

DIGITAL VNswitch 900 Series

Switch Management

Part Number: AA-R2LDE-TE

June 1998

This manual describes how to configure, monitor, and manage a VNswitch 900 series switch.

Revision/Update Information:	This is a revised manual.
Software and Version	VNswitch 900 series Version 3.0

Digital Equipment Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

Possession, use, or copying of the software described in this publication is authorized only pursuant to a valid written license from DIGITAL or an authorized sublicensor.

© Digital Equipment Corporation 1996, 1997, 1998. All rights reserved. Printed in U.S.A.

The following are trademarks of Digital Equipment Corporation:

clearVISN, DECconnect, DEChub, DECreeper, DEC, DECnet, DIGITAL, the DIGITAL logo, LAT

The following are third-party trademarks:

Apple and AppleTalk are registered trademark of Apple Computer, Inc.

Apollo is a registered trademark of Apollo Computer, Inc.

Banyan is a registered trademark of Banyan Systems, Inc.

NetBIOS is a trademark of Micro Computer Systems, Inc.

NetWare and Novell are registered trademarks of Novell, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Ltd.

Windows is a registered trademark and Windows 95, Windows NT, and Internet Explorer are trademarks of Microsoft Corporation.

Netscape is a registered trademark of Netscape Communications Corporation.

All other trademarks and registered trademarks are the property of their respective holders.

Contents

Preface

Overview	xv
Purpose of This Manual	xv
Intended Audience	xv
Organization	xvi
Conventions	xxviii
Associated Documents	xxix
Correspondence	xxii
Documentation Comments	xxii
Online Services	xxii
How to Order Additional Documentation	xxiii

1 Introduction

Overview	1-1
Introduction	1-1
In This Chapter	1-1
Concepts and Terminology	1-2
Bridge Plug and Play	1-2
Switch Console Sessions	1-2
Web-based Management Application	1-3
Transparent Bridging	1-3
Understanding Network Interfaces and Ports	1-3
User Interface	1-11
Types of Switch Memory	1-13
Configuration Commands Requiring Restart	1-13
Address Types	1-14

System Security	1-15
Overview of Your Responsibilities	1-16

2 Operational Basics

Overview	2-1
Introduction	2-1
In This Chapter.	2-1
Starting and Terminating Console Sessions	2-2
Starting and Terminating Local Sessions.	2-3
Starting and Terminating Remote Sessions	2-5
Accessing CLI Prompts.	2-7
Accessing the Main Prompt	2-7
Accessing the Config Prompt	2-8
Accessing the Monitor Prompt	2-8
Adding the Hostname to a Prompt.	2-9
Exiting a Prompt	2-10
Using the Command Line Interface	2-11
Using Command Line Editing	2-11
Using Command Line Recall.	2-12
Using Command Line Completion	2-12
Entering Commands and Command Shortcuts	2-14
Entering Subsystem Commands	2-15
Using Auto-Prompts	2-16
Using the Web-Based Management Application	2-18
Accessing VNswitch Modules Over the Web	2-18
Managing VNswitch Modules Over the Web	2-20
Accessing Web Help	2-20
Disabling and Enabling the VNswitch Web Server.	2-20
Displaying CLI Help.	2-21

3 Adding and Managing Users

Overview	3-1
Introduction	3-1
In This Chapter.	3-1
Adding Users.	3-2
Displaying a List of All Users.	3-4
Changing a User's Name, Password, and Security Level.	3-5
Changing Your Own Password	3-5
Changing Another User's Password or Security Level	3-6
Enabling and Disabling Prompting for ID and Password	3-8
Deleting Users.	3-9

Deleting a Single User	3-9
Deleting (Clearing) All Users	3-9

4 Configuring and Monitoring Switchwide Parameters

Overview	4-1
Introduction	4-1
In This Chapter	4-1
Resetting (Clearing) NVRAM to Default Values	4-2
Setting the Session Inactivity Timer	4-5
Assigning a Host Name to the VNswitch Module	4-6
Displaying General Information About the Switch	4-7
Information Displayed from the Config Prompt	4-7
Information Displayed from the Monitor Prompt	4-8
Setting and Viewing Clock Time	4-11
Setting the Time	4-11
Setting the Time Zone Offset	4-13
Setting Time Host Synchronization	4-14
Viewing Clock Time Parameters	4-15
Monitoring Switch Memory	4-16
Monitoring Crash Counts and Restart or Reload Data	4-18

5 Configuring Network Interfaces

Overview	5-1
Introduction	5-1
In This Chapter	5-1
Displaying the Interface Number and Type	5-3
Enabling and Disabling an Interface	5-4
Disabling an Interface	5-5
Accessing and Exiting an Interface Prompt	5-7
Accessing an Interface Prompt	5-7
Exiting an Interface Prompt	5-7
Configuring an Ethernet Logical Interface	5-8
Setting the Interface for Full-Duplex or Half-Duplex Mode	5-8
Redirecting Interfaces to the DEChub 900 Backplane	5-9
Redirecting Interface 12 to the DEChub 900 Backplane ThinWire	5-9
Redirecting Interface 11 to a DEChub ONE AUI Port	5-9
Displaying Current Ethernet Configuration Parameters	5-10
Configuring a Fast Ethernet Logical Interface	5-11
IEEE Compliance	5-11
Enabling or Disabling Auto-Negotiation	5-11
Restarting Auto-Negotiation	5-13

Setting Advertised Connection Capabilities	5-13
Setting Full-Duplex or Half-Duplex Mode Manually	5-14
Setting the Line Speed Manually	5-15
Displaying Current Fast Ethernet Configuration Parameters	5-16
Configuring an FDDI Logical Interface	5-18
Supported Station Types	5-18
PHY Port Identification	5-19
Setting the Station Type	5-22
Setting the Link Error Rate Alarm	5-22
Automatically Disconnecting Nodes Causing Excessive Link Errors	5-23
Enabling and Disabling SMT Notification	5-24
Setting Token Passing and Frame Timing Parameters	5-24
Configuring the Interface to Purge Bad Frames From the Ring	5-26
Setting the Interface for Full-Duplex or Half-Duplex Mode	5-27
Resetting All Configuration Parameters to Default Values	5-28
Displaying Current FDDI Interface Configuration Parameters	5-29
Configuring ATM Physical and Logical Interfaces	5-31
Configuring the Physical Interface	5-31
Configuring OC3 Interfaces	5-32
Configuring DS1 Interfaces for T1/E1 Links	5-38
Configuring DS3 Interfaces for T3/E3 Links	5-41
Configuring a Logical Interface	5-47
Displaying ATM Physical Interface and Logical Port Information	5-72
Configuring the Address Resolution Protocol	5-73
Adding and Changing ARP Cache Entries Manually	5-73
Deleting a Manually Entered ARP Entry	5-74
Managing the Time That Learned Entries Are Retained	5-75
Displaying Information About the Current Configuration	5-77

6 Monitoring Network Interfaces

Overview	6-1
Introduction	6-1
In This Chapter	6-1
Displaying the Interface Number and Type	6-2
Monitoring an Ethernet Interface	6-5
Monitoring a Fast Ethernet Interface	6-10
Monitoring an FDDI Interface	6-15
Monitoring an ATM Interface	6-20
Displaying Interface Test Results and MAC Type	6-20
Monitoring the Physical Interface	6-20
Monitoring a Logical Interface	6-25
Monitoring the LAN Emulation ARP Database	6-29
Monitoring Packet Statistics and Error Counts	6-31

Displaying Packet Buffer Data	6-31
Packet Error Statistics	6-34
Input and Output Queues	6-36
Packet Statistics	6-38
Displaying Interface Test Results	6-40
Clearing Interface Counters	6-41
Testing an Interface	6-42
Monitoring the Address Resolution Protocol	6-43
Monitoring ICMP Counters	6-46

7 Configuring the Transparent Bridge

Overview	7-1
Introduction	7-1
In This Chapter	7-1
Accessing and Exiting the Bridge Configuration Prompt	7-2
Accessing the Bridge Configuration Prompt	7-2
Exiting the Bridge Configuration Prompt	7-2
Setting and Enabling Switch-wide Rate Limiting	7-3
Setting Maximum Frames Per Second	7-3
Enabling and Disabling Rate Limiting	7-4
Configuring Permanent Address Filters	7-5
Source Address Filtering	7-5
Destination Address Filtering	7-5
Creating and Modifying a Permanent Address Filter	7-6
Deleting Permanent Address Filters	7-8
Configuring Protocol Filters	7-10
Creating and Modifying Protocol Filters	7-11
Deleting Protocol Filters	7-14
Creating and Modifying Default Protocol Filters	7-18
Deleting Default Protocol Filters	7-20
Configuring the Spanning Tree Protocol	7-21
Influencing Selection of the Root Bridge	7-22
Influencing Selection of the Root Port	7-22
Detecting Changes in Network Topology	7-23
Enabling and Disabling STP	7-26
Setting a Port to Start Forwarding Immediately (Fast Start)	7-27
Forwarding Using Only Manually Created Address Filters	7-30
Enabling Manual Mode	7-30
Disabling Manual Mode	7-31
Enabling and Disabling a Bridge Port	7-32
Enabling a Port	7-32
Disabling a Port	7-32
Bridging Ethernet and FDDI Networks	7-33

IP Fragmentation	7-33
Enabling and Disabling IPX Translation	7-33
Auto-Testing of Ports Inactive for Extended Periods	7-39
Setting the Time That Unused Addresses Are Retained	7-40
Displaying Current Bridge Configuration Parameters	7-41
Duplicate MAC Addresses on Separate VSDs	7-45
Configuring Duplicate MAC Addresses	7-46
Displaying Duplicate MAC Addresses	7-47
Deleting A Duplicate MAC Address	7-48

8 Monitoring the Transparent Bridge

Overview	8-1
Introduction	8-1
In This Chapter.	8-1
Accessing and Exiting the Bridge Monitor Prompt	8-2
Accessing the Bridge Monitor Prompt.	8-2
Exiting the Bridge Monitor Prompt	8-2
Monitoring the Bridge.	8-3
General Bridging Operation.	8-3
Port Activity Counters	8-4
Bridge Ports	8-6
MAC Address Database.	8-8
Protocol Filters.	8-10
Spanning Tree Protocol	8-15
Configuring a MAC Address As a Static Entry	8-18
Source Address Filtering	8-18
Destination Address Filtering	8-18
Creating and Modifying a Static MAC Address Filter	8-19
Deleting Static MAC Address Filters	8-21

9 Configuring Virtual LANs

Overview	9-1
Introduction	9-1
In This Chapter.	9-1
VLAN Secure Domains	9-3
VNbus	9-3
VNbus Tags	9-3
Default VSD.	9-3
Spanning Tree Protocol Support	9-4
Accessing and Exiting the VSD Configuration Prompt	9-5
Accessing the VSD Configuration Prompt	9-5

Exiting the VSD Configuration Prompt	9-5
Creating VSDs	9-6
VSDs Within a Single Module	9-8
VSDs Across Modules in a Single Hub	9-8
VSDs Across ATM Emulated LANs and Bridge Tunnels	9-9
Reserving VSDs	9-11
Modifying VSDs	9-12
Deleting VSDs	9-15
Displaying Information About VSDs	9-16
Warning Messages	9-17
Configuring the VNbus to Support VLANs	9-18
Establishing a LAN Segment on the VNbus	9-18
Disabling and Enabling a VNbus Port	9-19
Assigning a VNswitch IP End Node to a VSD	9-20
Restrictions	9-20
Assigning an IP Address and Subnet Mask	9-20
Selecting the VSD on Which the Host Resides	9-21

10 Performing Routine Maintenance

Overview	10-1
Introduction	10-1
In This Chapter	10-1
Accessing and Exiting the Boot Config Prompt	10-2
Accessing the Boot Config Prompt	10-2
Exiting the Boot Config Prompt	10-2
Restarting the Switch	10-3
How to Restart the Switch	10-3
Upgrading and Reinstalling Switch Software	10-4
Installing the Software	10-4
Configuring Installation File Locations	10-8
Displaying File Names and Server Locations	10-9
Modifying Installation File Locations	10-10
Canceling the Installation Procedure	10-12
Backing Up and Restoring the Switch	10-13
Automatic Image Recovery	10-13
Configuring Automatic Image Recovery	10-14
Backing Up Configuration Settings	10-14
Checking Available RAM	10-19
How to Check Available RAM	10-19
Capturing Restart or Crash Messages and Diagnostic Data	10-20
Displaying and Managing Restart or Crash Error Messages	10-20
Downloading Diagnostic Data for Problem Analysis	10-24
Displaying All Boot Config Settings	10-34

11 Event Logging and Reporting

Overview	11-1
Introduction	11-1
In This Chapter.	11-1
Event Messages and Related Concepts	11-2
Types of Events That Are Logged	11-2
Elements of an Event Message	11-2
Logging Levels and Event Types.	11-5
Preconfigured Logging Criteria (Groups)	11-7
Selecting Which Events Are Logged	11-8
Modes of Configuration	11-8
Commands Used to Log Events.	11-8
Configuring ELS in Nonvolatile Memory	11-10
Configuring ELS in Volatile Memory	11-23
Displaying the Event Log	11-41
Choosing the Method of Display	11-41
Displaying and Exiting the Event Log.	11-42
Advanced Methods for Viewing Events	11-43
Printing ELS Output	11-44

12 Configuring Remote Management

Overview	12-1
Introduction	12-1
In This Chapter.	12-1
Configuring the Module for Remote Console Sessions	12-2
Configuring TCP/IP Host Services	12-2
Configuring and Viewing OBM Port Connections	12-2
Redirecting Ethernet Interface 11	12-5
Configuring the Module for MIB-Based Management Systems	12-6
Configuring TCP/IP Host Services	12-6
Configuring SNMP	12-6
Configuring TCP/IP Host Services	12-20
Configuring the Module's In-Band IP Address and Subnet Mask	12-21
Assigning a VNswitch IP End Node to a VSD	12-21
Configuring a Default Gateway	12-22
Enabling and Disabling Router Discovery and RIP Listening	12-22
Enabling and Disabling Host Services.	12-23
Displaying TCP/IP Host Services Settings	12-25
Monitoring and Managing TCP/IP Host Services.	12-26
Displaying the Routing Table	12-26
Displaying the Bridge's Interface Addresses.	12-27

Testing a Network Connection Using the Ping Command	12-27
Displaying a List of Routers	12-28
Displaying the Path to a Destination Device	12-29
Displaying the VSD on Which TCP/IP Host Services Is Available	12-31
Connecting to Other Devices Using Telnet	12-33

13 Monitoring Network Activity

Overview	13-1
Introduction	13-1
In This Chapter	13-1
About the VNswitch Mirror Feature	13-2
Mirror Restrictions and Mirroring Effect on VNswitch Performance	13-3
Mirror Restrictions	13-3
Mirroring Effect on VNswitch Performance	13-3
Configuring the Mirror Interface	13-4
Accessing the Mirror Prompt	13-4
Disabling Bridging on the Destination Interface	13-4
Setting the Mirror Interface	13-5
Setting Multiple Mirror Interfaces	13-5
Changing the Mirror Destination Port	13-6
Listing Mirror and Destination Interfaces	13-6
Clearing Mirror and Destination Interfaces	13-6
Mirroring Multiple VNswitch Modules	13-8
Mirror Commands	13-10
About the VNswitch RMON Agent	13-11
RMON Alarm and Event Groups	13-11
RMON Command Line Interface	13-14
Accessing the RMON Configuration Process	13-14
Accessing the RMON Monitor Process	13-15
Clearing RMON Configuration Information	13-16
Displaying RMON Statistics	13-17
RMON Example	13-18
Configuring an SNMP Community	13-18
Configuring an RMON Trap	13-18

A Advanced Console Management

Overview	A-1
Introduction	A-1
In This Appendix	A-1
Process Aliases and IDs	A-2
Viewing Process Status and PIDs	A-3

Process List Output	A-3
Viewing Output from Multiple Processes	A-5
Canceling Display of Output to a Console Session	A-6
Terminating Process Output	A-7

B Plug and Play Default Settings

Overview	B-1
Introduction	B-1
In This Appendix	B-1
Switch-wide Default Settings	B-3
Nonspecific Interface Default Settings	B-4
Ethernet Interface Default Settings	B-5
Fast Ethernet Interface Default Settings	B-6
FDDI Interface Default Settings	B-7
ATM Interface Default Settings	B-8
VNbus Default Settings	B-11
Bridge Default Settings	B-12
Maintenance Default Settings	B-14
ELS Configuration Default Settings	B-15
Remote Management Default Settings	B-17

C VNswitch Packet Counters

Overview	C-1
Introduction	C-1
In This Appendix	C-1
Packet Counter Overview	C-2
Counter Relationships	C-5
Interface Counters	C-6
Bridge Port Counters	C-9
Counter Relationships	C-13
Management Interfaces	C-14
CLI	C-14

Index

Figures

1-1 Physical Interfaces	1-4
1-2 Logical Interfaces and Bridge Ports	1-6

1-3	VLAN Logical Interfaces	1-8
1-4	TCP/IP Host Services Logical Interface	1-9
2-1	VNswitch 900EA Installation Menu	2-4
2-2	DIGITAL MultiSwitch 900 Installation Menu	2-4
2-3	Web Management Display	2-19
5-1	VNswitch PHY Ports	5-19
5-2	Sample FDDI Configuration and Associated PHY Ports	5-21
5-3	ATM Bridge Tunnels	5-43
7-1	VNswitch 900EF Modules Requiring IPX Translation	7-35
9-1	ATM ELAN and Bridge Tunnel VSDs	9-10
10-1	Sample Crash Log	10-22
11-1	Sample Message Generated by an Event	11-3
13-1	Using clearVISON to Connect a Mirror Destination Interface	13-9
C-1	Packet Flow	C-3

Tables

1-1	MAC Address Type Descriptions	1-13
2-1	Command Line Editing Keys	2-11
2-2	Command Line Recall Keys	2-12
4-1	Clear Command Options and Descriptions	4-4
4-2	Description of the Configuration Report from the Monitor Prompt	4-10
4-3	Description of the Memory Report	4-16
5-1	Station Types and PHY Ports	5-20
5-2	List Command Options and Descriptions for FDDI Configurations	5-30
5-3	ATM Interface Defaults at Startup for EA Module	5-38
5-4	Guidelines for Determining Bridge Tunnel Type	5-42
5-5	ELAN Connection Phases	5-49
5-6	List Command Options and Descriptions for ATM Logical Configurations	5-60
5-7	List Command Options and Descriptions for ARP Configurations	5-67
6-1	Description of the Interface Information Displayed	6-3
6-2	Description of the Monitoring Information Displayed for Ethernet	6-6
6-3	Description of the Monitoring Information Displayed for Fast Ethernet	6-11
6-4	Description of the Monitoring Information Displayed for FDDI	6-16
6-5	ATM Physical Interface Command Options	6-22
6-6	ATM Physical Interface MIB-Statistics Command Options and Descriptions	6-24
6-7	ATM Logical Interface Command Options	6-25
6-8	Description of the Buffer Information Displayed	6-28
6-9	Description of the Packet Error Information Displayed	6-31
6-10	Description of the Input and Output Queue Information Displayed	6-33
6-11	Description of the Packet Statistics Displayed	6-35
6-12	ARP Monitor Command Options	6-39
7-1	Hexadecimal Values for Common Ethernet-II (IEEE 802.3) Protocols	7-14
7-2	Hexadecimal Values for Common SNAP OUI/IP Protocols	7-14

