used on T1 circuits. SF consists of 12 frames of 192 bits each, with the 193rd bit providing error checking and other functions. SF has been superseded by ESF, but is still widely used. Also called *D4 framing*. See also ESF.

SGMP (Simple Gateway Management Protocol) - the predecessor to SNMP.

Shaping Descriptor - *n* ordered pairs of GCRA parameters (I,L) used to define the negotiated traffic shape of an APP connection. The traffic shape refers to the load-balancing of a network. In this context, load-balancing means configuring the data flows to maximize the efficiency of the network.

SIR (Sustained Information Rate) - the long-term average data transmission rate across the User-to-Network Interface.

SMDS (Switched Multimegabit Data Service) - a high-speed, datagram-based, public data network service expected to be widely used by telephone companies in their data networks.

SMTP (Simple Mail Transfer Protocol) - the Internet electronic mail protocol used to transfer electronic mail between hosts.

SNAP - SubNetwork Access Protocol

SNMP (Simple Network Management Protocol) - the Internet standard protocol for managing nodes on an IP network.

snmpd - an SNMP agent for a given adapter card.

SONET (Synchronous Optical Network) - a new and growing body of standards that defines all aspects of transporting and managing digital traffic over optical facilities in the public network.

Source Traffic Descriptor - a set of traffic parameters belonging to the ATM Traffic Descriptor used during the connection set-up to capture the intrinsic traffic characteristics of the connection requested by the source.

Spanning Tree Protocol - provides loop-free topology in a network environment where there are redundant paths.

SPANS (Simple Protocol for ATM Network Signalling) - FORE Systems' proprietary signalling protocol used for establishing SVCs between FORE Systems equipment.

SPARC (Scalable Processor Architecture Reduced instruction set Computer) - a powerful workstation similar to a reduced-instruction-set-computing (RISC) workstation.

SPE (Synchronous Payload Envelope) - the payload field plus a little overhead of a basic SONET signal.

SPVC (Smart PVC) - a generic term for any communications medium which is permanently provisioned at the end points, but switched in the middle. In ATM, there are two kinds of SPVCs: smart permanent virtual path connections (SPVPCs) and smart permanent virtual channel connections (SPVCCs).

Static Route - a route that is entered manually into the routing table.

Statistical Multiplexing - a technique for allowing multiple channels and paths to share the same link, typified by the ability to give the bandwidth of a temporarily idle channel to another channel.

STM (Synchronous Transfer Mode) - a transport and switching method that depends on information occurring in regular and fixed patterns with respect to a reference such as a frame pattern.

STP (Shielded Twisted Pair) - two or more insulated wires that are twisted together and then wrapped in a cable with metallic braid or foil to prevent interference and offer noise-free transmissions.

STS (Synchronous Transport Signal) - a SONET electrical signal rate.

Sublayer - a logical subdivision of a layer.

Super User - a login ID that allows unlimited access to the full range of a device's functionality, including especially the ability to reconfigure the device and set passwords.

SVC (Switched Virtual Circuit (or Channel)) - a channel established on demand by network signalling, used for information transport between two locations and lasting only for the duration of the transfer; the datacom equivalent of a dialed telephone call.

Switched Connection - a connection established via signalling.

Symmetric Connection - a connection with the same bandwidth value specified for both directions.

Synchronous - signals that are sourced from the same timing reference and hence are identical in frequency.

Systems Network Architecture (SNA) - a proprietary networking architecture used by IBM and IBM-compatible mainframe computers.

T1 - a specification for a transmission line. The specification details the input and output characteristics and the bandwidth. T1 lines run at 1.544 Mbps and provide for 24 data channels. In common usage, the term “T1” is used interchangeably with “DS1.”

T3 - a specification for a transmission line, the equivalent of 28 T1 lines. T3 lines run at 44.736 Mbps. In common usage, the term “T3” is used interchangeably with “DS3.”

Tachometer - in *ForeView*, the tachometer shows the level of activity on a given port. The number in the tachometer shows the value of a chosen parameter in percentage, with a colored bar providing a semi-logarithmic representation of that percentage.

TAXI (Transparent Asynchronous Transmitter/Receiver Interface) - Encoding scheme used for FDDI LANs as well as for ATM; supports speeds of up to 100 Mbps over multimode fiber.

TC (Transmission Convergence) - generates and receives transmission frames and is responsible for all overhead associated with the transmission frame. The TC sublayer packages cells into the transmission frame.

TCP (Transmission Control Protocol) - a specification for software that bundles and unbundles sent and received data into packets, manages the transmission of packets on a network, and checks for errors.

TCP/IP (Transmission Control Protocol/Internet Protocol) - a set of communications protocols that has evolved since the late 1970s, when it was first developed by the Department of Defense. Because programs supporting these protocols are available on so many different computer systems, they have become an excellent way to connect different types of computers over networks.

TDM (Time Division Multiplexing) - a method of traditional digital multiplexing in which a signal occupies a fixed, repetitive time slot within a higher-rate signal.

Telnet - a TCP/IP protocol that defines a client/server mechanism for emulating directly-connected terminal connections.

Token Ring - a network access method in which the stations circulate a token. Stations with data to send must have the token to transmit their data.

topology - a program that displays the topology of a FORE Systems ATM network. An updated topology can be periodically re-displayed by use of the interval command option.

Traffic - the calls being sent and received over a communications network. Also, the packets that are sent on a data network.

Trailer - the protocol control information located at the end of a PDU.

Transit Delay - the time difference between the instant at which the first bit of a PDU crosses one designated boundary, and the instant at which the last bit of the same PDU crosses a second designated boundary.

trap - a program interrupt mechanism that automatically updates the state of the network to remote network management hosts. The SNMP agent on the switch supports these SNMP traps.

UAS (Unavailable Seconds) - a measurement of signal quality. Unavailable seconds start accruing when ten consecutive severely errored seconds occur.

UBR (Unspecified Bit Rate) - a type of traffic that is not considered time-critical (e.g., ARP messages, pure data), allocated whatever bandwidth is available at any given time. UBR traffic is given a “best effort” priority in an ATM network with no guarantee of successful transmission.

UDP (User Datagram Protocol) - the TCP/IP transaction protocol used for applications such as remote network management and name-service access; this lets users assign a name, such as “RVAX*2,S,” to a physical or numbered address.

Unassigned Cells - a generated cell identified by a standardized virtual path identifier (VPI) and virtual channel identifier (VCI) value, which does not carry information from an application using the ATM Layer service.

UNI (User-to-Network Interface) - the physical and electrical demarcation point between the user and the public network service provider.

UNI 3.0 - the User-to-Network Interface standard set forth by the ATM Forum that defines how private customer premise equipment interacts with private ATM switches.

UPC (Usage Parameter Control) - the mechanism that ensures that traffic on a given connection does not exceed the contracted bandwidth of the connection. UPC is responsible for policing or enforcement. UPC is sometimes confused with congestion management, to which it is functionally related on the *CellPath 300* (see congestion management).

UTP (Unshielded Twisted Pair) - a cable that consists of two or more insulated conductors in which each pair of conductors are twisted around each other. There is no external protection and noise resistance comes solely from the twists.

V.35 - ITU-T standard describing a synchronous, physical layer protocol used for communications between a network access device and a packet network. V.35 is most commonly used in the United States and Europe, and is recommended for speeds up to 48 Kbps.

VBR (Variable Bit Rate) - a type of traffic that, when sent over a network, is tolerant of delays and changes in the amount of bandwidth it is allocated (e.g., data applications).

VC (Virtual Channel (or Circuit)) - a communications path between two nodes identified by label rather than fixed physical path.

VCC (Virtual Channel Connection) - a unidirectional concatenation of VCLs that extends between the points where the ATM service users access the ATM Layer. The points at which the ATM cell payload is passed to, or received from, the users of the ATM Layer (i.e., a higher layer or ATMM-entity) for processing signify the endpoints of a VCC.

VCI (Virtual Channel Identifier) - the address or label of a VC; a value stored in a field in the ATM cell header that identifies an individual virtual channel to which the cell belongs. VCI values may be different for each data link hop of an ATM virtual connection.

VCL (Virtual Channel Link) - a means of unidirectional transport of ATM cells between the point where a VCI value is assigned and the point where that value is translated or removed.

VINES (Virtual Network Software) - Banyan's network operating system based on UNIX and its protocols.