7-3	Hexadecimal Values for Common DSAP Protocols	7-15
7-4	IPX Translation Rules	7-33
7-5	List Command Options and Descriptions	7-41
8-1	List Counters Command Options and Descriptions	8-5
8-2	List Database Command Options and Descriptions	8-9
8-3	List Protocol Filter Command Options and Descriptions	8-11
8-4	Hexadecimal Values for Common Ethernet-II Protocols	8-13
8-5	Hexadecimal Values for Common DSAP Protocols	8-13
8-6	Hexadecimal Values for Common SNAP OUI/IP Protocols.	8-14
8-7	List STP Command Options and Descriptions	8-16
9-1	Create VSD Command Options.	9-6
9-2	Modify VSD Command Options.	9-13
10-1	LED Lighting Sequence During Load/Reload.	10-8
11-1	Event Subsystems and Associated Short Names	11-4
11-2	Logging Levels and Associated Event Types	11-6
11-3	Commands Used to Log and Trap Event Messages	11-9
11-4	List Command Options and Descriptions	11-22
11-5	List Command Options and Descriptions	11-37
12-1	Community Access Options	12-11
12-2	Trap Type Options	12-13
12-3	Community Command Options.	12-15
13-1	RMON Set Alarm Command Parameters	13-12
13-2	RMON Set Event Command Parameters.	13-13
13-3	RMON Configuration Commands.	13-15
13-4	RMON Monitor Commands.	13-16
B-1	Switch-wide Defaults.	B-3
B-2	Nonspecific Interface Defaults	B-4
B-3	Ethernet Defaults	B-5
B-4	Fast Ethernet Defaults	B-6
B-5	FDDI Defaults	B-7
B-6	Interface Defaults at Startup	B-9
B-7	ATM Physical Interface Default Parameters.	B-9
B-8	ATM Logical Interface Default Parameters	B-10
B-9	VNbus Defaults	B-11
B-10	Bridge Defaults	B-12
B-11	Maintenance Defaults	B-14
B-12	ELS NVRAM Configuration Defaults	B-15
B-13	ELS Volatile Memory Configuration Defaults	B-15
B-14	Remote Management Configuration Defaults	B-17
C-1	Interface Counters and Descriptions	C-7
C-2	Bridge Port Counters and Descriptions.	C-10

Preface

Overview

Purpose of This Manual

This manual provides instructions for configuring, monitoring, and managing the DIGITAL VNswitch 900 series of modules.

Intended Audience

This manual is intended for persons who install, configure, and manage computer networks. Although experience with computer networking hardware and software is helpful, you do not need programming experience to configure, monitor, and manage VNswitch modules.

Organization

This manual is organized as follows:

Section	Description
Chapter 1	Provides general information about the VNswitch 900 Series of products and an overview of what your responsibilities are as a VNswitch administrator. This chapter also provides an introduction to concepts and terminology with which you should be familiar.
Chapter 2	Describes operational basics that are common to many of the configuration and management tasks described throughout the manual.
Chapter 3	Provides instructions about assigning login permission and security levels to users who are to manage a VNswitch module.
Chapter 4	Provides information about how to manage and monitor functions that affect switchwide operation.
Chapter 5	Provides instructions about how to reconfigure Ethernet, Fast Ethernet, FDDI, ATM, and VNbus interface default settings to maximize network performance, and to accommodate requirements unique to your network environment.
Chapter 6	Provides information about how to monitor Ethernet, Fast Ethernet, FDDI, and ATM network interfaces.
Chapter 7	Provides information about how to configure the transparent bridge and how to display information about the current configuration.
Chapter 8	Provides information about how to monitor the transparent bridge and how to configure static addresses. Monitoring includes the ability to display specific information about operational states, activity counters, and various bridge configuration settings.

Section	Description
Chapter 9	Describes how to create, modify, and delete Virtual LANs (VLANs).
Chapter 10	Describes maintenance procedures you may need to perform periodically, and those that you may need to perform regularly.
Chapter 11	Provides information about how to configure an event log to record specific types of operational events and errors and to eliminate others, depending on the level of detail you require.
Chapter 12	Describes how to configure the VNswitch module so you can manage the module from a remote device.
Chapter 13	Explains how to configure the mirror port and RMON agent so that you can monitor network activity.
Appendix A	Discusses local and remote console session management tasks that may be of interest to advanced VNswitch module users.
Appendix B	Lists the factory default settings that are used when the switch is first installed. You may need to alter default settings to maximize network performance, to enable special functions such as Virtual Local Area Networks (VLANs), and to set IP addresses for remote management.
Appendix C	Describes VNswitch counters and their relationships.

Conventions

This manual uses the following conventions:

Convention	Description
Special Type	This special type in examples indicates system output.
Boldface	Boldface type indicates user input.
<i>Boldface Italics</i>	Boldface type in italics indicates variables for which the user or the system supplies a value
<u>Boldface underscore</u>	Underscored boldface characters indicate the least number of characters you must enter to identify a command. The underscored characters are referred to as command shortcuts. For example, the commands for listing users is <u>list users</u> , and can be entered as <u>l u</u> . Similarly, the command for viewing error statistics is <u>error</u> and can be entered as <u>er</u> .
Return	Indicates that you should press the Return key.
Ctrl/ <i>keystroke</i>	Indicates you should press the key specified by <i>keystroke</i> while holding down the Control key. For example, Ctrl/P indicates you should press the P key while holding down the Control key.
MultiSwitch 900	This term refers to the DIGITAL MultiSwitch 900, formerly the DEChub 900 MultiSwitch.

Associated Documents

The following documents provide information relating to the VNswitch. For online copies of these documents, refer to the [Online Services](#) section of this manual (or browse your CD). All documents beginning with a “EK-” order number can be obtained in printed form directly from DIGITAL, as described in the section [How to Order Additional Documentation](#).

Title and Order Number	Description
<i>DIGITAL VNswitch 900 Series Technical Overview</i> AA-R2LC*-TE	Provides a technical overview of the VNswitch 900 family of high-density switching products. It also includes product specifications and ordering information.
<i>DIGITAL VNswitch 900 Series Router Management</i> AA-R87C*-TE	Describes how to use the VNswitch routing software to configure, monitor, and manage VNswitch routing functions.
<i>DIGITAL VNswitch 900EE Installation and Configuration</i> EK-DVNEE-IN	Describes the VNswitch 900EE module, including features, installation, and configuration information.
<i>DIGITAL VNswitch 900EX Installation and Configuration</i> EK-DVNEX-IN	Describes the VNswitch 900EX module, including features, installation, and configuration information.
<i>DIGITAL VNswitch 900EF Installation and Configuration</i> EK-DVNEF-IN	Describes the VNswitch 900EF module, including features, installation, and configuration information.
<i>DIGITAL VNswitch 900EA Installation and Configuration</i> EK-DVNEA-IN	Describes the VNswitch 900EA module, including features, installation, and configuration information.
<i>DIGITAL VNswitch 900LL Installation and Configuration</i> EK-DVNLL-IN	Describes the VNswitch 900LL module, including features, installation, and configuration information.
<i>DIGITAL VNswitch 900XX Installation and Configuration</i> EK-DVNXX-IN	Describes the VNswitch 900XX module, including features, installation, and configuration information.

Associated Documents

Title and Order Number	Description
<i>DIGITAL VNswitch 900XA Installation and Configuration</i> EK-DVNXA-IN	Describes the VNswitch 900XA module, including features, installation, and configuration information.
<i>DIGITAL VNswitch 900FA Installation and Configuration</i> EK-DVNFA-IN	Describes the VNswitch 900FA module, including features, installation, and configuration information.
<i>DIGITAL ATM Modular PHY Cards Installation</i> EK-DAGGM-IN	Provides installation and operating guidelines for installing, verifying, and removing modular PHY cards. Describes cabling and LED information.
<i>DIGITAL MultiSwitch 900 Owner's Manual</i> EK-DH2MS-OM	Provides installation, use, security, and troubleshooting information for the DIGITAL MultiSwitch 900.
<i>DEChub ONE Installation</i> EK-DEHU2-IN	Provides installation and operation guidelines for standalone module configuration, including mounting options and cabling.
<i>DEChub ONE-MX Installation</i> EK-DEF1H-IN	Provides installation and operation guidelines for standalone module configuration, including mounting options and cabling.
<i>clearVISN Installation</i> AA-QX86*-TK	Provides pre- and post-installation information, as well as actual installation procedures for each application.
<i>clearVISN Overview</i> AA-QX87*-TK	Provides an overview of clearVISN, an explanation of each application, and descriptions of all concepts necessary to understand and use the application efficiently.
<i>clearVISN RMON Manager User's Guide</i> AA-R2RM*-TH	Provides an overview of the RMON Manager, procedures for monitoring and viewing network segment traffic, and information on logging and reporting the statistics collected.
<i>clearVISN User's Guide</i> AA-QX88*-TK	Provides information for starting each application, configuring them, and general use information.

Associated Documents

Title and Order Number	Description
<i>OPEN DECconnect Applications Guide</i> EC-G6387-42	Provides information to help plan and install networking systems based on DIGITAL OPEN DECconnect System and networking products.
<i>Event Logging System Messages Guide</i> AA-QL2A*-TE	Describes messages logged by the Event Logging System.
<i>Bridge and Extended LAN Reference</i> EK-DEBAM-HR	Describes how bridges are used to create extended local area networks (LANs). The descriptions include the use of bridges in extended LAN configurations, information on LAN interconnections, overall bridge operation, spanning tree, bridge management, and solving bridge-related problems in a network.
<i>A Primer on FDDI: Fiber Distributed Data Interface</i>	Provides general introductory information about the features, topologies, and components of the FDDI local area network standard.
<i>DEChub Network Modules 900-Series Concentrator Reference</i> EK-CONTR-HR	Provides detailed reference information about the various DIGITAL 900-series concentrators supported by the DIGITAL MultiSwitch 900.

Correspondence

Documentation Comments

If you have comments or suggestions about this manual, send them to DIGITAL Network Products.

Attn.: Documentation Project Manager

E-MAIL: doc_quality@lkg.mts.dec.com

Online Services

To locate product-specific information, refer to the DIGITAL Network Products Home Page on the World Wide Web located at the following addresses:

Americas <http://www.networks.digital.com>

Europe <http://www.networks.europe.digital.com>

Asia Pacific <http://www.networks.digital.com.au>

How to Order Additional Documentation

To order additional printed documentation, use the following information:

To Order:	Contact:
By Telephone	USA (except Alaska, New Hampshire, and Hawaii): 1-800-DIGITAL (1-800-344-4825) Alaska, New Hampshire, and Hawaii: 1-603-884-6660 Canada: 1-800-267-6215
Electronically (USA only)	Dial 1-800-DEC-DEMO (For assistance, call 1-800-DIGITAL)
By Mail (USA and Puerto Rico)	DIGITAL EQUIPMENT CORPORATION P.O. Box CS2008 Nashua, New Hampshire 03061 (Place prepaid orders from Puerto Rico with the local DIGITAL subsidiary: 809-754-7575)
By Mail (Canada)	DIGITAL EQUIPMENT of CANADA LTD. 940 Belfast Road Ottawa, Ontario, Canada K1G 4C2 Attn.: A&SG Business Manager
Internationally	DIGITAL EQUIPMENT CORPORATION Attn.: A&SG Business Manager c/o local DIGITAL subsidiary or approved distributor
Internally	U.S. Software Supply Business (SSB) DIGITAL EQUIPMENT CORPORATION 8 Cotton Road Nashua, New Hampshire 03063

Chapter 1

Introduction

Overview

Introduction

This chapter gives an overview of your responsibilities as a VNswitch administrator. This chapter also provides an introduction to concepts and terminology with which you should be familiar.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Concepts and Terminology	1-2
Overview of Your Responsibilities	1-16

Concepts and Terminology

This section presents basic VNswitch module concepts and terminology with which you should become familiar. Refer to the *DIGITAL VNswitch 900 Series Technical Overview* for a more detailed discussion about specific concepts and terms.

Bridge Plug and Play

The VNswitch family of modules features plug-and-play operation. Once installed, a VNswitch module begins handling network traffic and running transparent bridging on each port, by default. You may need to alter default settings to maximize network performance, to enable special functions such as Virtual Local Area Networks (VLANs), and to set IP addresses for remote management.

The default settings used at installation are listed in [Appendix B](#). The default setting for each configurable parameter is also given when that parameter is discussed in the following chapters.

Switch Console Sessions

You can configure, monitor, and manage a VNswitch module by establishing either a local or remote console session with the module or by using the clearVISN network management product from Digital Equipment Corporation. With this release, you can also use the VNswitch web-based management application to manage the VNswitch.

Local Sessions

A local session is established by connecting a terminal (a workstation or PC running terminal emulation, for example) directly to the setup port of a DIGITAL MultiSwitch 900 or a DEChub ONE docking station and selecting Go to Local Console from the setup menu. Local sessions are used to configure and manage those modules on the hub to which the workstation or PC is attached.

Remote Sessions

A remote session is established by running Telnet on a remote system and entering the VNswitch module's IP Host Services IP address or the IP address of the switch's Out of Band Management (OBM) port. The Telnet client program can be run on a workstation, a PC, or a terminal server, for example. A maximum of two remote sessions can be established with a switch at the same time.

Web-based Management Application

This version of the VNswitch provides a web-based management application that allows you to perform many of the tasks needed to configure, monitor, and manage VNswitch modules over the Internet. Currently, not all tasks necessary to manage VNswitch modules can be performed using the web-based management application. You will still need to use the CLI to fully manage a module. This manual focuses on the CLI management interface. Documentation for the web-based management application can be found in the online help for the application.

Transparent Bridging

VNswitch modules support transparent bridging. Transparent bridging is the ability of the switch to automatically “learn” the network addresses and locations of other network devices. The switch organizes the addresses in bridge tables in such a way that it is able to determine on which switch port the device is located. Transparent bridges are also referred to as *learning* or *adaptive bridges*.

Understanding Network Interfaces and Ports

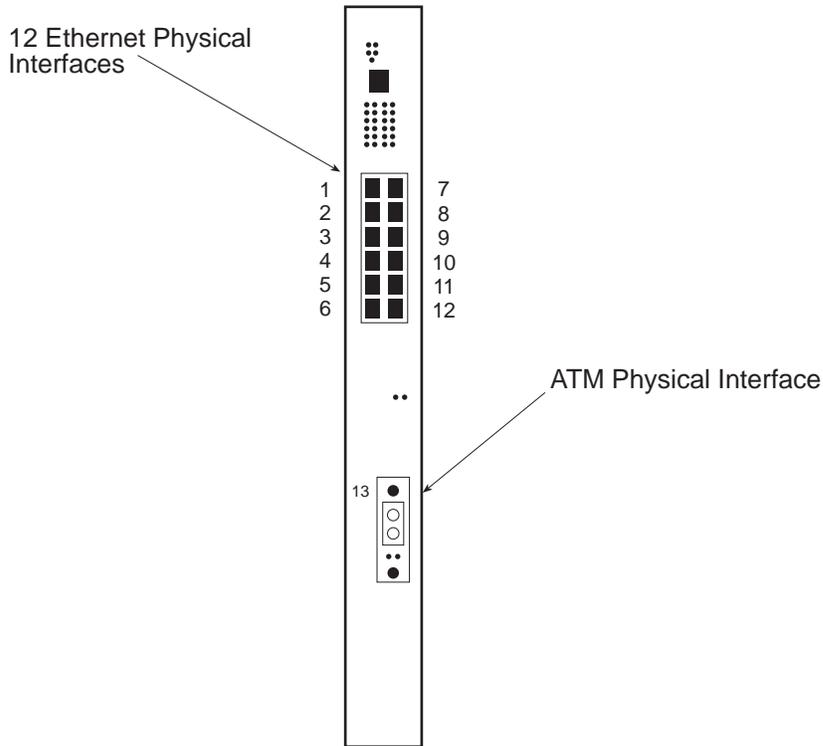
VNswitch module architectural design applies different definitions to the terms *interface* and *port*. The design further distinguishes among two types of interface: physical and logical.

Physical Interface

A physical interface is the physical point on the module to which a network transmission medium (cable or fiber, for example) is connected. Physical interfaces on VNswitch modules include Ethernet, Fast Ethernet, FDDI, ATM, and VNbus. Ethernet, Fast Ethernet, FDDI, and ATM physical interfaces are identified by a unique number next to each interface on the module’s front panel. The VNbus physical interface is located on the module’s back panel. The VNbus interface number (0) is next to the VNbus Light Emitting Diodes (LEDs) at the top of the front panel. (Refer to your module’s installation and configuration documentation to determine the exact location of these LEDs.) [Figure 1-1](#) shows the physical interfaces on the front panel of a VNswitch 900EA module.