Virtual Channel Switch - a network element that connects VCLs. It terminates VPCs and translates VCI values. The Virtual Channel Switch is directed by Control Plane functions and relays the cells of a VC.

Virtual Connection - an endpoint-to-endpoint connection in an ATM network. A virtual connection can be either a virtual path or a virtual channel.

Virtual Path Switch - a network element that connects VPLs, it translates VPI (not VCI) values and is directed by Control Plane functions. The Virtual Path Switch relays the cells of a Virtual Path.

VPT (Virtual Path Terminator) - a system that unbundles the VCs of a VP for independent processing of each VC.

VP (Virtual Path) - a unidirectional logical association or bundle of VCs.

VPC (Virtual Path Connection) - a concatenation of VPLs between virtual path terminators (VPTs). VPCs are unidirectional.

VPDN (Virtual Private Data Network) - a private data communications network built on public switching and transport facilities rather than dedicated leased facilities such as T1s.

VPI (Virtual Path Identifier) - the address or label of a particular VP; a value stored in a field in the ATM cell header that identifies an individual virtual path to which the cell belongs. A virtual path may comprise multiple virtual channels.

VPL (Virtual Path Link) - a means of unidirectional transport of ATM cells between the point where a VPI value is assigned and the point where that value is translated or removed.

VPN (Virtual Private Network) - a private voice communications network built on public switching and transport facilities rather than dedicated leased facilities such as T1s.

VT (Virtual Tributary) - a structure used to carry payloads such as DS1s that run at significantly lower rates than STS-1s.

WAN (Wide-Area Network) - a network that covers a large geographic area.

Warm Start Trap - a *CellPath* 300 SNMP trap that indicates that SNMP alarm messages or agents have been enabled.

Yellow Alarm - an alarm that occurs on a device when the signal from the device is not received at the far-end.

X.21 - ITU-T standard for serial communications over synchronous digital lines. The X.21 protocol is used primarily in Europe and Japan.

X.25 - ITU-T standard that defines how connections between DTE and DCE are maintained for remote terminal access and computer communications in PDNs. X.25 specifies LAPB, a data link protocol, and PLP, a network layer protocol. Frame Relay has, to some degree, superseded X.25. See also Frame Relay, LAPB, and PLP.

Glossary

Index

A

- address database
 - aging of learned addresses 3 - 8
 - and security 1 - 8
 - displaying 3 - 26, 5 - 19
 - learning process 1 - 7
 - modifying 3 - 27, 5 - 19
- atm uplink
 - attaching to packetbus 3 - 33
 - configuring 3 - 33
 - enabling 3 - 33
- atmuplink disable* 5 - 2
- atmuplink enable* 5 - 2
- atmuplink help* 5 - 2

B

- backplane, description of 1 - 6
- bandwidth, using parallel cables to increase 1 - 23
- broadcast packets, forwarding of 3 - 30
- bus assignment 5 - 15

C

- caution statement, definition of v
- chassis configuration, testing 4 - 9
- chassis show* 4 - 2
- check config* 4 - 9
- Command Help 2 - 2
- commands, operator console 5 - 29
 - atmuplink disable* 5 - 2
 - atmuplink enable* 5 - 2
 - atmuplink help* 5 - 2
 - card default* 5 - 4

- card disable* 5 - 4
- card enable* 5 - 4
- card help* 5 - 4
- card reset* 5 - 4
- card serial* 5 - 4
- card show* 5 - 4
- chassis help* 5 - 6
- chassis show* 4 - 2, 5 - 6
- check config* 4 - 9, 5 - 7
- check help* 5 - 7
- community help* 2 - 2, 5 - 9
- community ip* 2 - 8, 5 - 9
- community name* 2 - 7, 5 - 9
- community priv* 2 - 7, 5 - 9
- community show* 2 - 6, 5 - 9
- ebpseg show* 2 - 14
- esbpbus help* 5 - 12
- esbpbus memo* 5 - 12
- esbpbus show* 5 - 2, 5 - 12
- escam age* 5 - 13
- escam counters* 5 - 13
- escam default* 5 - 13
- escam disable* 5 - 13
- escam enable* 5 - 13
- escam help* 5 - 13
- escam reset* 5 - 13
- escam show* 5 - 13
- esgroup age* 3 - 8, 5 - 15
- esgroup bus* 5 - 15
- esgroup default* 3 - 7, 5 - 15
- esgroup disable* 3 - 6, 5 - 15
- esgroup enable* 3 - 6, 5 - 15

- esgroup help* 5 - 15
- esgroup reset* 3 - 7, 5 - 15
- esgroup show* 3 - 4, 5 - 15
- esport addr* 3 - 27, 5 - 19
- esport camshow* 5 - 19
- esport default* 3 - 15, 5 - 19
- esport disable* 3 - 11, 5 - 17, 5 - 19
- esport duplex* 3 - 23, 5 - 19
- esport enable* 3 - 11, 5 - 17
- esport forwarding* 3 - 21, 3 - 30, 3 - 37,
5 - 20
- esport group* 3 - 16, 5 - 19
- esport help* 5 - 17, 5 - 19
- esport learning* 3 - 29, 5 - 19
- esport memo* 3 - 17, 5 - 19
- esport priority* 3 - 14, 3 - 18, 5 - 19
- esport reset* 3 - 15, 5 - 19
- esport show* 3 - 9, 5 - 17, 5 - 19
- esport showaddr* 3 - 26, 5 - 19
- esport showx* 5 - 19
- esport sniff* 3 - 21, 5 - 19
- esport speed* 3 - 24, 5 - 19
- esport vlan* 3 - 32, 5 - 19
- group default* 5 - 25
- group disable* 5 - 25
- group enable* 5 - 25
- group help* 5 - 25
- group reset* 5 - 25
- group segment* 5 - 25
- group show* 5 - 25
- help* 2 - 2
- if address* 2 - 11, 5 - 27
- if disable* 5 - 27
- if enable* 5 - 27
- if help* 5 - 27
- if mode* 2 - 15, 5 - 27
- if netmask* 2 - 13, 5 - 27
- if segment* 2 - 13, 5 - 27
- if show* 2 - 10, 5 - 27
- if stats* 4 - 8, 5 - 27
- port default* 5 - 29
- port disable* 5 - 29
- port enable* 5 - 29
- port help* 5 - 29
- port memo* 5 - 29
- port priority* 5 - 29
- port reset* 5 - 29
- port show* 5 - 29
- quit* 2 - 6
- reboot* 2 - 33
- rmon default* 4 - 4, 5 - 31
- rmon help* 5 - 31
- rmon host* 4 - 5, 5 - 33
- rmon long* 4 - 6, 5 - 36
- rmon macip* 4 - 6, 5 - 39
- rmon matrix* 4 - 5, 5 - 42
- rmon short* 4 - 6, 5 - 45
- rmon show* 4 - 4, 5 - 31
- rmon stats* 4 - 6, 5 - 48
- route add* 2 - 17, 5 - 51
- route delete* 2 - 18, 5 - 51
- route delete default* 2 - 18, 5 - 51
- route help* 5 - 51
- route show* 5 - 51
- setup* 2 - 2, 5 - 54
- snmp auth-key* 2 - 29, 5 - 55
- snmp default* 2 - 27, 2 - 28, 5 - 55
- snmp help* 2 - 26, 5 - 55
- snmp party-init* 2 - 27, 2 - 28
- snmp re-sync* 2 - 31
- snmp show* 2 - 27, 2 - 28, 5 - 55
- snmp v1 disable* 2 - 30, 5 - 55
- snmp v1 enable* 2 - 30, 5 - 55
- stbridge age* 5 - 58