Concepts and Terminology

Figure 1-1: Physical Interfaces



LKG-10261-96F

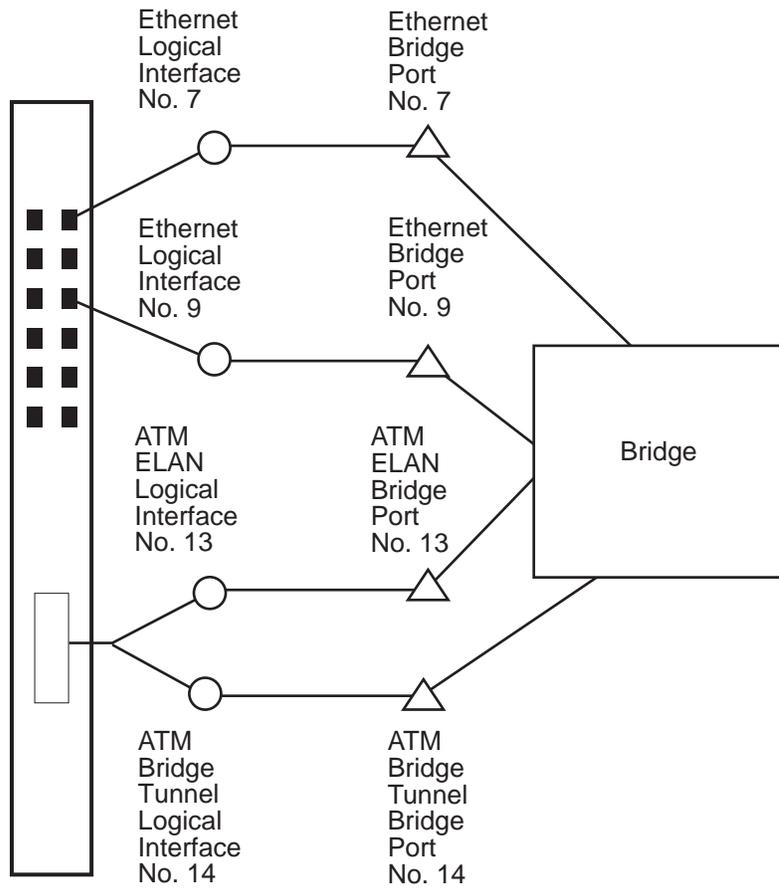
Logical Interface

A logical interface is an abstract connection point, within the system's software, between a physical interface and a bridge port. Ethernet, Fast Ethernet, and FDDI physical interfaces are each associated with one logical interface. An ATM physical interface is associated with 1 to 16 logical interfaces, each of which is the connection point to either an ATM emulated LAN (ELAN), or an ATM bridge tunnel. Each logical interface on a switch is identified by a unique number. [Figure 1-2](#) shows examples of logical interfaces on a VNSwitch 900EA module.

Bridge Port

A bridge port is an abstract connection point, within the system's software, to a transparent bridge. The transparent bridge forwards data to, or receives data from, bridge ports, based on the MAC address associated with the data. Each bridge port on a switch is identified by a unique number. [Figure 1-2](#) shows a VNSwitch 900EA module that includes two ports from two Ethernet interfaces, one port from an ELAN interface, and one port from an ATM bridge tunnel interface.

Figure 1-2: Logical Interfaces and Bridge Ports



LKG-10262-96F

Interface and Bridge Port Numbering Scheme

Each physical and logical interface is assigned an interface number. For Ethernet, Fast Ethernet, and FDDI interfaces, the physical interface number printed on the module's front panel is the same as the logical interface number. For modules that provide an ATM interface, for example a DIGITAL VNswitch 900EA, the physical interface number (13) is the same as only one of the sixteen ATM logical interface numbers (13 through 28). The VNbus interface, located on the module's back panel, is assigned a physical and logical interface number of zero (0). The VNbus interface number (0) is printed next to the VNbus LEDs at the top of the front panel.

A bridge port number has the same value as that of its associated logical interface. For example, FDDI logical interface number 13 is associated with FDDI bridge port number 13. Similarly, ATM logical interface number 15 has a bridge port number of 15, and so on. Refer to [Figure 1-2](#) for an example of interface and port numbering schemes.

Concepts and Terminology

VLANs and VLAN Secure Domains

A VLAN is a group of bridge ports logically linked to define a LAN on one or more hubs. This network configuration scheme enables you to configure a set of devices so they logically appear to be on the same LAN segment, although they may be physically on different segments.

A VLAN Secure Domain (VSD) is a logical set of one or more VLANs that operate with one spanning tree. A VLAN consists of a set of distinct bridge ports. Each set of bridge ports is isolated from other ports on the same switch by blocking all unicast and multicast traffic between VSDs. VNswitch modules presently support one VLAN per VSD, but the VSD concept provides for expanded support of multiple VLANs within a single VSD.

In most respects, each VSD operates as a separate logical bridge within the module. For example, a separate instance of the spanning tree protocol is run on each VSD.

Note

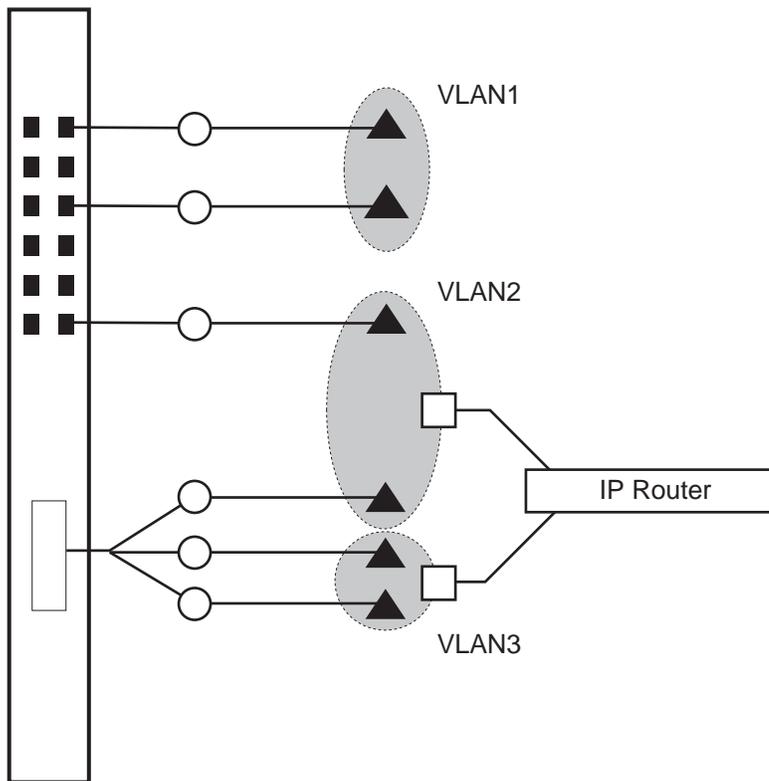
A VLAN is currently equivalent to a VSD because the present implementation supports one VLAN per VSD.

Refer to [Chapter 9](#) for more information about VLANs and VSDs.

VLAN Logical Interface

A VLAN logical interface is an abstract connection between a VLAN and a router, enabling you to connect multiple VLANs through the router. [Figure 1-3](#) shows examples of VLAN logical interfaces on a VNSwitch 900EA module.

Figure 1-3: VLAN Logical Interfaces



- ▲ = Bridge Port
- = Logical Interface
- = VLAN Logical Interface

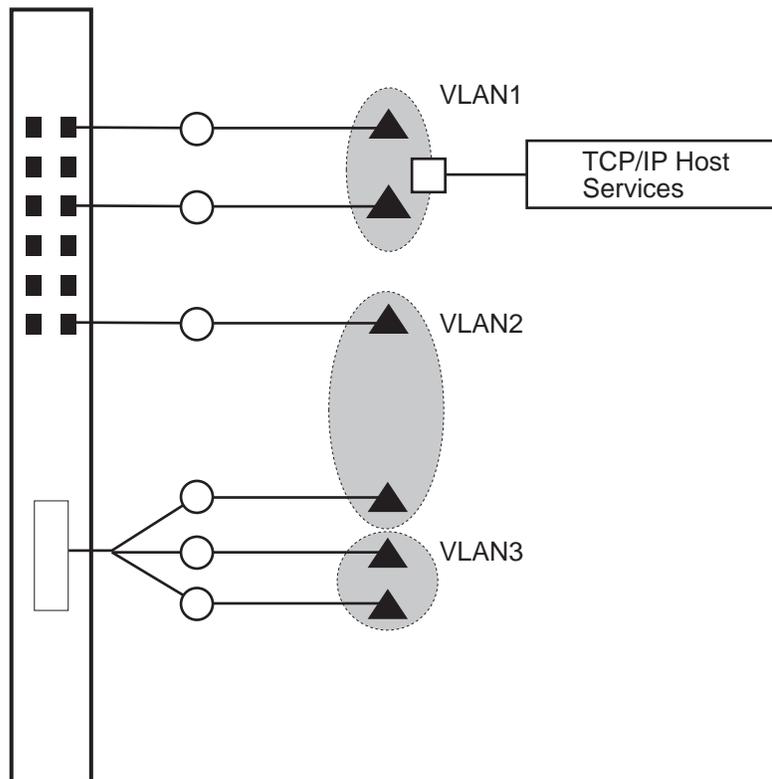
LKG-10263-96F

Concepts and Terminology

TCP/IP Host Services Logical Interface

A TCP/IP Host Services logical interface is an abstract connection between a VLAN and TCP/IP Host Services. [Figure 1-4](#) shows examples of VLAN logical interfaces on a VNswitch 900EA module.

Figure 1-4: TCP/IP Host Services Logical Interface



▲ = Bridge Port

○ = Logical Interface

□ = TCP/IP Host Services Logical Interface (BDG/0)

LKG-10368-97F

User Interface

You can configure, monitor, and manage VNswitch modules using a command line interface (CLI). You can access the CLI through either a local or remote console session.

You can also configure, monitor, and manage VNswitch modules through one of the following:

- A graphical interface if you install a MIB-based management system such as MultiChassis Manager (optional), a component of the clearVISN network management product from Digital Equipment Corporation. Refer to the documentation accompanying your MIB-based management system for additional information.
- The VNswitch web-based management application

With Version 2.0 of the VNswitch 900 series switch, certain commands from previous versions have become obsolete. These commands are still part of the software and are still functional; however, they do not appear in the help display. DIGITAL recommends that you avoid using these commands as they may not be supported in future releases.

The remaining chapters in this manual describe what commands you enter at the CLI to perform specific tasks. The initial steps for most of those tasks involve accessing a CLI prompt.

The VNswitch CLI consists of three components:

- Main operator's process
- Configuration
- Monitor

To perform a specific task, you access the prompt associated with the task you want to perform. The following describes the VNswitch components:

Component	Prompt	Description
Main operator's process	Main>	This component performs certain backup, upgrade, and restoration tasks; manages console sessions; and restarts the switch. You access all lower level prompts from the Main prompt.

Concepts and Terminology

Component	Prompt	Description
Configuration	Config>	<p>This component configures parameters that apply to the entire switch, rather than to a specific interface or port. These parameters include, for example, adding users who can manage the switch, and enabling or disabling prompting for ID and password.</p> <p>You use the Config prompt to access lower level prompts that enable you to configure the switch's interfaces, protocols, ports, and parameters pertaining to boot and dump configuration, event logging, error logging, VLANs, remote monitoring, and the mirror port.</p>
Monitor	Monitor>	<p>This component monitors parameters that apply to the entire switch, rather than to a specific interface or port. These include, for example, displaying a list of users who can manage the switch, and determining whether prompting for ID and password is enabled.</p> <p>You use the Monitor prompt to access lower level prompts that enable you to monitor parameters associated with switch interfaces, protocols, and ports.</p>

Instructions about how to access prompts are presented in [Chapter 2](#). Instructions about how to access lower level prompts within each component are presented in the relevant chapters in this manual. Refer to [Chapter 11](#) for information about how to display the event log.

Auto-Prompts

The VNSwitch CLI enables you to perform certain tasks either by entering commands and parameter variables as a string, or by responding to instructional prompts that walk you through each item you must enter. Instructions about how to use auto-prompts are presented in [Chapter 2](#).

Types of Switch Memory

VNswitch modules store data in the following two types of random access memory (RAM):

Type of RAM	Description
Volatile RAM	<p>Volatile RAM is used to store parameters you enter from the Monitor prompt. It is also used to store various buffers, and information such as network addresses learned by the switch, and static address entries entered by you or another network manager. (Refer to the Address Types section for information about the types of addresses stored in Volatile RAM.)</p> <p>Information stored in volatile RAM is lost when the switch is powered down or restarted.</p>
Nonvolatile RAM (NVRAM)	<p>NVRAM is used to store parameters you configure from the Config prompt. It is also used to store the switch's executable software (also known as the boot or image file). User-configured parameters may include network addresses, referred to as permanent address entries. (Refer to the Address Types section for information about the types of addresses stored in NVRAM.)</p> <p>The configuration database and image remain intact when the switch is powered down or power is lost. Refer to Chapter 10 for information about backup and restore.</p>

Configuration Commands Requiring Restart

Some of the commands you use to specify or change configuration settings on the VNswitch module require that you restart the module for the changes to take effect. The procedures described in this manual indicate when a restart is required. If the procedure does not indicate a restart is required, the configuration parameter you specify takes affect immediately. This is sometimes referred to as dynamic configuration.

All parameters configured from the Config prompt survive restarts and loss of power. The configuration parameters you specify from the Monitor prompt, or lower level prompts under the Monitor prompt, are dynamic. These settings are stored in volatile memory and, therefore, are lost when there is a power outage or the module is restarted.

Concepts and Terminology

Address Types

The MAC address of network devices can be added to bridge tables automatically through bridge learning, manually by you or another network manager, or by one or more network devices. [Table 1-1](#) describes the different address types and where they are stored.

Table 1-1: MAC Address Type Descriptions

Address Type	Description
Dynamic address	A MAC address automatically learned by the VNswitch module. Dynamic address entries are stored in volatile RAM and, therefore, do not survive power cycles or system resets. Dynamic entries are affected by address aging.
Permanent address	A MAC address entered manually by a network manager from the Config prompt. Permanent address entries survive power cycles or system resets. Permanent entries are not affected by address aging.
Static address	A MAC address entered manually by a network manager from the Monitor prompt. Static address entries do not survive power cycles or system resets. Static entries are not affected by address aging.
Reserved address	A MAC address reserved by the IEEE 802.1d standard.
Registered address	A unicast MAC address that belongs to communications hardware attached to the module, or a multicast address enabled by protocol forwarders.
Unknown address	A MAC address that does not fall into any of the other address type categories listed in this table.

System Security

The switch software can, optionally, require users to enter a user name and password when logging in at a switch console. (Refer to [Chapter 3](#) for information about how to enable and disable ID and password prompts.) It further distinguishes between the following three types of users, each of which is associated with a different level of access privilege to configuration, monitoring, and management functions:

Administrative users — Can access any configuration, monitoring, or management function, including adding and managing users. Only a user with Administrative access can change configuration in NVRAM.

Operations users — Can view any network configuration parameter or statistic, run potentially disruptive tests, dynamically change switch operation by reconfiguring parameters via volatile RAM, and restart the switch.

Monitor users — Can only view configuration parameters and network statistics.

Overview of Your Responsibilities

A VNswitch 900 manager is responsible for the following general activities:

Preinstallation planning — As a VNswitch module user, you should be involved in network design or expansion planning. If you are not closely involved in such planning activities, you should receive detailed site preparation instructions as well as instructions about how the various network devices are to be configured. Refer to the *DIGITAL VNswitch 900 Series Technical Overview* for background information about switching and bridging concepts, and examples of network configurations that are based on those concepts. The information in the Technical Overview may be helpful during network design and planning. If you plan to use routing, refer to the *DIGITAL VNswitch 900 Series Router Management* manual.

Installing VNswitch modules — Hardware installation documentation is shipped with each VNswitch 900 module. Refer to the appropriate installation document for instructions about how to install each module and run diagnostics. The installation documentation also includes information about hardware requirements such as supported cabling and connectors.

Connecting the switch console — Local switch consoles must be connected to the switching hub after the switch modules are installed. Refer to the appropriate installation document for instructions about connecting consoles.

Adding and managing users — You can allow other users to access certain configuration, monitoring, and management functions by adding them to a user list. You can restrict the types of activities the user can perform by assigning them to one of three categories of user. You can also change user passwords, disable the ability to log in from a remote console, and delete users. Refer to [Chapter 3](#) for instructions about adding and managing users.

Overview of Your Responsibilities

Configuring a switch — VNSwitch modules feature bridge plug-and-play configuration. Once installed, a module automatically configures itself with default settings, and begins switching and bridging traffic over the network.

You may typically need to alter default settings to maximize network performance, to enable special functions such as Virtual Local Area Networks (VLANs), and to configure remote management. Refer to [Chapter 4](#), [Chapter 5](#), [Chapter 7](#), [Chapter 9](#), [Chapter 12](#), and [Chapter 13](#) for instructions about the following topics, respectively:

- Configuring switch-wide parameters
- Configuring interfaces and the Address Resolution Protocol (ARP)
- Configuring transparent bridging
- Configuring Virtual LANs
- Configuring remote management, the Simple Network Management Protocol (SNMP), and TCP/IP Host Services
- Configuring the mirror interface and the RMON agent

Monitoring and managing a switch — You are responsible for ongoing monitoring and maintenance of VNSwitch modules once they are operational. This responsibility includes such activities as monitoring traffic collision statistics and error counters, and backing up switch configurations. Refer to [Chapter 4](#), [Chapter 6](#), [Chapter 8](#), [Chapter 10](#), [Chapter 12](#) and [Chapter 13](#) for instructions about the following topics, respectively:

- Monitoring switch-wide parameters
- Monitoring interfaces and the Address Resolution Protocol (ARP)
- Monitoring and managing bridging
- Performing routine maintenance procedures such as backup and restoration, detecting crashes that may have occurred (followed by automatic recovery), and upgrading VNSwitch software
- Monitoring remote management, the SNMP, and TCP/IP Host Services
- Monitoring the mirror interface and the RMON agent

Chapter 2

Operational Basics

Overview

Introduction

This chapter describes the basic procedures for working with a DIGITAL VNswitch module. They are the operational basics that are common to many of the configuration and management tasks described in the following chapters.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Starting and Terminating Console Sessions	2-2
Accessing CLI Prompts	2-7
Exiting a Prompt	2-10
Using the Command Line Interface	2-11
Entering Commands and Command Shortcuts	2-14
Using Auto-Prompts	2-16
Using the Web-Based Management Application	2-18
Displaying CLI Help	2-21

These topics are common to most of the procedures described throughout this manual and are frequently referenced by those procedures. For example, you must start a console session and access a CLI prompt before you begin to configure a network interface. Similarly, once logged in, you may want to display Help information about command options you can use.

Starting and Terminating Console Sessions

You can configure, monitor, and manage a VNswitch module by establishing either a local or remote console session with a module.

Local sessions are used to configure and manage only those modules on the hub to which the terminal is attached. A local session is established by connecting a terminal directly to the setup port of a DIGITAL MultiSwitch 900 (formerly known as the DEChub 900) or a DEChub ONE docking station and selecting Go to Local Console from the setup menu. Remote sessions are used to configure and manage any module on the network. Remote sessions are established by running Telnet (either from the command line or from the web-based management application) on a remote system and connecting through any of the following interfaces:

- The OBM port on the hub
- The AUI port (DEChub ONE docking station only)
- Any VNswitch module network interface (Ethernet, Fast Ethernet, FDDI, ATM ELAN, ATM Bridge Tunnel, or the VNbus)

The setup port and the OBM port provide out-of-band management. Access through the AUI port or through a network interface provides in-band management. In-band management cannot be used when the network is down, and can use up a portion of the network's bandwidth. Out-of-band management remains operational when the network is down and does not affect bandwidth. In-band management traffic is subject to bridge and VLAN filters. Out-of-band management is not subject to bridge and VLAN filters. Refer to [Chapter 9](#) for additional information about VLANs.

The following instructions about starting and terminating sessions assume your terminal is already physically connected and configured for access to VNswitch modules. Refer to the VNswitch module installation and configuration documentation for information about installing devices to connect to a setup, OBM, or AUI port. Refer to [Chapter 12](#) of this document for information about how to configure the VNswitch module for Telnet, OBM, and AUI connections.

Starting and Terminating Local Sessions

To start a local console session through a setup port, perform the following steps:

Step	Action
1	Turn on the power to your terminal. If your terminal is attached to a DEChub ONE, a menu similar to the one shown in Figure 2-1 is displayed. If your terminal is attached to a DIGITAL MultiSwitch 900, the menu shown in Figure 2-2 is displayed.
2	If your terminal is connected to a DEChub ONE, go to step 4. If your terminal is connected to a DIGITAL MultiSwitch 900, enter 9 (Start Redirect Mode), and press Return. You are prompted for the number of the slot that contains the switch module with which you want to establish a session.
3	Enter the number of the slot that contains the module with which you want to establish a session, then press Return. A menu similar to the one shown in Figure 2-1 is displayed.
4	Enter 6 (Go to Local Console) and press Return. A Local Console session is established and the Main prompt (Main>) is displayed. Note: If ID and password prompting is enabled by the switch administrator, you are prompted for your ID and password before the Main prompt (Main>) is displayed.

To terminate a console session, enter **Ctrl/C** or **logout** at the Main prompt (Main>) and press Return. If you are connected to a DIGITAL MultiSwitch 900, a message is displayed indicating that your terminal is disconnected from the module. Press Return to redisplay the main installation menu ([Figure 2-2](#)). If you are connected to a DEChub ONE, the installation menu shown in [Figure 2-1](#) is displayed.

Starting and Terminating Console Sessions

Figure 2-1: VNswitch 900EA Installation Menu

```
VNswitch 900EA
=====
                VNswitch 900EA INSTALLATION MENU

                [1] Restart with Factory Defaults
                [2] Restart with Current Settings
                [3] Show Current Settings
                [4] Configure IP ...
                [5] Configure Out-of-Band Port ...
                [6] Go to Local Console
                [7] Product-Specific Options ...

=====

                Enter selection:
```

Figure 2-2: DIGITAL MultiSwitch 900 Installation Menu

```
DIGITAL MultiSwitch 900
=====
                MultiSwitch 900 INSTALLATION MENU

                [1] Reset with Factory Defaults
                [2] Reset with Current Settings
                [3] Show Current Settings
                [4] Configure IP
                [5] Dump Error Log
                [6] Downline Upgrade
                [7] Configure Out-of-Band Port
                [8] Start Event Display Mode
                [9] Start Redirect Mode
                [10] Product-Specific Options

=====

                Enter selection number:
```

Starting and Terminating Remote Sessions

Remote console sessions can be established only after configuring the appropriate network connections. Refer to [Chapter 12](#) for information about how to do so. A maximum of two remote sessions can be established with a switch at the same time.

To start a session from a remote device (workstation, PC, or terminal server, for example) to the module through an OBM port, AUI port, or a module's network interface, perform the following steps:

Step	Action
1	<p>Access the operating system prompt from your terminal.</p> <p>Note: If you are running a local session through the setup port of a hub, you can also access the Main prompt and run Telnet on the module. This enables you to establish a remote session from one VNswitch module to another.</p>
2	<p>Enter telnet ip-address, where <i>ip-address</i> is one of the following addresses.</p> <ul style="list-style-type: none"> • If you are using the OBM port, enter the OBM-IP address for the module you want to access. • If you are using the AUI port, enter the IP address for the AUI port on the DEChub ONE you want to access. • If you Telnet through a network server, enter the IP address for the VNswitch you want to access. <p>Refer to Chapter 12 for information about how to determine the IP address of an OBM port and the module's IP address.</p>
3	<p>Press Return. The Main prompt (Main>) of the remote VNswitch module is displayed.</p> <p>Note: If ID and password prompting is enabled by the switch administrator, you are prompted for your ID and password before the Main Process prompt is displayed.</p>

To terminate a console session, enter **Ctrl/C** or **logout** at the Main Process prompt (Main>) and press Return. The network operating system prompt is displayed.

Obtaining an IP Address Automatically

If you have a BootP or DHCP server running on your network, and your VNswitch is not assigned an IP address, the VNswitch takes advantage of BootP client software to automatically obtain an IP address for itself during power-up or restart. Refer to the vendor's BootP or DHCP documentation for configuration information.

Starting and Terminating Console Sessions

A VNswitch (as a BootP client) that does not have an IP address assigned, sends out a BootP (broadcast) request to a BootP or DHCP server. When the server replies with an IP address, the VNswitch configures the IP address for HST dynamically. This IP address is stored permanently, so power-cycling the switch does not have any impact on the IP address. To change the IP address, you use the configuration menu.

An IP address is required for the VNswitch if you plan to manage it using an SNMP tool such as clearVISN.

Accessing CLI Prompts

The initial steps for most of the tasks discussed throughout this manual involve accessing the CLI prompts (Main, Config, and Monitor). Instructions about how to access the prompts are presented here, rather than repeating them for each task covered later in this manual.

Only one user at a time can access the Main, Config, and Monitor prompts or the error log. If another user attempts to access the same prompt you are currently using, the message `Current Process has been Redirected` is displayed and you are returned to the Main prompt (Main>). If, for example, you access the Monitor prompt from the Main prompt and another user then accesses the Monitor prompt, the message `Current Process has been Redirected` is displayed and you are returned to the Main prompt. The user who accessed the prompt you were using receives all redirected output from those tasks you initiated, but that did not yet display on your screen. Any task you initiated is completed unless the user to whom the output is redirected cancels it.

Accessing the Main Prompt

The Main prompt (Main>) is automatically displayed each time you start a console session. (Refer to the [Starting and Terminating Console Sessions](#) session.)

Commands available from the Main> prompt are:

```
Main>?  
CONFIG  
DIVERT output from process  
DUMP contents of memory  
EVENTS  
FLUSH output from process  
HALT output from process  
INTERCEPT character is  
LOGOUT  
MEMORY statistics  
MONITOR  
RESTART  
RELOAD  
STATUS of process(es)  
TELNET to IP-Address <this terminal type>
```

Accessing CLI Prompts

Accessing the Config Prompt

To access the Config prompt, perform the following steps:

Step	Action
1	At the Main prompt (Main>), enter C onfig.
2	Press Return. The Config prompt (Config>) is displayed.
3	If the prompt is not displayed, press Return a second time.

Commands available from the Config> prompt are:

```
Config>?  
ADD  
BOOT subsystem  
CHANGE  
CLEAR  
DELETE  
DISABLE  
ELS subsystem  
ENABLE  
ERR-LOGS subsystem  
INTERFACE subsystem  
LIST  
MONITOR switch to Monitor process  
SET  
TIME of day params  
VLANS subsystem  
EXIT  
Protocol subsystems:  
    IP, ARP, SNMP, OSPF, BRIDGE, HST  
Feature subsystems:  
    RMON, Mirror
```

Accessing the Monitor Prompt

To access the Monitor prompt, perform the following steps:

Step	Action
1	At the Main prompt (Main>), enter M onitor.

Step	Action
2	Press Return. The Monitor prompt (Monitor>) is displayed.
3	If the prompt is not displayed, press Return a second time.

Commands available from the Monitor> prompt are:

```
Monitor>?
BUFFER statistics
CLEAR statistics
CONFIG switch to Config process
ELS subsystem
ERROR counts
ERR-LOGS subsystem
INTERFACE subsystem, commands and statistics
LIST
MEMORY statistics
QUEUE lengths
STATISTICS of network
UPTIME of gateway
VLANS subsystem
EXIT
Protocol subsystems:
    IP, ARP, SNMP, OSPF, BRIDGE, HST
Feature subsystems:
    RMON, Mirror
```

Adding the Hostname to a Prompt

You can prefix prompts with the switch's hostname. For example, if the hostname is defined as `cmtsrv.lkg.dec.com`, the new prompt might look like this:

```
cmtsrv Monitor>
```

To enable the hostname prefix, enter **enable prompting-with-hostname** at the Config> prompt. To disable the hostname prefix, enter **disable prompting-with-hostname** at the Config> prompt.

Exiting a Prompt

The Config prompt (`Config>`) and the Monitor prompt (`Monitor>`) are accessed from the Main prompt (`Main>`). Tasks you may want to perform using these prompts typically require that you access lower level prompts before you can enter the appropriate commands.

To exit any lower level prompt and return to the next higher level prompt, enter **exit** and press Return. To return to the Config or Monitor prompt from any lower level prompt, enter **exit** and press Return repeatedly until the Config or Monitor prompt is displayed.

To return to the Main prompt from any lower level prompt, press Ctrl/P (the default intercept character). If you use Ctrl/P to skip two or more lower-level prompts and jump to the Main prompt directly, the next time you access either the Config or Monitor prompt, you are returned directly to that lower level prompt from which you previously exited. This can be convenient if you are performing a task at some lower level prompt under `Config>`, for example, and to move back and forth from that prompt level to the Monitor prompt.

The Ctrl/P key combination is called the intercept character. You can change the intercept character, if necessary. (Refer to the [Changing the Intercept Character](#) section for instructions.)

Changing the Intercept Character

The intercept character is used to return to the Main prompt from another prompt. The default intercept character is Ctrl/P. To change the intercept character, perform the following steps:

Step	Action
1	At the Main prompt, enter intercept .
2	Press Return. The following message is displayed: <code>Enter character []:</code>
3	Enter the desired character (x , for example).
4	Press Return. The intercept character is changed.

Using the Command Line Interface

The command line interface provides features that make entering commands to the VNSwitch module quicker and easier. You can:

- Edit commands on the command line
- Recall commands previously entered
- Complete partially entered commands automatically

For information on using command shortcuts, see *Entering Commands and Command Shortcuts* on page 2-14.

Using Command Line Editing

Command line editing allows you to correct or change your entries on the command line. Table 2-1 lists the command line editing keys.

Table 2-1: Command Line Editing Keys

To...	Enter...
Move left one character	Ctrl/b or left arrow
Move right one character	Ctrl/f or right arrow
Restore the line as it was before editing	Ctrl/r
Delete the character to the left	Ctrl/h, Delete or Backspace
Delete the character at the cursor	Ctrl/d
Move to the beginning of the line	Ctrl/a
Move to the end of the line	Ctrl/e
Delete to the end of the line	Ctrl/k
Transposes the characters at the cursor and the character to the left	Ctrl/t

Using the Command Line Interface

Using Command Line Recall

Command line recall stores up to 10 previously entered commands in one session. You can redisplay those commands on the command line, one at a time, to re-enter them or edit and then re-enter them. Table 2-2 lists the command line recall keys.

Table 2-2: Command Line Recall Keys

To...	Enter...
Display the command that you entered after the currently displayed command	Ctrl/u or up arrow
Display the command that you entered before the currently displayed command	Ctrl/n or down arrow

Using Command Line Completion

With command line completion (CLC), you can enter part of a command then press the space bar for automatic completion of the command. Depending on the ambiguity of your entry, CLC completes as much of the command as possible, or displays a list of options.

Examples: The following examples show how CLC works. The underscore () in these examples represents pressing the space bar. The vertical bar (|) represents the cursor position after command line completion.

Command Line Completion Conditions

Command Line Entry	Resulting Command Line Completion
MAIN>_	There are 14 available options... CONFIG DIVERT . TELNET
	MAIN>
IP Config>l_	IP Config>list
IP Config>add ac_	IP Config>add acce
IP Config>c_	There are 2 available options... CHANGE CONFIG
	IP Config>c

Using the Command Line Interface

Command Line Entry	Resulting Command Line Completion
IP Config> add a_	There are 3 available options... ACCEPT-RIP-ROUTE ACCESS-CONTROL ADDRESS IP Config>add a
IP Config> add _	There are 7 available options... ACCEPT-RIP-ROUTE ACCESS-CONTROL ADDRESS BOOTP-SERVER BROADCAST-FORWARDER ENHANCED-PROXY-ARP FILTER ROUTE IP Config>add

Entering Commands and Command Shortcuts

You perform tasks by entering commands at a CLI prompt. For example, if you want to view a list of all users, you access the Config prompt (`Config>`) and enter **list users**. Similarly, to see how long the switch has been running since the last reboot, you access the Monitor prompt (`Monitor>`) and enter **uptime**.

Most tasks can also be initiated by entering only part of a command as a shortcut, rather than entering the whole command. In the following chapters, that portion of the command that can be entered as a shortcut is indicated with an underscore. For example, the commands for listing users is shown as **list uers**, and can be entered at the Config prompt as **l u**. Similarly, the command for viewing uptime is shown as **uptime** and can be entered at the Monitor prompt as **u**.

Be aware that some commands are quite similar, so more of the command needs to be entered to indicate uniqueness. For example, the command for viewing error statistics is shown as **error** and can be entered at the Monitor prompt as **er**. However, the command to view the error logs is shown as **err-logs**, and can be entered as **err-**.

Entering Subsystem Commands

Although the CLI is tree-structured, you can bypass that structure when you are familiar with the commands for various subsystems. For example, if you are at the Monitor (Monitor>) prompt, you can check bridge information without going first to the Bridge console prompt (Bridge>). At the Monitor prompt you enter:

```
Monitor>monitor bridge list all
```

The Monitor prompt remains, but the data displayed is from the Bridge console subsystem. This shortcut allows you to execute a single command for a subsystem without leaving the Monitor or Config prompt.

You can also use this shortcut to enter commands for other VNswitch subsystems without leaving the current subsystem. For example, while at the Monitor IP> prompt, you can find out about bridge counters by entering the following command:

```
Monitor IP> monitor bridge list counters summary
```

Note

You cannot use this shortcut to execute a Config command while at the Monitor prompt. Nor can you execute a Monitor command while at the Config prompt. In either of these cases, you must enter the Config or Monitor component first, then enter the desired command.

Using Auto-Prompts

You can perform certain tasks either by entering commands and parameter variables as a string, or by responding to instructional prompts that walk you through each item you must enter.

Examples

To back up a switch's configuration database by entering the command and variables as a string, you enter a string similar to the following example:

```
Boot config>tftp put config 11.22.33.44 /usr/local/tftp/switch11.cfg
```

Alternatively, to back up a switch's configuration database in response to auto-prompts, you perform the following steps:

Step	Action
1	At the Boot config prompt, enter tftp put .
2	Press Return. The following message is displayed: local filename [CONFIG]? CONFIG is the default file name. It cannot be changed.
3	Press Return. The following message is displayed: remote host [0.0.0.0]?
4	Enter the IP address of the host to which you want to copy the configuration database. The default (0 . 0 . 0 . 0) is an invalid address, and is meant only to show the format of the address.
5	Press Return. The following message is displayed: host filename [0B070706.cfg]?
6	Enter the path, followed by a unique file name, to identify the location on the remote system where you want to back up the configuration database. The default is 0B070706 . cfg. It is recommended that you give the file a name that is descriptive of the switch from which it originates. By doing so, you should more easily be able to distinguish between backup files derived from multiple modules. Example: /usr/local/tftp/switch11.cfg
7	Press Return. The following message is displayed when the transfer is completed: TFTP transfer complete, status: OK

Note

Auto-prompts are not supported for all VNswitch commands.

Using the Web-Based Management Application

The VNswitch module, with V3.0 firmware, includes a built-in web server and management application that allow you to configure and monitor the module over the Internet. You can use either of the following web browsers:

- Netscape V4.0
- Internet Explorer V4.0

Accessing VNswitch Modules Over the Web

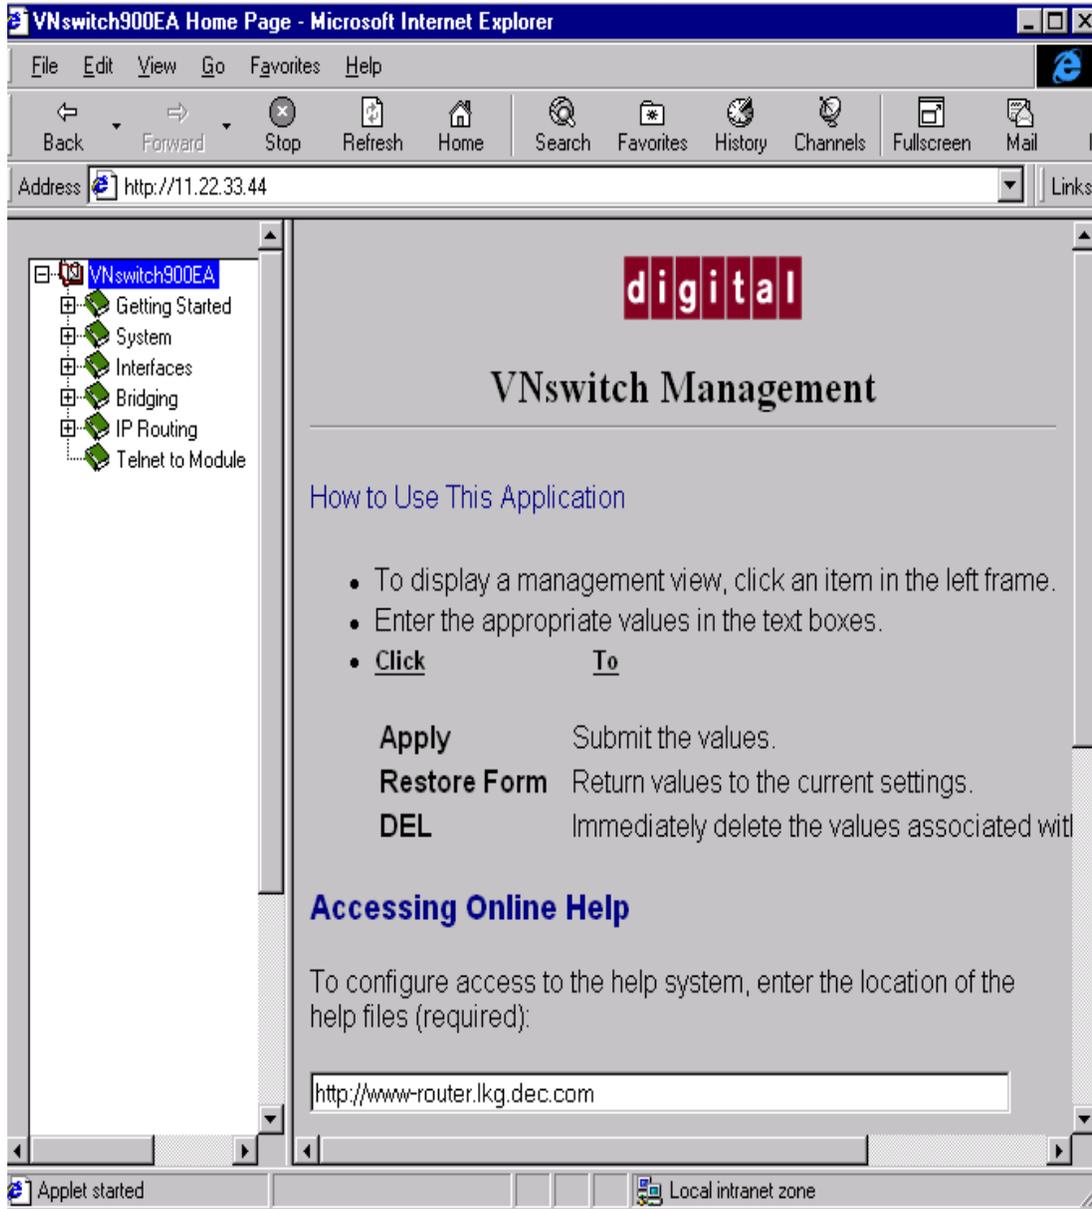
To access a VNswitch module, open your browser and enter the module's IP address in the Location field.