- stbridge default* 5 - 58
 - stbridge disable* 5 - 58
 - stbridge enable* 5 - 58
 - stbridge forwarddelay* 5 - 58
 - stbridge hello* 5 - 58
 - stbridge help* 5 - 58
 - stbridge newroot* 5 - 58
 - stbridge priority* 5 - 58
 - stbridge reset* 5 - 58
 - stbridge topchange* 5 - 58
 - stport default* 5 - 63
 - stport disable* 5 - 63
 - stport enable* 5 - 63
 - stport help* 5 - 63
 - stport pathcost* 5 - 63
 - stport priority* 5 - 63
 - stport reset* 5 - 63
 - stport show* 5 - 63
 - stport status* 5 - 63
 - system contact* 4 - 3, 5 - 67
 - system help* 5 - 67
 - system location* 4 - 3, 5 - 67
 - system name* 4 - 3, 5 - 67
 - system show* 4 - 2, 5 - 67
 - trap add* 2 - 20, 5 - 69
 - trap community* 2 - 21, 5 - 69
 - trap delete* 2 - 21, 5 - 69
 - trap help* 5 - 69
 - trap show* 2 - 20, 5 - 69
 - version* 4 - 2, 5 - 71
 - community ip* 2 - 8
 - community name* 2 - 7
 - community priv* 2 - 7
 - community show* 2 - 6
 - community strings
 - assigning 5 - 7
 - modifying 2 - 6
 - same name 5 - 8
 - configuration
 - controller module 2 - 1
 - errors 5 - 7
 - connecting switches 3 - 35
 - with parallel cables 1 - 23
 - with VLAN tagging 3 - 37
 - controller module
 - configuration 2 - 1
 - reboot 2 - 33
 - setup** command 2 - 2
 - controllers
 - NMM-1 1 - 3
 - NMM-2 1 - 4
 - NMM-SEG-1 1 - 5
- ## D
- default** commands, reasons for use 3 - 2
 - default gateway, defining 2 - 16
 - default netmask 2 - 13
 - defaults, setting RMON 4 - 4
 - destination address 1 - 13
 - diagnostics
 - checking the configuration 5 - 7
 - Self-Test-Diagnostic (STD) test 4 - 9
 - disable
 - groups 3 - 6, 5 - 15
 - History table creation 4 - 6
 - Hosts table creation 4 - 5
 - Matrix table creation 4 - 5
 - SNMPv1, displaying 2 - 6
 - Statistics table creation 4 - 6
 - user ports 3 - 11, 5 - 15
 - duplex mode 1 - 7, 3 - 23, 5 - 19
- ## E
- enable
 - groups 3 - 6, 5 - 15

- History table creation 4 - 6
- Hosts table creation 4 - 5
- Matrix table creation 4 - 5
- SNMPv1 2 - 30
- Statistics table creation 4 - 6
- user ports 3 - 11, 5 - 15
- error checking 1 - 13
- errors in configuration 5 - 7
- ES-4810
 - main components 1 - 1
 - user module features 1 - 7
- esbpbus help** 5 - 12
- esbpbus memo** 5 - 12
- esbpbus show** 5 - 2, 5 - 12
- esgroup age** 3 - 8, 5 - 15
- esgroup bus** 5 - 15
- esgroup default** 3 - 7, 5 - 15
- esgroup disable** 3 - 6, 5 - 15
- esgroup enable** 3 - 6, 5 - 15
- esgroup help** 5 - 15
- esgroup reset** 3 - 7, 5 - 15
- esgroup show** 3 - 4, 5 - 15
- esp 5 - 20
- esport addr** 3 - 27, 5 - 19
- esport camshow** 5 - 19
- esport default** 3 - 15, 3 - 18, 5 - 19
- esport disable** 3 - 11, 5 - 17, 5 - 19
- esport duplex** 3 - 23, 5 - 19
- esport enable** 3 - 11, 5 - 17
- esport forwarding** 3 - 21, 3 - 30, 3 - 37, 5 - 20
- esport group** 3 - 16, 5 - 19
- esport help** 5 - 17, 5 - 19
- esport learning** 3 - 29, 5 - 19
- esport memo** 3 - 17, 5 - 19
- esport priority** 3 - 14
- esport reset** 3 - 15, 5 - 19
- esport show** 3 - 9, 5 - 17, 5 - 19
- esport showaddr** 3 - 26, 5 - 19
- esport showx** 5 - 19
- esport sniff** 3 - 21, 5 - 19
- esport speed** 3 - 24, 5 - 19
- esport vlan** 3 - 32, 5 - 19
- exiting the console interface 2 - 6
- F**
- forwarding mode
 - See also* uplink
 - setting 3 - 21, 3 - 30, 3 - 37, 5 - 20
- G**
- groups 3 - 4
 - number of 3 - 4
 - specifying number in commands .. 3 - 3
 - stand-alone mode 3 - 4
- H**
- help** command 2 - 2
- History table, RMON, creation 4 - 6
- Hosts table, RMON, creation 4 - 5
- I**
- if address** 2 - 11
- if mode** 2 - 15, 5 - 27
- if netmask** 2 - 13
- if segment** 2 - 13
- if show** 2 - 10, 5 - 27
- if stats** 4 - 8, 5 - 27
- interswitch link 1 - 11
- IP netmask
 - description of 2 - 11
 - setting 2 - 11
- IP netmask, default value 2 - 13
- L**
- learning mode, setting 3 - 29, 5 - 19
- learning process 1 - 7
- log table, RMON, displaying configuration of 4