NOTE

For web access, you must first assign an IP address to the module using the CLI. See [Chapter 12](#) for the procedure.

The management application displays with the VNswitch Management window on the right, and an application tree on the left (Figure 2-3). The VNswitch Management window is the first (or top) item in the application tree, which expands with a menu of parameters that you can use to manage the module.

Figure 2-3: Web Management Display



Using the Web-Based Management Application

Managing VNSwitch Modules Over the Web

Once you access the VNSwitch web-based management application, you can configure the module with a limited set of system, interface, bridge, and IP router parameters. Choose a parameter from the application tree, and enter the appropriate information in the related application window. For parameters that you cannot configure with the web-based management application, the application tree contains a Telnet feature that allows you to access the CLI.

Accessing Web Help

The VNSwitch web-based management application includes a comprehensive help system that provides information related to the application windows, plus links to online documentation. To conserve module memory, the web help is made available on the VNSwitch 900 Series Information Library CD. DIGITAL recommends installing the help on a web server. For access to the help, you must specify its location in the VNSwitch Management window (Figure 2-3).

Help is accessible from any application window by clicking the Help button. You can access an overview of the help system from the VNSwitch Management window, or by clicking the Main Contents button in any help window.

Disabling and Enabling the VNSwitch Web Server

The web server in the VNSwitch module is enabled by default. However, you have the option of disabling the server by entering a single command using the CLI. Disabling the server disables the VNSwitch web-based management application.

To disable the web server, enter the following command:

```
Http Config> disable
```

To enable the web server, enter the following command:

```
Http Config> enable
```

Displaying CLI Help

You can obtain help at any of the VNswitch prompts (Main, Config, or Monitor) and at most of the subsystem prompts (Bridge Config>, VSD Config>, and so on) by entering `?`, followed by pressing Return. Help is displayed as a list of the commands available at that prompt level. Use `? (Help)` to list the commands that are available from the current prompt level. You can also enter `? (Help)` after a specific command name to list its options.

Chapter 3

Adding and Managing Users

Overview

Introduction

This chapter provides instructions about assigning login permission and security levels to users who are to manage a DIGITAL VNswitch module.

Users who are assigned administrative permission control access to switch configuration and monitoring tasks. In addition, a user with any of the three security levels (administrative, operations, or monitor) can display a list of users and change their own password.

Caution

Login IDs and passwords are not required at installation. You can, however, configure the VNswitch module to require users to enter a login ID and password. Refer to the [Enabling and Disabling Prompting for ID and Password](#) section for additional information.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Adding Users	3-2
Displaying a List of All Users	3-4
Changing a User's Name, Password, and Security Level	3-5
Enabling and Disabling Prompting for ID and Password	3-8
Deleting Users	3-9

Adding Users

If a VNSwitch administrator enables prompting for ID and password, an individual cannot access switch configuration and management functions unless the user's name is added to a user list, assigned a password, and given one of three security levels (administrative, operations, or monitor).

Refer to the [Enabling and Disabling Prompting for ID and Password](#) section for information about requiring the entry of an ID and password. Refer to [Chapter 1](#) for a description of the three security levels and the access privileges that each provides.

To add an individual to the user list, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter add user .
2	Press Return. The following message is displayed: Enter user name: []?
3	Enter the name of the user (Mary , for example). The name can be a maximum of eight characters and is case sensitive. Spaces are permitted. If the maximum of eight characters is exceeded, the entry is truncated.
4	Press Return. The following message is displayed: Password:
5	Enter a password for the user. The password can be a maximum of 79 characters and is case sensitive.
6	Press Return. The following message is displayed: Enter password again:
7	Reenter the password you entered in step 5 to confirm that it is correct.
8	Press Return. The following message is displayed: Enter permission: (A)dmin, (O)perations, or (M)onitor [A]?
9	Enter either an A , an O , or an M (for administrative, operations, or monitor, respectively) to designate one of the three possible permission levels you can assign. (Refer to Chapter 1 for a description of the three security levels and the access privileges that each provides.)

Step	Action
10	<p>Press Return.</p> <p>The following message and the Config prompt are displayed:</p> <pre>User 'Mary' has been added</pre> <p>If prompting for ID and password mode (enable or disable console login) was changed prior to this procedure, a notification message to this effect is also displayed.</p>

Note

You cannot enable ID and password prompting unless at least one user having administrative permission is added to the user list on the module you are configuring. The message, `Warning: Console login is disabled until an administrative user is added`, is displayed if you attempt to enable password prompting on a module for which there is no administrative user. Refer to the [Enabling and Disabling Prompting for ID and Password](#) section for additional information.

Displaying a List of All Users

You can display a list of all users, including the security (permission) levels to which they are assigned.

To view the list of users, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter list users .
2	Press Return. A list of users is displayed, including the security (permission) level to which they are assigned.

Example:

```
USER  PERMISSION
joe   Operations
mary  Admin
peter Monitor
Console login-prompting is enabled
```

Changing a User's Name, Password, and Security Level

You can change the name, password, and security level of another user, or change your own password.

Changing Your Own Password

You can change your own password, regardless of the security level to which you are assigned. To change your password, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter change password .
2	Press Return. The following message is displayed: Enter current password:
3	Enter your current password.
4	Press Return. The following message is displayed: Password:
5	Enter your new password. The password can be a maximum of eight characters and is case sensitive.
6	Press Return. The following message is displayed: Enter password again:
7	Reenter the password you entered in step 5 to confirm that it is entered correctly. If the confirmation entry does not match the password you entered in step 5, your old password remains in effect.

Changing a User's Name, Password, and Security Level

Changing Another User's Password or Security Level

You must be assigned to the administrative security level to change another user's password or security level. To change another user's password or security level, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter change user .
2	Press Return. The following message is displayed: Enter user name: []?
3	Enter the name of the user whose information you want to change.
4	Press Return. The following message is displayed: Change password? (Yes or [No]):
5	Enter Yes if you want to change the user's password. Enter No if you do not want to change the user's password. No is the default.
6	Press Return. If you entered No in step 5, the following message is displayed: Change permission? (Yes or [No]): Go to step 11. If you entered Yes in step 5, the following message is displayed: Password: Go to step 7.
7	Enter the user's new password. The password can be a maximum of eight characters and is case sensitive.
8	Press Return. The following message is displayed: Enter password again:
9	Reenter the password you entered in step 7 to confirm that it is entered correctly.
10	Press Return. The following message is displayed: Change permission? (Yes or [No]):
11	Enter Yes if you want to change the permission level. Enter No if you do not want to change the permission level.

Changing a User's Name, Password, and Security Level

Step	Action
12	<p>Press Return.</p> <p>If you entered No, the Config prompt (Config>) is displayed.</p> <p>If you entered Yes, the following message is displayed:</p> <pre>Enter permission: (A)dmin, (O)perations, or (M)onitor [A]?</pre>
13	<p>Enter either an A, an O, or an M (for administrative, operations, or monitor, respectively) to designate one of the three possible permission levels you can assign. (Refer to Chapter 1 for a description of the three security levels and the access privileges that each provides.)</p>
14	<p>Press Return. The specified changes are made and the Config prompt (Config>) is displayed.</p>

Enabling and Disabling Prompting for ID and Password

You can configure the switch so that users are required to enter an ID and password before the CLI or the web-based management application is displayed. You can also choose to disable ID and password prompts. If ID and password prompting is disabled, full access to all functions is available to any individual who logs in. That is, there are no restrictions to access of functions based on administrative, operations, and monitor privileges. Disabled is the default.

To enable and disable prompting for ID and password, perform the following steps:

Step	Action
1	If you want to enable ID and password prompting, at the Config prompt (Config>), enter enable console-login-prompting . If you want to disable ID and password prompting, at the Config prompt (Config>), enter disable console-login-prompting .
2	Press Return. ID and password prompting is enabled or disabled as specified.
3	Restart the module for the new setting to take effect.

Note

You cannot enable ID and password prompting unless at least one user having administrative permission is added to the user list on the module you are configuring. The message, `Warning: Console login is disabled until an administrative user is added`, is displayed if you attempt to enable password prompting on a module for which there is no administrative user.

Deleting Users

You can delete individuals from the list of users who have access to switch configuration and management functions. You must be assigned to the administrative security level to do so. You can delete users either by deleting one user at a time using the **delete user** command, or by deleting all users from the list using the **clear users** command.

Deleting a Single User

To delete a single user, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter delete user .
2	Press Return. The following message is displayed: Enter user name: []?
3	Enter the name (Mary , for example) of the user you want to delete.
4	Press Return. The following message is displayed: Delete 'Mary'? (Yes or [No]):
5	Enter Yes if you want to delete the user. Enter No if you do not want to delete the user.
6	Press Return. If you entered Yes , the following message and the Config prompt are displayed: User 'Mary' has been deleted If you entered No , the Config prompt is displayed.

Deleting (Clearing) All Users

Deleting all users resets the user list to its factory default, deleting the names, passwords, and associated security levels of all individuals from the list. You must log in using a local console after you clear all users. If ID and password prompting (enabling and disabling remote console login) is currently enabled, the console login setting is temporarily disabled until a new user with administrative privileges is added.

Deleting Users

To delete all individuals from the list of users, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt, enter <code>clear user</code> .
2	Press Return. The following message is displayed: <code>You are about to clear all User configuration information</code> <code>Are you sure you want to do this (Yes or [No]):</code>
3	Enter y if you want to clear all user configuration information to the factory default. Enter n if you do not want to clear all user configuration information to the factory default.
4	Press Return. If you entered y , the following message and the <code>Config></code> prompt are displayed: <code>User configuration cleared</code> If you entered n , the following message and the <code>Config</code> prompt (<code>Config></code>) are displayed: <code>Aborted</code>

Chapter 4

Configuring and Monitoring Switchwide Parameters

Overview

Introduction

DIGITAL VNswitch functional components can be divided into those that provide and affect specific network operations such as transparent bridging and interface-dependent functions, and those that provide and affect switchwide operation. This chapter provides information about how to manage and monitor functions that affect switch-wide operation.

In This Chapter

This chapter discusses the following topics.

Topic	Page
Resetting (Clearing) NVRAM to Default Values	4-2
Setting the Session Inactivity Timer	4-5
Assigning a Host Name to the VNswitch Module	4-6
Displaying General Information About the Switch	4-7
Setting and Viewing Clock Time	4-11
Monitoring Switch Memory	4-16
Monitoring Crash Counts and Restart or Reload Data	4-18

Resetting (Clearing) NVRAM to Default Values

You can reset nonvolatile memory (NVRAM) configuration parameters to their factory default values for one or all of the following functional components:

- Transparent bridging, including VLANs
- Backup (installation and dump) file locations
- Protocols (ARP, SNMP, and TCP/IP Host Services)
- ATM interface configuration settings
- Event Logging System (ELS)
- Users and passwords
- Time of day

You cannot reset Ethernet, Fast Ethernet, and FDDI interface configuration settings to factory defaults using the CLI. You must use the Reset with Factory Defaults option from the VNswitch module installation menu. Refer to your VNswitch module installation and configuration documentation for information about how to do so.

Refer to [Appendix B](#) for complete lists of default values for all functional components.

To reset the current configuration parameters for one or all functional components to their default values, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter <code>clear command-option</code> , where <i>command-option</i> is the command you must enter to specify the components for which you want to reset defaults. Refer to Table 4-1 for a list of commands, and for a description of the functional component for which you reset defaults when you enter the command.

Resetting (Clearing) NVRAM to Default Values

Step	Action
2	<p>Press Return.</p> <p>If you use either the bridge or the all command option, the following message is displayed:</p> <pre>You are about to clear all Bridge configuration information *** WARNING *** This will invoke an automatic RESTART Are you sure you want to do this (Yes or [No]):</pre> <p>If you use any of the other command options, the following message is displayed:</p> <pre>You are about to clear all SNMP configuration information Are you sure you want to do this (Yes or [No]):</pre>
3	<p>Enter y if you want to reset parameters for the specified functional component to their defaults.</p> <p>Enter n if you want to abort the command.</p>
4	<p>Press Return.</p> <p>If you are resetting parameters for all functional components or for boot configurations, the module is automatically restarted and the Main prompt (Main>) is displayed.</p> <p>If you are resetting parameters for any functional component other than for all functional components or for boot configurations, the parameters are reset in memory and the Config prompt (Config>) is again displayed. The operational state of the module may not immediately reflect a cleared configuration. You should restart the module to ensure that the changes take effect.</p>

Resetting (Clearing) NVRAM to Default Values

Table 4-1: Clear Command Options and Descriptions

Command	Description
<u>bridge</u>	Resets all NVRAM configuration parameters to factory defaults, such that transparent bridging is run on each port. All static or permanent address filters, as well as protocol filters are removed. All ports are reassigned to the default VLAN.
<u>boot</u>	Resets NVRAM parameters for the backup file locations of all image (installation), configuration, and dump files.
<u>arp</u>	Resets the Address Resolution Protocol.
<u>snmp</u>	Resets the SNMP Protocol.
<u>tcp/ip</u>	Resets Host Services (HST). IP addresses are removed, disabling remote management of the switch.
<u>atm</u>	Resets all ATM physical and logical interface settings.
<u>els</u>	Resets all Event Logging System NVRAM configuration parameters.
<u>time</u>	Resets the time of day clock (switchwide parameter).
<u>user</u>	Resets user access and passwords. If ID and password prompting (enabling and disabling remote console login) is currently enabled, the console login setting is temporarily disabled until a new user with administrative privileges is added. (Refer to Chapter 3 for information about enabling and disabling prompting for ID and password.)
<u>all</u>	Resets NVRAM configuration parameters for most functional components including transparent bridging, VLANs, protocols, backup file locations, ATM interface settings, users and passwords, time of day, and ELS. You cannot reset Ethernet, Fast Ethernet, and FDDI interface configuration settings to factory defaults using the CLI. You must use the Reset with Factory Defaults option from the VNswitch module installation menu. Refer to your VNswitch module installation and configuration documentation for information about how to do so.

Setting the Session Inactivity Timer

You can set the amount of time a local or remote console is inactive before the module automatically logs out a user. This setting affects only those consoles linked to modules on which an ID and password is required to log in. (Refer to [Chapter 3](#) for information about enabling and disabling prompting for ID and password.) The default setting of 0 (zero) turns off the inactivity timer, so that an established session remains active no matter how long a console is inactive.

To set the maximum amount of time a console can remain inactive, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter set inactivity-timer .
2	Press Return. The following message is displayed: <code>Console inactivity timer in minutes [0]?</code>
3	Enter the maximum number of minutes you want a session to remain inactive before a user is automatically logged out. The value entered must be a whole number, and must be in the range of 1 through 65535. The default setting of 0 (zero) turns off the inactivity timer, so that an established session remains active no matter how long a console is inactive.
4	Press Return. The value you entered is set.
5	Restart the switch. Refer to Chapter 10 for information about how to restart the switch.

Example

```
Config>set inactivity-timer 6
```

Assigning a Host Name to the VNswitch Module

You can identify the VNswitch module by assigning it a host name. The name is used only for descriptive or informational purposes and does not affect or change the address of the switch.

To assign a host name to the switch, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter set hostname .
2	Press Return. The following message is displayed: What is the new host name []?
3	Enter a name for the switch. The name can be composed of alphanumeric characters and can be up to 80 characters long. Spaces are permitted.
4	Press Return. The following message is displayed: Do you want to include host name in prompts ([Yes] or [No]) If you choose yes, only the portion of the name up to the first period is displayed as the port of the prompt.
5	Press Return. The name you specified is assigned to the switch and the Config> prompt is displayed.
6	Restart the switch. Refer to Chapter 10 for information about how to restart the switch.

Example

```
Config>set hostname Bldg 3-2
```

Displaying General Information About the Switch

You can display two reports that include general information about the switch. One report is accessed from the Config prompt (`Config>`), and the other is accessed from the Monitor prompt (`Monitor>`).

Information Displayed from the Config Prompt

The report you display from the Config prompt (`Config>`) includes the name of the responsible contact person and the location of the module. It also includes such switchwide parameters and values as the maximum packet size, whether remote access using a modem is enabled, and the amount of available configuration memory.

To view general information about the switch from the Config prompt, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter list configuration .
2	Press Return. A report that includes general information about the switch is displayed.

Example

```
Config>list configuration
```

```
Hostname: [none]
Maximum packet size: [autoconfigured]
Maximum number of global buffers: [autoconfigured]
Logging disposition: detached
Console inactivity timer (minutes): 0
Physical console login prompting: enabled
Modem control: disabled
Contact person for this node: [none]
Location of this node: [none]
```

```
Configurable Protocols:
```

```
Num Name Protocol
3 ARP Address Resolution
11 SNMP Simple Network Management Protocol
23 Bridge Transparent Bridging
24 HST TCP/IP Host Services
```

```
127616 bytes of configuration memory free
```

Displaying General Information About the Switch

Information Displayed from the Monitor Prompt

The report you display from the Monitor prompt (`Monitor>`) includes information about switch interfaces and protocols and the setup port baud rate used by local consoles.

To view general information about the switch from the Monitor prompt, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter list all .
2	Press Return. A report that includes general information about the switch is displayed.

Displaying General Information About the Switch

Example

The following is an example of the information displayed for an FDDI module. The report is divided into three sections. The first section includes switch identification information and the setup port baud rate used by the local console. Refer to [Table 4-2](#) for a description of the last two sections of the report (protocols and interfaces).