- 6
- loops 1 - 20, 3 - 35
- M**
- management interfaces 1 - 10
- Matrix table, RMON, creation 4 - 5
- memo
 - user port 3 - 17
- MIB support 1 - 10
- modes, normal and promiscuous 2 - 15, 5 - 28
- monitoring mode 1 - 9, 3 - 21, 5 - 19
- N**
- name, community
 - assigning IP addresses 2 - 8
 - assigning 2 - 6
 - check to see if assigned 5 - 7
 - checking for same name 5 - 8
 - defining privilege level 2 - 7
 - displaying 2 - 6
 - modifying 2 - 7
- netmask
 - default value 2 - 13
 - description of 2 - 11
 - setting 2 - 11
- netmask parameter definition 5 - 28
- normal mode, frame copy 2 - 15, 5 - 28
- note statement, definition of v
- P**
- packet bus 1 - 6
 - assigning group to 3 - 4
 - capturing traffic on 1 - 9
- port
 - default 5 - 19
 - priority 3 - 18, 5 - 19
 - reset 3 - 15, 5 - 19
 - setting to default configuration ... 3 - 15
 - specifying number in commands .. 3 - 3
- speed 3 - 24, 5 - 19
- status, displaying 3 - 9
- types 1 - 7
- port group* 5 - 29
- promiscuous mode 2 - 15, 5 - 28
- Q**
- quit* 2 - 6
- R**
- reboot* 2 - 33
- Remote Network Monitoring (RMON)
 - configuration 4 - 4
 - see RMON
- reset* commands, reasons for use 3 - 2
- RMON configuration, displaying 4 - 4
- rmon default* 4 - 4
- rmon show* 4 - 4
- route delete default* 2 - 18, 5 - 51
- routes
 - default 2 - 16, 5 - 51
 - static 2 - 16, 5 - 51
- S**
- security 1 - 8, 3 - 25
- segment switch manager
 - description 1 - 5
 - forwarding mode 1 - 19
 - limitations 1 - 5, 1 - 7
- Self-Test-Diagnostic (STD) test 4 - 9
- setup* command 2 - 2
- snmp auth-key* 2 - 29
- snmp default* 2 - 27, 2 - 28
- snmp help* 2 - 26
- snmp party-init* 2 - 27, 2 - 28
- snmp re-sync* 2 - 31
- snmp show* 2 - 27, 2 - 28
- snmp v1 disable* 2 - 30
- snmp v1 enable* 2 - 30

Index

- source address 1 - 7
- Spanning Tree Algorithm 3 - 35
- speed 3 - 24, 5 - 19
- stand-alone mode 3 - 4
- station migration 1 - 8
- statistics counters 1 - 9
- Statistics table, RMON, creation 4 - 6
- store-and-forward process 1 - 13
- super* access level 2 - 6, 5 - 7
- switch monitoring 1 - 9
- system contact* 4 - 3
- system location* 4 - 3
- system name* 4 - 3
- system show* 4 - 2

- T**
- TFTP 2 - 6
- Transparent bridging 1 - 12
- trap add* 2 - 20
- trap community* 2 - 21
- trap delete* 2 - 21
- trap show* 2 - 20
- troubleshooting 5 - 7

- U**
- unknown unicasts, forwarding of 3 - 30
- uplink
 - between ES-4810 and non-ES-4810 switch
3 - 36
 - between ES-4810 Ethernet modules 1 - 9
 - between ES-4810s 3 - 37
 - configuring 3 - 35
 - description 1 - 9, 1 - 11, 3 - 35

- V**
- version* 4 - 2, 5 - 71
- VLAN tagging 1 - 16, 3 - 37
- VLANs
 - configuring 3 - 32, 5 - 19

- description 1 - 14

- W**
- warnings, definition of v