Monitor> **list all**

```
900EF-MM, 12 Eth 1 FDDI, HW=04f1E1, RO=V1.5.5, #2278, SW=V2.0
Hostname: VN1
Console baud rate: 0
```

```
Num Name Protocol
0 IP DOD-IP
3 ARP Address Resolution Protocol
11 SNMP Simple Network Management Protocol
23 BRIDGE Adaptive Source Routing Transparent Enhanced Bridge
```

```
Num Name Feature
5 RMON Remote Monitoring
6 Mirror Port Mirror
```

46 Interfaces:

Ifc Name	MAC/Data-Link	Hardware	State
0 VNbus/0	VNbus-Encapsulation	VNbus	Up
1 Eth/1	Ethernet/IEEE 802.3	Ethernet/802.3	Up
2 Eth/2	Ethernet/IEEE 802.3	Ethernet/802.3	Up
3 Eth/3	Ethernet/IEEE 802.3	Ethernet/802.3	Up
4 Eth/4	Ethernet/IEEE 802.3	Ethernet/802.3	Down
5 Eth/5	Ethernet/IEEE 802.3	Ethernet/802.3	Down
6 Eth/6	Ethernet/IEEE 802.3	Ethernet/802.3	Down
7 Eth/7	Ethernet/IEEE 802.3	Ethernet/802.3	Down
8 Eth/8	Ethernet/IEEE 802.3	Ethernet/802.3	Down
9 Eth/9	Ethernet/IEEE 802.3	Ethernet/802.3	Down
10 Eth/10	Ethernet/IEEE 802.3	Ethernet/802.3	Down
11 Eth/11	Ethernet/IEEE 802.3	Ethernet/802.3	Down
12 Eth/12	Ethernet/IEEE 802.3	Ethernet/802.3	Down
13 FDDI/13	FDDI/IEEE 802.2	FDDI	Up
14 VLAN/14	VLAN Interface	None	Up
15 VLAN/15	VLAN Interface	None	Up
16 VLAN/16	VLAN Interface	None	Up
17 VLAN/17	VLAN Interface	None	Down
...			
45 VLAN/45	VLAN Interface	None	Down

Displaying General Information About the Switch

Table 4-2: Description of the Configuration Report from the Monitor Prompt

Information Displayed	Description
Num	Number associated with the protocol.
Name	Protocol's short name.
Protocol	Full name of the protocol.
Ifc	Physical and logical interface number.
Name	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
MAC/Data Link	The type of MAC/data link configured for the interface.
Hardware	Interface hardware type.
State	Current state of the interface. The following states are possible: <ul style="list-style-type: none">• <code>Up</code> — The interface is connected and operational.• <code>Down</code> — The interface is not operational and failed self-test.• <code>Disabled</code> — The interface is either temporarily or permanently disabled.• <code>Testing</code> — The interface is in the process of completing a self-test.• <code>Not present</code> — Power-up (restart) diagnostics detected a fault on the interface.

Setting and Viewing Clock Time

You can set the following time-related parameters for the VNswitch module:

- Current time
- Time zone offset (from GMT)
- Frequency with which the module synchronizes with the time host when using a Network Time Protocol (NTP) server

You can also view the current settings to verify that they are set as desired.

Setting the Time

You can set the current time on the VNswitch module clock in one of the following two ways:

- Manually
- Using a Network Time Protocol (NTP) server

Setting the Time Manually

You can set the VNswitch module clock manually, based on the current wall clock time. To set the time, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter time set .
2	Press Return. The following message is displayed: <code>year [0]?</code>
3	Enter the numeric representation for the current year using the full four digits. Example: 1996
4	Press Return. The following message is displayed: <code>month [0]?</code>
5	Enter the numeric representation for the current month. Do not use leading zeros. Examples: 5 (for May) or 11 (for November)
6	Press Return. The following message is displayed: <code>date [0]?</code>

Setting and Viewing Clock Time

Step	Action
7	Enter the current day of the month. Example: 22
8	Press Return. The following message is displayed: hour [0]?
9	Enter the current hour of the day in military time. Examples: 5 (for 5 AM) or 17 (for 5 PM)
10	Press Return. The following message is displayed: minute [0]?
11	Enter the current minute of the hour. Example: 45
12	Press Return. The following message is displayed: second [0]?
13	Enter the current second of the minute. Example: 00
14	Press Return. The clock is set.

Note

Alternatively, the above time can be entered using the following syntax:
time set 1996 11 22 17 45 00

Setting the Time Using an NTP Server

An NTP server provides a single-source location from which multiple devices can derive their clock settings. This helps ensure time synchronization among the various participating devices. The NTP server time is based on Greenwich Mean Time (GMT).

To set the time using an NTP server, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter time host .

Setting and Viewing Clock Time

Step	Action
2	Press Return. The following message is displayed: IP address of time host [0.0.0.0]?
3	Enter the IP address of the NTP server from which you want to derive the current time.
4	Press Return. The clock is set and the Config prompt is displayed.
5	Restart the module if you want the new configuration settings to take effect.

Setting the Time Zone Offset

You can specify the time zone by setting the number of minutes that the current time is offset from Greenwich Mean Time (GMT). Values west of GMT are entered as negative values. For example, Eastern Standard Time (EST) in the United States is 5 hours earlier than GMT. Therefore, the offset for EST is -300 (5 hours x 60 minutes/hour). The default of 0 (zero) indicates no offset from GMT.

To set the time zone offset, perform the following steps:

Step	Action
1	At the Config> prompt, enter time offset .
2	Press Return. The following message is displayed: minutes from GMT (-720 to 720) [0]?
3	Enter the number of minutes the time zone is offset from GMT.
4	Press Return. The time zone offset is configured and the Config> prompt is displayed.

Note

Alternatively, the time zone offset can be entered using the following syntax:
time offset -300

Setting and Viewing Clock Time

Setting Time Host Synchronization

You can set the frequency, in seconds, with which the VNswitch module polls the time host for the current time. This function is applicable only if you are using an NTP server to synchronize the VNswitch module clock. The default of 0 (zero) indicates the switch is not to poll the time host to synchronize clocks. You must configure TCP/IP Host Services for this functionality to be operational. Refer to [Chapter 12](#) for information about configuring TCP/IP Host Services.

To set the frequency with which the VNswitch module polls the time host for the current time, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt, enter <code>time sync</code> .
2	Press Return. The following message is displayed: <code>seconds between time syncs [0]?</code>
3	Enter the number of seconds that is to pass between polling attempts.
4	Press Return. The frequency is set and the <code>Config></code> prompt is displayed.
5	Restart the module if you want the new configuration settings to take effect.

Note

Alternatively, the time host polling frequency can be entered using the following syntax: `Config>time sync 10`

Viewing Clock Time Parameters

You can view the current time settings to verify that they are set as desired. The information includes information about who or what last set the time (either the operator, or the IP address of the time host).

To view the current settings, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt, enter <code>time list</code> .
2	Press Return. A report of the current time settings is displayed, and the <code>Config></code> prompt is displayed.

Example

```
Config> time list
```

```
11:21:21 Monday July 7, 1997
```

```
Set by: 110.110.110.157
```

```
Time Host: 110.110.110.157
```

```
Sync Interval: 1200 seconds
```

```
GMT Offset: -240 minutes
```

Monitoring Switch Memory

You can display a report that contains information about VNswitch module CPU memory, number of buffers, and packet sizes. The information provided by the memory report reflects activity occurring on the VNswitch module's management processor only. It does not include data about memory in the VNswitch module's forwarding subsystem.

To display the report, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter memory .
2	Press Return. The report and the Monitor prompt are displayed.

Example

Monitor>memory

```

                Total  Reserve   Never   Perm   Temp   Prev
                Alloc   Alloc   Alloc   Alloc   Alloc   Alloc
Heap memory    6232191   24784  2926207  3271888  28464   5632
Number of global buffers: Total = 400, Free = 367, Fair = 77, Low = 80
Global buff size: Data = 4478, Hdr = 25, Wrap = 0, Trail = 8, Total = 4516

```

Refer to [Table 4-3](#) for a description of the information.

Table 4-3: Description of the Memory Report

Information Displayed	Description
Heap memory	Memory available for dynamically allocated data structures including, for example, ARP database entries.
Heap memory: Total	Original amount of space available for allocation of memory (heap size total).
Heap memory: Reserve	Minimum amount of memory required by the currently configured protocols.
Heap memory: Never Alloc	Amount of memory that has never been allocated.

Monitoring Switch Memory

Information Displayed	Description
Heap memory: Perm Alloc	Amount of memory requested for permanent allocation.
Heap memory: Temp Alloc	Amount of memory temporarily allocated.
Heap memory: Prev Alloc	Amount of memory temporarily allocated and then returned.
Number of Global Buffers: Total	Original number of global buffers on the system.
Number of Global Buffers: Free	Current number of global buffers available.
Number of Global Buffers: Fair	Fair number of buffers for each interface. (Refer to Number of Global Buffers: Low in this table.)
Number of Global Buffers: Low	Number of free buffers at which the allocation strategy changes to conserve buffers. If the value of Free is less than Low, then buffers are not placed on any queue containing more than the Fair number of buffers.
Global buff size: Data	Maximum data link packet size of all interfaces.
Global buff size: Hdr	Sum of the maximum hardware, MAC, and data link headers of all interfaces.
Global buff size: Wrap	This parameter is always zero and is not currently used by the VNswitch module.
Global buff size: Trail	Sum of the largest MAC and hardware trailers for interfaces.
Global buff size: Total	Size of the largest packet buffer.

Monitoring Crash Counts and Restart or Reload Data

You can display the following information about switch crashes and switch restarts or reloads:

- Number of restarts
- Number of known crashes
- Whether the switch was last reloaded or restarted
- Time elapsed since the last reload
- Time elapsed since the last restart

To display information about crash counts and switch restarts or reloads, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter uptime .
2	Press Return. Information about crash counts and switch restarts or reloads is displayed, and the Monitor prompt is redisplayed.

Example

```
Monitor>uptime
```

```
1 Start, (0 known crashes)
Last Power Up: 41 minutes ago
Last Restart: 41 minutes ago
```

Chapter 5

Configuring Network Interfaces

Overview

Introduction

DIGITAL VNswitch 900 modules support Ethernet, Fast Ethernet, FDDI, and ATM network interfaces, and the VNbus interface. All interfaces are automatically added and configured when a VNswitch module is installed. The configuration settings used are the factory defaults listed in the following sections and in [Appendix B](#). You may need to alter the default settings to maximize network performance, and to accommodate requirements unique to your network environment. This chapter provides information about how to do so.

Note

The only VNbus parameter you can configure is enabling and disabling the VNbus logical interface.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Displaying the Interface Number and Type	5-3
Enabling and Disabling an Interface	5-4
Accessing and Exiting an Interface Prompt	5-7
Configuring an Ethernet Logical Interface	5-8
Configuring a Fast Ethernet Logical Interface	5-11
Configuring an FDDI Logical Interface	5-18

Topic	Page
Configuring ATM Physical and Logical Interfaces	5-31
Configuring the Address Resolution Protocol	5-63

Displaying the Interface Number and Type

You can display the number assigned to an interface and the interface type (Ethernet, Fast Ethernet, FDDI, or ATM). You may typically need to do so before you perform other tasks discussed in this chapter.

To list a switch's interface numbers and types, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter list interfaces .
2	Press Return. A list of the interfaces on the module is displayed, including the interface's assigned number.

Example

The following example applies to the VNswitch 900EF module:

Config>**list interfaces**

```
Ifc  Type
0   VNswitch VBus
1   VNswitch Ethernet
2   VNswitch Ethernet
3   VNswitch Ethernet
4   VNswitch Ethernet
5   VNswitch Ethernet
6   VNswitch Ethernet
7   VNswitch Ethernet
8   VNswitch Ethernet
9   VNswitch Ethernet
10  VNswitch Ethernet
11  VNswitch Ethernet
12  VNswitch Ethernet
13  VNswitch FDDI
14  VLAN      Name: DEFAULT  VNBUS Tag: 65  Ports: 4,5,7-12
15  VLAN      Name: vsd1   Ports: 1,2,6,13
16  VLAN      Name: vsd2   Ports: 3
17  VLAN
...
45  VLAN
```

Enabling and Disabling an Interface

You can enable and disable a logical interface (Ethernet, Fast Ethernet, FDDI, ATM, or VNbus). You may need to disable an interface, for example, if you want to logically disconnect a portion of a network while making significant changes to the network segment's physical environment. Enabled is the default.

Enabling and disabling an interface does not require that you restart the module for the settings to take effect. The setting is effective immediately. These settings survive power outages and module restarts.

Note

You can also enable and disable the ATM logical interface through volatile memory from the `ATM/n Config>` prompt, as described in the [Configuring ATM Physical and Logical Interfaces](#) section later in this chapter. Enabling or disabling the ATM interface from the `ATM/n Config>` prompt allows you to do so without accessing the `Config` prompt and then returning to `ATM/n Config>`.

Disabling an Interface

To disable an interface, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt, enter <code>disable interface interface#</code> , where <i>interface#</i> is the number of the interface you want to disable. Refer to the Displaying the Interface Number and Type section for information about how to determine the number assigned to an interface.
2	Press Return. The interface is disabled as specified.

Enabling and Disabling an Interface

Enabling an Interface

To enable an interface, perform the following steps:

Step	Action
1	<p>At the Config prompt (<code>Config></code>), enter enable interface interface#, where <i>interface#</i> is the number of the interface you want to enable.</p> <p>Refer to the Displaying the Interface Number and Type section for information about how to determine the number assigned to an interface.</p>
2	<p>Press Return.</p> <p>The state of the interface is first tested and the following message is displayed:</p> <pre>Testing interface 13 FDDI/1...</pre> <p>If the test is successful, the following message is displayed and the interface is enabled:</p> <pre>Testing interface 13 FDDI/1...successful</pre> <p>If the test failed, the following message is displayed:</p> <pre>Testing interface 13 FDDI/1...failed</pre> <p>If testing takes longer than 30 seconds, the Config process may time out and display the following message:</p> <pre>Testing interface 13 FDDI/1...still testing</pre>

Accessing and Exiting an Interface Prompt

Each type of interface (Ethernet, Fast Ethernet, FDDI, and ATM) has its own configuration prompt that you must access to configure the interface or to view configuration information associated with the interface. The prompts are accessed from the Config prompt (Config>).

Accessing an Interface Prompt

To access an interface prompt, perform the following steps:

Step	Action
1	Determine the network interface numbers for which the switch is currently configured. Refer to the Displaying the Interface Number and Type section.
2	Record the desired interface numbers.
3	At the Config prompt, enter interface , followed by a space and the number of the logical interface you want to configure. Example: Config>interface 1
4	Press Return. The configuration prompt for the selected interface is displayed on the console. The <i>n</i> refers to the number of the selected interface. The following is a list of possible interface configuration prompts: <ul style="list-style-type: none"> • ETH/<i>n</i> config> (for an Ethernet interface) • FAST/<i>n</i> ETH config> (for a Fast Ethernet interface) • FDDI/<i>n</i> config> (for an FDDI interface) • ATM/<i>n</i> LEC Config> (for an ATM LAN Emulated Client LEC interface) • ATM/<i>n</i> BT Config> (for an ATM Bridge Tunnel interface)

Exiting an Interface Prompt

To exit an interface prompt, type **exit** and then press Return. You exit an interface prompt to return to the next higher level prompt. For example, to return to the Config prompt (Config>) from the Ethernet interface configuration prompt (ETH/*n* Config>), enter **exit**, and then press Return.

Configuring an Ethernet Logical Interface

VNswitch 900EX, 900EF, and 900EA modules each include twelve, 10BaseT Ethernet physical interfaces. The VNswitch 900EE module includes twenty-four, 10BaseT Ethernet physical interfaces. The VNswitch 900LL module includes 12 front panel 10BaseFL 802.3/Ethernet ports and 12 backplane 802.3/Ethernet ports.

You can perform the following tasks when configuring an Ethernet interface:

- Set the interface for a full- or half-duplex circuit.
- Redirect a front panel port to a DEChub ONE docking station AUI port.
- Redirect a front panel port to the DEChub 900 backplane.
- Redirect a front panel port to the DEChub 900 backplane ThinWire segment.
- Display current Ethernet interface configuration parameters.

Setting the Interface for Full-Duplex or Half-Duplex Mode

You can set an Ethernet interface to function in either full-duplex mode or half-duplex mode. On any one VNswitch 900EX, 900EF, 900EA, or 900LL module, a maximum of nine Ethernet interfaces can be set to full-duplex. On any one VNswitch 900EE module, a maximum of 18 Ethernet interfaces can be set to full-duplex, any of the nine interfaces numbered 1 through 12, and any of the nine interfaces numbered 13 through 24.

An interface is typically set for full-duplex mode if it is established as a point-to-point connection to a single device, and for half-duplex if connected to a shared medium supporting multiple devices.

Caution

The 10 Mb/s Ethernet interfaces do not perform auto-negotiation with the device to which it is attempting to establish a connection. Therefore, you must make sure the Ethernet interface you are configuring, and the device with which it is to connect, are both set for the same circuit type (full- or half-duplex mode).

Configuring an Ethernet Logical Interface

To set the Ethernet interface for either full-duplex or half-duplex mode, perform the following steps:

Step	Action
1	At the <code>ETH/n config></code> prompt, enter <code>set duplex-mode-default state</code> , where <i>state</i> is either full (for full-duplex mode) or half (for half-duplex mode). The default is half .
2	Press Return. The interface is set for either full-duplex or half-duplex mode, as specified. The <code>ETH/n config></code> prompt and a message similar to the following are displayed: <code>Default duplex mode set to FULL DUPLEX.</code>

Redirecting Interfaces to the DEChub 900 Backplane

Each of the 12 front panel Ethernet interfaces on a VNSwitch 900EX, 900EF, and 900EA module can be individually redirected to a DEChub 900 backplane LAN segment to create a backplane LAN. Similarly, 12 of the 24 Ethernet interfaces on a VNSwitch 900EE module can be individually redirected to a DEChub 900 backplane LAN segment.

To redirect Ethernet interfaces to the backplane, you must install the optional MultiChassis Manager, a component of the clearVISN network management product. Refer to the documentation accompanying the product for information about how to use MultiChassis Manager.

Redirecting Interface 12 to the DEChub 900 Backplane ThinWire

When connected to a DEChub 900, you can connect a module's Ethernet interface 12 to the ThinWire segment of the DEChub backplane. The ThinWire segment provides backward compatibility for DEChub 90 modules.

To redirect Ethernet interface 12 to the ThinWire segment on the backplane, you must install the optional MultiChassis Manager, a component of the clearVISN network management product. Refer to the documentation accompanying the product for information about how to use MultiChassis Manager.

Redirecting Interface 11 to a DEChub ONE AUI Port

When connected to a DEChub ONE docking station, you can connect the module's front panel interface 11 to the docking station's Attachment Unit Interface (AUI) port. The AUI port can then be used to connect interface 11 to a LAN (backbone) using an 802.3 Medium Attachment Unit (MAU). The MAU, in turn, can be connected to one of several supported media types including unshielded twisted pair, ThinWire, or fiber optic cable, depending on the model installed. For instructions about installing MAUs in a DEChub ONE, refer to the appropriate DEChub ONE installation documentation.

Configuring an Ethernet Logical Interface

To redirect Ethernet interface 11 from the front panel to the AUI port of a DEChub ONE docking station, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter interface 11 . The Ethernet interface configuration prompt (Eth/n config>) is displayed.
2	If you want to redirect Ethernet interface 11 to the AUI port on the module, enter connect back . If Ethernet interface 11 is already redirected to the AUI port and you want to redirect it to the front panel, enter connect front .
3	Press Return. Interface 11 is redirected as you specified in step 2, and the Ethernet configuration prompt is displayed.

Displaying Current Ethernet Configuration Parameters

You can display information about current Ethernet interface configuration settings. To display the information, perform the following steps:

Step	Action
1	At the Eth/n config> prompt, enter list .
2	Press Return. The report and a new Eth/n config> prompt is displayed.

Example

The following is an example of a typical display for an Ethernet interface having 10BaseT MMI. A speed value of 10 Mbits indicates the interface is set for half-duplex mode. A speed value of 20 Mbits indicates the interface is set for full-duplex mode.

```
Eth/1 config>list
```

```
Connect:                FRONT (Ports 11 and 12 only)
Default duplex mode:    HALF DUPLEX
```

Configuring a Fast Ethernet Logical Interface

The VNswitch 900EX module includes twelve, 10BaseT Ethernet physical interfaces and two 100BaseTX or 100BaseFX Fast Ethernet interfaces. You can perform the following tasks when configuring a Fast Ethernet interface:

- Enable, disable, or restart auto-negotiation (100BaseTX only).
- Set the interface for a full- or half-duplex circuit.
- Enable or disable advertising of supported operational modes (100BaseTX only).
- Set the default line speed (100BaseTX only).
- Display current Fast Ethernet interface configuration parameters.

In the following section, information about 100BaseTX ports also applies to the 100BaseTX/FX ports of the VNswitch 900XA.

IEEE Compliance

The VNswitch 900EX module Fast Ethernet interface complies with the IEEE 802.3u standard.

Enabling or Disabling Auto-Negotiation

Auto-negotiation is supported only on 100BaseTX Modular Media Interface (MMI) PHY cards. 100BaseFX MMIs must be set to full-duplex mode manually.

When first installed, a 100BaseTX Fast Ethernet interface negotiates with the destination device to which it is attached. The negotiation process occurs automatically, and involves each device notifying (advertising to) the other what connection capabilities it can support, including line speed and whether the circuit is half- or full-duplex. (Refer to the [Setting Advertised Connection Capabilities](#) section for information about how to set the capabilities that are advertised.) The devices determine what common set of capabilities both support and then, from that common set, select the connection that provides the greatest throughput. The common set of capabilities is typically referred to as the “highest common denominator.” No connection is established if the two devices do not support at least one common line speed and duplex setting combination.

For example, a Fast Ethernet interface on a VNswitch module may support only a full-duplex circuit on either 10BaseT or 100BaseX cable. The device to which it is connected supports a full-duplex circuit on only 100BaseX cable. The two devices “agree” to use a full-duplex circuit at 100 Mb/s, since both support 100BaseX cable, which provides greater bandwidth.

Configuring a Fast Ethernet Logical Interface

You can disable auto-negotiation if, for example, you want to set the line speed and full- or half-duplex settings manually. You can then reenable auto-negotiation at any time. Enabled is the default for auto-negotiation.

Disabling Auto-Negotiation

To disable auto-negotiation, perform the following steps:

Step	Action
1	At the <code>FAST ETH/n config></code> prompt, enter set auto-negotiation disabled .
2	Press Return. The following message is displayed: <pre>This forces the default speed and duplex mode settings to be used. Continue? (Yes or [No]):</pre>
3	Enter y if you want to disable auto-negotiation. The interface will use the default parameters you specify in the Setting Full-Duplex or Half-Duplex Mode Manually section and Setting the Line Speed Manually section, respectively. Enter n if you do not want to disable auto-negotiation.
4	Press Return. If you entered y , auto-negotiation is disabled and the following message and the <code>FAST ETH/n config></code> prompt are displayed: <pre>Auto-Negotiation DISABLED.</pre> If you entered n , auto-negotiation remains enabled and the following message and the <code>FAST ETH/n Config></code> prompt are displayed: <pre>Command aborted. Auto-Negotiation setting not applied.</pre>

Enabling Auto-Negotiation

To enable auto-negotiation, perform the following steps:

Step	Action
1	At the <code>FAST ETH/n config></code> prompt, enter set auto-negotiation enabled .
2	Press Return. Auto-negotiation is enabled and the following message and the <code>FAST ETH/n config></code> prompt are displayed: <pre>Auto-Negotiation ENABLED.</pre>

Restarting Auto-Negotiation

Auto-negotiation is initiated when either of the following events occur:

- The Unshielded Twisted Pair (UTP) cable is first plugged in to the MMI
- One of the devices at either end of the link is powered up, restarted, or a renegotiation request is made.

You may want to restart auto-negotiation if, for example, you change the advertised connection capabilities, or the connection to the device with which the switch is communicating is lost.

To restart auto-negotiation, perform the following steps:

Step	Action
1	At the <code>FAST ETH/n config></code> prompt, enter <code>set auto-negotiation restart</code> .
2	Press Return. The following message and the <code>FAST ETH/n Config></code> prompt are displayed: <code>Auto-negotiation RE-STARTED.</code>

Setting Advertised Connection Capabilities

You can select the supported line speed and duplex setting combinations advertised during auto-negotiation. You can do so only for Modular Media Interface (MMI) PHY cards supporting 100BaseTX and 10BaseT interfaces. The specified connections can include one or more line speed and duplex setting combinations. For example, you may choose to advertise that the interface supports only a 100 Mb/s half-duplex line. Alternatively, you may choose to advertise that the interface supports either a 10 Mb/s full-duplex line or a 100 Mb/s full-duplex line.

Note

The specific line speed and duplex settings that you can set vary, according to the particular Modular Media Interface (MMI) connector installed on the Fast Ethernet interface.

Configuring a Fast Ethernet Logical Interface

To set the connection capabilities that are to be advertised during auto-negotiation, perform the following steps:

Step	Action
1	At the <code>FAST ETH/n config></code> prompt, enter set advertise-capability .
2	Press Return. A message similar to the following is displayed: <code>100Base-TX Full Duplex, (E)nable or (D)isable [ENABLED]?</code> Note: The specific messages displayed vary, depending on the particular MMI installed on the Fast Ethernet interface.
3	Enter enable if you want to advertise support for a line having the speed and duplex setting specified in the message. Enter disable if you do not want to advertise support for a line having the speed and duplex setting specified in the message. The factory default is enable for each line speed and duplex setting for which you are prompted.
4	Press Return. If a message similar to the one shown in step 2 is displayed asking if you want to advertise support for another line speed and duplex setting combination, go to step 3. If the following message and the <code>FAST ETH/n Config></code> prompt are displayed, the selected connection capabilities to be advertised are set: <code>Advertised capability setting applied.</code>

Setting Full-Duplex or Half-Duplex Mode Manually

You can set the Fast Ethernet duplex setting manually, if you do not want to establish the setting automatically through auto-negotiation on a 100BaseTX MMI. The duplex setting must be set manually on a 100BaseFX MMI. An interface is typically set for full-duplex mode if it is established as a point-to-point connection to a single device, and for half-duplex mode if connected to a shared medium supporting multiple devices.

Note

Manually set duplex modes take effect only when auto-negotiation is disabled or not supported. Therefore, if auto-negotiation is supported, you must disable auto-negotiation before the duplex mode setting takes effect.

Configuring a Fast Ethernet Logical Interface

To set the Fast Ethernet interface for either full-duplex or half-duplex mode, perform the following steps:

Step	Action
1	Disable auto-negotiation according to the instructions earlier in this section, if you have not already done so.
2	At the <code>FAST ETH/n config></code> prompt, enter <code>set duplex-mode-default state</code> , where <i>state</i> is either full (for full-duplex mode) or half (for half-duplex mode). The factory default is half .
3	Press Return. The interface is set for either full-duplex or half-duplex mode, as specified. If auto-negotiation is supported, the <code>FAST ETH/n config></code> prompt and a message similar to the following are displayed: <code>Default duplex mode set to FULL DUPLEX.</code> <code>This setting will take effect only when auto-negotiation is disabled.</code> If auto-negotiation is not supported, the <code>FAST ETH/n config></code> prompt and a message similar to the following are displayed: <code>Default duplex mode set to FULL DUPLEX.</code> <code>This setting will take effect immediately.</code>

Setting the Line Speed Manually

You can set the Fast Ethernet line speed manually, if you do not want to establish the speed automatically through auto-negotiation.

Note

Manually set line speeds take effect only when auto-negotiation is disabled or not supported. Therefore, if auto-negotiation is supported, you must disable auto-negotiation before the line speed setting takes effect.

To set the Fast Ethernet interface line speed manually, perform the following steps:

Step	Action
1	Disable auto-negotiation according to the instructions earlier in this section, if you have not already done so.
2	At the <code>FAST ETH/n config></code> prompt, enter <code>set speed-default</code> .

Configuring a Fast Ethernet Logical Interface

Step	Action
3	Press Return. The following message is displayed: Speed Default in Mb/s, 10M or 100M: [100M]?
4	Enter the desired line speed, where the speed is either 10 (for 10 Mb/s) or 100 (for 100 Mb/s). The factory default is 100 .
5	Press Return. The interface is set at the desired line speed, as specified. If auto-negotiation is supported, the FAST ETH/ <i>n</i> config> prompt and a message similar to the following are displayed: Default speed set to 100 Mb/s. This setting will take effect only when auto-negotiation is disabled. If auto-negotiation is not supported, the FAST ETH/ <i>n</i> config> prompt and a message similar to the following are displayed: Default speed set to 100 Mb/s. This setting will take effect immediately.

Displaying Current Fast Ethernet Configuration Parameters

You can display information about current Fast Ethernet interface configuration settings. To display the information, perform the following steps:

Step	Action
1	At the FAST ETH/ <i>n</i> config> prompt, enter list .
2	Press Return. The report and a new FAST ETH/ <i>n</i> config> prompt is displayed.

Configuring a Fast Ethernet Logical Interface

Example

The following is an example of a typical display for a Fast Ethernet interface having a 100BaseTX MMI:

```
FAST ETH/n config>list
```

```
Modular media interface: 100BASE-TX (DEXVU-AA)
Default speed:           100 Mb/s
Default duplex mode:     HALF DUPLEX
Auto-Negotiation:       ENABLED
Auto-Neg Capabilities:
  100BASE-TX Full Duplex:  ENABLED
  100BASE-TX Half Duplex:  ENABLED
  10BASE-T Full Duplex:    ENABLED
  10BASE-T Half Duplex:    ENABLED
```

Configuring an FDDI Logical Interface

The VNswitch 900EF module includes twelve, 10BaseT Ethernet physical interfaces and 1 FDDI interface pair. You can perform the following tasks when configuring an FDDI logical interface:

- Set the station type.
- Set Link Error Rate alarm and cutoff values.
- Enable and disable SMT notification.
- Set token passing and frame timing parameters.
- Enable and disable purging of bad frames.
- Set the interface for a full- or half-duplex circuit.
- Reset all configuration parameters to default values.
- Display current FDDI interface configuration parameters.

Refer to the [Configuring an Ethernet Logical Interface](#) section for information about how to configure 10 Mb/s Ethernet interfaces on a VNswitch 900EF.

Caution

It is recommended that the only configuration tasks you perform on an FDDI interface is setting the station type, setting the interface for full- or half-duplex mode, resetting parameters to defaults, and displaying current parameters. It is recommended that you alter the remaining settings *only* if you fully understand the effect the change will have on the entire network.

Supported Station Types

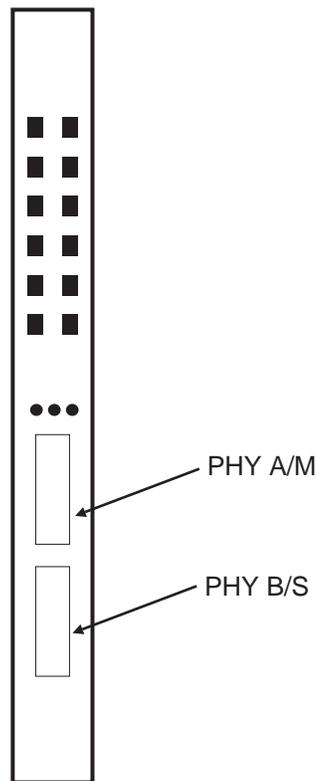
VNswitch 900EF modules can serve as any one of the following FDDI station types:

- Dual Attachment Station (DAS)
- Single Attachment Concentrator (SAC)

PHY Port Identification

The VNswitch FDDI interface is composed of two PHY ports labeled A/M and B/S (Figure 5-1).

Figure 5-1: VNswitch PHY Ports



LKG-10303-96F

The PHY A/M port functions as either an A port or an M port, depending on the module's station type. Similarly, the PHY B/S port functions as either a B port or an S port. Table 5-1 shows the relationship between station type and PHY ports.

Configuring an FDDI Logical Interface

Note

The FDDI logical interface on a VNswitch 900 EF is assigned interface number 13, and bridge port number 13. The FDDI interface number is different on other models.

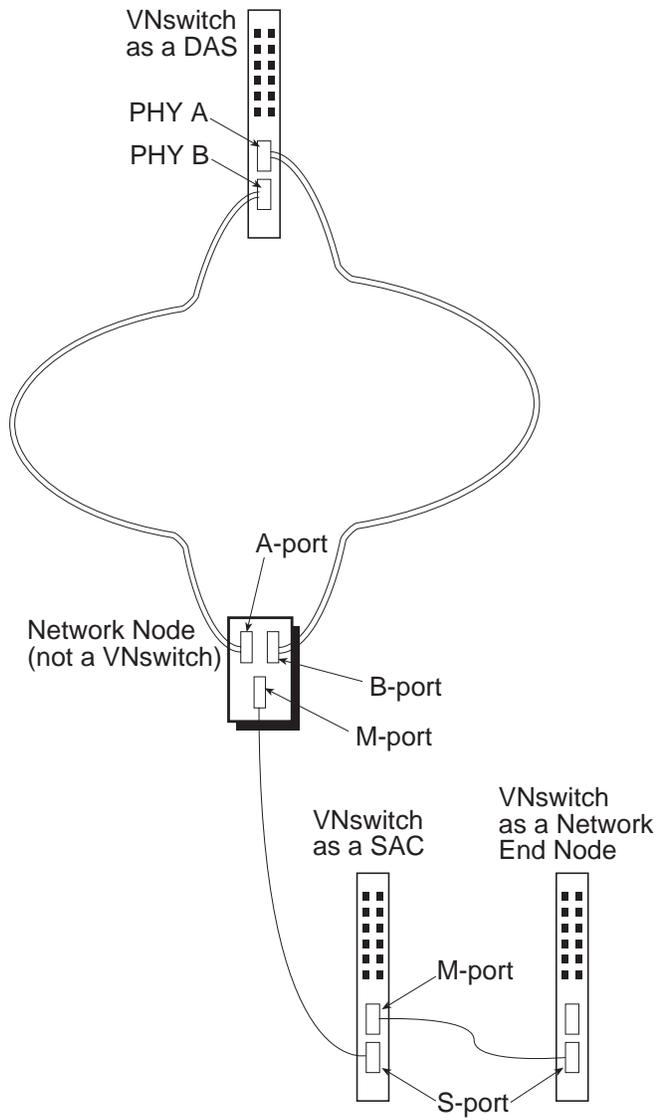
Table 5-1: Station Types and PHY Ports

If . . .	Then . . .
The VNswitch is cabled as a DAS	PHY A/M can be used either as an A port connected to the FDDI ring, or to the M port of a DAC if in a dual homing topology. PHY B/S can be used either as a B port connected to the FDDI ring, or to the M port of a DAC if in a dual homing topology.
The VNswitch is cabled as a SAC	PHY A/M is used as an M port connected to the S port of a SAS. PHY B/S is an S port connected to the M port of a Dual Attachment Concentrator (DAC).
The VNswitch is cabled as a SAS	PHY B/S is an S port connected to the M port of either a SAC or DAC. PHY A/M is not used.

[Figure 5-2](#) illustrates the PHY port connections when the VNswitch serves as a DAS, SAC, and SAS.

Configuring an FDDI Logical Interface

Figure 5-2: Sample FDDI Configuration and Associated PHY Ports



LKG-10305-96F

Configuring an FDDI Logical Interface

Setting the Station Type

The FDDI interface must be configured to reflect the station type for which the physical cables are connected. For example, if the VNswitch 900EF module is cabled to function as a Single Attachment Concentrator, you must set the station type for the interface to SAC. The default installation setting is DAS.

To set or change the station type, perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter set station <i>station-type</i> , where <i>station-type</i> is das for a Dual Attachment Station, or sac for either a Single Attachment Concentrator or a Single Attachment Station. The default is das .
2	Press Return. The interface is set as specified, and the <code>FDDI/n config></code> prompt is displayed.

Setting the Link Error Rate Alarm

The VNswitch 900EF module monitors the number of link errors that occur over each PHY port's link. The types of link errors it monitors includes, for example, line state violations. Notification messages are sent to the SMT management station when the link error rate exceeds a number referred to as the alarm value. You can set the alarm value that triggers notification messages sent to the management station. The value is set independently for each of the module's PHY ports. For example, the value for PHY port A is set independently of the value for PHY port B on a DAS.

Note

Link error rate is expressed as a negative exponential value of base 10. For example, an error rate of 10^{-4} is equal to 0.0001 bits per second, 10^{-5} is equal to 0.00001 bits per second, and so on.

To set the link error alarm, perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter set phy port <i>ler-alarm</i> , where <i>port</i> is either a/m for an A/M PHY port, or b/s for a B/S PHY port.
2	Press Return. The following message is displayed: Link error alarm rate in -Log10: [8]?

Configuring an FDDI Logical Interface

Step	Action
3	Enter a whole number from 4 through 12, where the number is a negative exponential value of base 10. Therefore, for example, entering a 4 sets the value at 10^{-4} (0.0001 bits per second), entering a 5 sets the value at 10^{-5} (0.00001 bits per second), and so on. The default value is 8.
4	Press Return. The link error rate alarm is set as specified, and the FDDI/ <i>n config</i> > prompt is displayed.

Automatically Disconnecting Nodes Causing Excessive Link Errors

The VNswitch 900EF module monitors the number of link errors that occur over each PHY port's link. The types of link errors it monitors includes, for example, line state violations. The neighbor node connected to a link is disconnected from the network if the error rate exceeds a number referred to as the cutoff value. You can set the cutoff value that triggers disconnection of a neighbor node. The value is set independently for each of the module's PHY ports. For example, the value for PHY port A is set independently of the value for PHY port B on a DAS.

If disconnection of a neighbor node is triggered, the FDDI network establishes a loopback (wrap) onto the secondary ring at the node that initiated the disconnect. The connection is automatically reestablished when the link is again active.

Note

Link error rate is expressed as a negative exponential value of base 10. For example, an error rate of 10^{-4} is equal to 0.0001 bits per second, 10^{-5} is equal to 0.00001 bits per second, and so on.

To set the link error cutoff, perform the following steps:

Step	Action
1	At the FDDI/ <i>n config</i> > prompt, enter set phy port ler-cutoff , where <i>port</i> is either a/m for an A/M PHY port, or b/s for a B/S PHY port.
2	Press Return. The following message is displayed: Link error cutoff rate in -Log10: [8]?

Configuring an FDDI Logical Interface

Step	Action
3	Enter a whole number from 4 through 15, where the number is the exponential value of base 10. Therefore, for example, entering a 4 sets the value at 10^{-4} (0.0001 bits per second), entering a 5 sets the value at 10^{-5} (0.00001 bits per second), and so on. The default value is 8.
4	Press Return. The link error rate cutoff is set as specified, and the <code>FDDI/n config></code> prompt is displayed.

Enabling and Disabling SMT Notification

FDDI stations use Status Report Frames (SRFs) to notify the SMT management station about certain events or conditions. The events and conditions can include, for example, ring wrap, exceeding thresholds such as link error alarm and cutoff values, and the detection of an illegal port type configuration.

To set the status report policy, perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter set status-report-policy state , where <i>state</i> is either on or off . Enter on if you want the module to send SRFs to the management station. Enter off if you do not want the module to send SRFs. The default is on .
2	Press Return. SMT notification is enabled or disabled, as specified, and the <code>FDDI/n config></code> prompt is displayed.

Setting Token Passing and Frame Timing Parameters

FDDI uses token passing to control which station has access to the network at any given time. When a station that is ready to transmit data receives the token, the station removes the token from the ring and transmits its data. When data transmission is complete, the station places the token back on the ring.

A token can, occasionally, become lost or corrupted. FDDI stations use a configurable setting called the Maximum Token Rotation Time (tmax-lower-bound) to regenerate a token if this situation occurs. The Maximum Token Rotation Time is the length of time a station waits to receive the token, since the last time it had possession. If the Maximum Token Rotation Time is exceeded, the station then uses a second value, the Requested Token Rotation Time (t_req) to negotiate for token possession. The node

Configuring an FDDI Logical Interface

with the lowest Requested Token Rotation Time wins the bid for the token. If more than one station has the same Requested Token Rotation Time, the station with the highest MAC address wins.

Similarly, FDDI stations use a configurable setting called the Valid Transmission Time (tvx-timer) to determine the maximum amount of time the interface waits to receive a *valid* data frame. The absence of valid data frames for an extended period of time might be caused by conditions such as babble errors, repeated partial frames, or loss of the token. If the Valid Transmission Time is exceeded, the token claim process is reinitiated.

Setting the Maximum Token Rotation Time

To set the Maximum Token Rotation Time (tmax-lower-bound), perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter set tmax-lower-bound #-of-milliseconds , where #-of-milliseconds is equal to the maximum number of milliseconds the interface waits for a token. The #-of-milliseconds can be a number from 8 through 1331 milliseconds. The default is 167 milliseconds. Note: The software rounds the specified value to the nearest multiple of 5.24288 milliseconds. For example, the default value of 167 milliseconds is interpreted by the software as 167.77216 milliseconds.
2	Press Return. The Maximum Token Rotation Time is set and the <code>FDDI/n config></code> prompt is displayed.

Setting the Requested Token Rotation Time

To set the Requested Token Rotation Time (t_req), perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter set t_req #-of-milliseconds , where #-of-milliseconds must be a value that is greater than the Minimum Token Rotation Time (6 milliseconds) and less than the Maximum Token Rotation Time (8 - 1331 milliseconds). The default is 8 milliseconds. Note: The software rounds the specified value to the nearest multiple of 20.48 microseconds. For example, the default value of 8 milliseconds is interpreted by the software as 7.9872 milliseconds.
2	Press Return. The Requested Token Rotation Time is set and the <code>FDDI/n config></code> prompt is displayed.

Configuring an FDDI Logical Interface

Setting the Valid Transmission Timer

To set the Valid Transmission Time (tvx-timer), perform the following steps:

Step	Action
1	<p>At the <code>FDDI/n config></code> prompt, enter <code>set tvx-timer #-of-milliseconds</code>, where <code> #-of-milliseconds</code> is equal to the maximum number of milliseconds the interface waits for a data frame. The <code> #-of-milliseconds</code> can be a number from 1 millisecond through 5 milliseconds. The default is 2.5 milliseconds.</p> <p>Note: The software rounds the specified value to the nearest multiple of 20.48 microseconds. For example, the default value of 2.5 milliseconds is interpreted by the software as 2.51904 milliseconds.</p>
2	<p>Press Return. The Valid Transmission Time is set and the <code>FDDI/n config></code> prompt is displayed.</p>

Configuring the Interface to Purge Bad Frames From the Ring

The VNSwitch FDDI interface is able to detect and remove frame fragments and No Owner Frames (NOFs) from the ring through a function referred to as the Ring Purger. You can turn the Ring Purger on or off using the CLI.

It is generally recommended that you leave the Ring Purger on. The purger uses available bandwidth on the ring to ensure that fragments and No-Owner Frames are removed. Only a very small percentage of available bandwidth may be lost with the purger running if the ring is operating at very high traffic levels. However, when the purger is running, the frame counter observed on any MAC on the ring shows a much larger count than what may be expected. The high number may be the result of including void frames in the total. The value may not, therefore, necessarily reflect a loss of bandwidth.

If these larger frame counts interfere with measurements you wish to make, then the ring purger can be disabled. Note that all nodes capable of being the ring purger must be disabled in order for there to be no ring purger on the ring (ring purging is an elective process). If you attempt to use the frame counter to make utilization measurements, a more accurate method of doing so is available using the MultiChassis Manager feature that provides ring latency and token counts. Refer to [Chapter 6](#) of this document for information about how to monitor the FDDI interface to view the frame count.

Configuring an FDDI Logical Interface

To turn the Ring Purger on or off, perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter set ring-purger state , where <i>state</i> is either on or off . The default is off .
2	Press Return. The Ring Purger is turned on or off, as specified, and the <code>FDDI/n config></code> prompt is displayed.

Setting the Interface for Full-Duplex or Half-Duplex Mode

You can set the FDDI interface to function in either full-duplex mode or half-duplex mode. If you set the interface for full-duplex mode, the link will function as a full-duplex circuit only if the device with which auto-negotiation occurs is also enabled for full-duplex mode. If not, half-duplex mode is used.

An interface is typically set for full-duplex mode if it is established as a point-to-point connection to a single device, and for half-duplex mode if connected to a shared medium supporting multiple devices.

Since the auto-negotiation process detects a point-to-point connection before entering full-duplex mode, it is unnecessary to disable it on half-duplex configurations.

To set the FDDI interface for either full-duplex or half-duplex mode, perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter set full-duplex state , where <i>state</i> is either on (for full-duplex mode) or off (for half-duplex mode). The default is on .
2	Press Return. The interface is set for either full-duplex or half-duplex mode, as specified, and the <code>FDDI/n config></code> prompt is displayed.

Configuring an FDDI Logical Interface

Resetting All Configuration Parameters to Default Values

You can reset the current interface configuration settings (station type, link error rate alarm, Ring Purger, and so on) to their default values. Refer to the appropriate section in this chapter for information about the default value for a specific parameter, or to [Appendix B](#) for a concise list of FDDI interface defaults.

Note

You must restart the VNswitch 900EF module after resetting the configuration parameters for the default values to take effect. Refer to [Chapter 10](#) for information about how to restart the module.

To reset the current FDDI interface configuration parameters to their default values, perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter clear .
2	Press Return. The following message is displayed: <code>For defaults to take effect, restart or reboot is required.</code>
3	Restart the VNswitch 900EF module according to the instructions provided in Chapter 10 .

Displaying Current FDDI Interface Configuration Parameters

You can display specific information about various FDDI interface configuration settings. To display configuration information, perform the following steps:

Step	Action
1	At the <code>FDDI/n config></code> prompt, enter list configuration , where configuration is the command you must enter to display the desired information. Refer to Table 5-2 for a list of commands, and for a description of the type of information that is displayed when you enter the command.
2	Press Return. The desired information and a new <code>FDDI/n config></code> prompt is displayed.

Example

```
FDDI/n config>list all
```

```
Network Interface: 13
STATION TYPE: DAS
Status Report Policy: On
T_REQ: 7.9872ms
TVX: 2.51904ms
T_MAX: 167.77216ms
LER Cutoff, PHY A/M: 10^-8
LER Cutoff, PHY B/S: 10^-8
LER Alarm, PHY A/M: 10^-8
LER Alarm, PHY B/S: 10^-8
Ring Purger: On
Full Duplex: Off
```

Configuring an FDDI Logical Interface

Table 5-2: List Command Options and Descriptions for FDDI Configurations

Command	Description
station-type	Lists the station type (DAS or SAC) assigned to the switch.
phy	Displays the Link Error Rate cutoff and alarm values for each PHY port on the interface. The link error rate is expressed as a negative exponential value of base 10.
status-report-policy	Displays the state (on or off) of the status report policy.
t_req	Displays the Requested Token Rotation Time, in milliseconds.
tvx-lower-bound	Displays the Valid Transmission Time, in milliseconds.
tmax-lower-bound	Displays the Maximum Token Rotation Time, in milliseconds.
ring-purger	Displays the status (on or off) of the Ring Purger.
full-duplex	Indicates whether the interface is set for full-duplex mode (on) or half-duplex mode (off).
all	Displays information about all of the above commands.

Configuring ATM Physical and Logical Interfaces

The ATM physical interface supports up to 16 logical interfaces, numbered as follows for each VNswitch module:

VNswitch Module	ATM Logical Interface Numbers
VNswitch 900EA	13 through 28
VNswitch 900FA	2 through 17
VNswitch 900XA	3 through 18

Each logical interface can be either an ATM Bridge Tunnel, or a LAN Emulation Client (LEC) connected to an Emulated LAN (ELAN). The ATM logical interfaces can be directed to either the ATM physical interface on the front panel of the module, or to the backplane if the module is installed in a MultiSwitch 900. (Refer to [Chapter 1](#) for information about distinguishing between physical and logical interfaces.)

Configuring the Physical Interface

This section explains the tasks for configuring an ATM physical interface. The tasks are organized according to the types of ATM ModPHY cards installed:

- OC3 for OC3 Interfaces (See [Configuring OC3 Interfaces](#).)
- E1/T1 for DS1 Interfaces (See [Configuring DS1 Interfaces for T1/E1 Links](#).)
- E3/T3 for DS3 Interfaces (See [Configuring DS3 Interfaces for T3/E3 Links](#).)
- Accessing the ATM Physical Configuration Prompt

ATM logical interfaces can be directed to either the ATM physical interface on the front panel of the module, or to the backplane if the module is installed in a MultiSwitch 900. To access the configuration prompt, perform the following steps:

Step	Action
1	To set parameters or display information about the physical interface, from the <code>Config></code> prompt, enter atm .
2	Enter physical at the ATM <code>Config></code> prompt.

Configuring ATM Physical and Logical Interfaces

Step	Action
3	Press Return. Depending on how the interface was set in the DIGITAL clearVISN network management product, information about either the front panel interface or the interface to the backplane is displayed.

Directing the ATM Physical Connection to the Front Panel or the Backplane

You can direct the ATM physical connection to either the module's front panel, or to the backplane of a MultiSwitch 900. For instructions about how to do so, you must install the optional MultiChassis Manager, a component of the DIGITAL clearVISN network management product. Refer to the documentation accompanying clearVISN for information about how to use MultiChassis Manager.

Configuring OC3 Interfaces

When configuring OC3 interfaces, you can perform the following tasks:

- Set the transmission type (SONET or SDH).
- Set the transmission timing.
- Set payload scrambling to help ensure data integrity.
- Set the ModPHY card to loopback mode.
- Display current physical interface settings.

Setting the Transmission Type and Timing

The ATM physical interface can be set to use either the Synchronous Optical Network (SONET) or the Synchronous Digital Hierarchy (SDH) transmission type. SONET is the frame format most commonly used in the United States. SDH is the frame format most commonly used in European countries. The transmission type used by the module must be the same as that used by other devices on an ATM network.

Transmission and receipt of SONET or SDH frames must be synchronized among the VNswitch ATM interface and devices with which the module communicates. The timing can be based on the system clock of either the VNswitch module or some other device on the network.

Configuring ATM Physical and Logical Interfaces

Setting the Transmission Type

To set the transmission type, perform the following steps:

Step	Action
1	Access the ATM Physical Config> prompt.
2	Enter set transmission-type option , where <i>option</i> is sonet or sdh . The factory default is sonet .
3	Press Return. The transmission type is set.

Setting Transmission Timing

To set (synchronize) transmission timing, perform the following steps:

Step	Action
1	Access the ATM Physical Config> prompt.
2	Enter set timing option , where <i>option</i> is local or loop : <ul style="list-style-type: none">• If you want to synchronize the devices according to the VNswitch module's clock, enter local.• If you want to synchronize the devices according to the remote device's clock, enter loop. The factory default is local for both the front panel interface and the DEChub backplane.
3	Press Return. Synchronization is set for the chosen interface.

Configuring ATM Physical and Logical Interfaces

Enabling and Disabling Payload Scrambling

Payload scrambling is a function used to help ensure data integrity at the level of the SONET or SDH payload envelope. This process is enabled, by default, on the VNswitch 900EA. If the VNswitch module is to communicate with a DIGITAL GIGAswitch/FDDI AGL2, you must disable payload scrambling because it is not supported on the GIGAswitch/FDDI AGL2.

To disable or enable payload scrambling for the physical interface, perform the following steps:

Step	Action
1	Access the ATM Physical Config> prompt.
2	Enter enable payload-scramble or disable payload-scramble .
3	Press Return. Payload scrambling is disabled or enabled for the chosen interface and the ATM Physical Config> prompt is displayed.

Enabling and Disabling Loopback Mode for the ModPHY Card

Loopback mode allows you to test the integrity of the ATM ModPHY card, and determine if it can establish framing. If the ModPHY card is connected to the ATM physical interface on the module's front panel, the PHY LED lights green if the card can establish framing. The LED does not light, or lights yellow, if the card cannot establish framing. If the ModPHY card is connected to the backplane, you must use the Monitor process command, **list modpmd-errors**, to determine whether the loopback is successful.

The PHY State field listed in the resulting report indicates whether the loopback is successful. A value of Run indicates the loopback is successful. A value of RCV Down, XMT Down, or R/X Down, indicates the loopback is not successful. Refer to [Chapter 6](#) for information about monitoring the ATM physical interface.

Caution

Communication is lost during a loopback mode. You should not, therefore, use loopback mode during critical business hours unless you notify network managers responsible for other nodes. Disabling loopback automatically reestablishes communication between the VNswitch ATM interface and the network.

To turn the loopback test on or off, perform the following steps:

Step	Action
1	Access the ATM Physical Config> prompt.
2	Enter enable loopback or disable loopback . By default, loopback is disabled.
3	Press Return. Loopback testing is turned on or off for the chosen interface and the ATM Physical Config> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Displaying Current Settings for OC3 Interfaces

You can display information about current ATM physical interface configuration settings. To display current configuration settings, perform the following steps:

Step	Action
1	Access the ATM Physical Config> prompt.
2	Enter list all .
3	Press Return. The desired information and the ATM Physical Config> prompt is displayed.

Example

```
ATM Physical Config>list all
```

```
Logical Port Numbers      : 13 - 28
Loopback                  : Disabled
Payload Scramble          : Enabled
Timing                    : Local
Transmission Type         : SONET
```

Configuring ATM Physical and Logical Interfaces

Configuration Commands for OC3 Interfaces

Table 5-3: Configuration Commands for OC3 Interfaces

Command	Description
enable/disable loopback	Enables or disables loopback testing. Loopback testing is used to test the integrity of the ATM ModPHY card.
enable/disable payload-scramble	Enables or disables payload scrambling. Payload scrambling helps ensure data integrity at the level of the SONET or SDH payload envelope.
list all	Displays all the settings for the ATM physical interface.
list loopback	Displays the status of loopback testing: enabled or disabled.
list payload-scramble	Displays the status of payload scrambling: enabled or disabled.
list timing	Displays the timing set for the interface: either local or loop.
list transmission-type	Displays the transmission type SONET or SDH used on the ATM physical interface.
set timing	Sets the source of the transmit clock. You can synchronize the devices according to either the VNswitch module's clock (local) or to the remote device's clock (loop). The default is local timing.
set transmission-type	Sets the ATM physical interface to use either the Synchronous Optical Network (SONET) or the Synchronous Digital Hierarchy (SDH) transmission type. SONET is the frame format most commonly used in the United States. SDH is the frame format most commonly used in European countries. The transmission type used by the module must be the same as that used by other devices on an ATM network.

Configuring ATM Physical and Logical Interfaces

Configuring DS1 Interfaces for T1/E1 Links

When configuring DS1 interfaces for T1/E1 links, you can perform the following tasks:

- Set the transmission timing (See [Setting Transmission Timing](#))
- Set payload scrambling to help ensure data integrity. (See [Enabling and Disabling Payload Scrambling](#).)
- Initiate a self-test of the ModPHY card. (See [Enabling and Disabling Loopback Mode for the ModPHY Card](#).)
- Set a value for the T1 cable length.
- Set a value for line type.
- Display current physical interface settings.

Except where noted, the tasks for configuring DS1 interfaces are described in this section. All other tasks can be found in the [Configuring OC3 Interfaces](#) section.

Setting the Cable Length and Line Type for DS1 Interfaces

To set the cable length for the T1 point-to-point cable, perform the following steps::

Step	Action
1	From the <code>ATM Physical Config></code> prompt, enter set cable-length .
2	Press Return. The following message is displayed: <code>Cable Length (feet) [0]</code>
3	Enter the actual length (in feet) of the T1 point-to-point cable.
4	Press Return. Based on your answer, the software determines a length category from the following: <code>DSX1_CAB_LEN_LSS_110</code> <code>DSX1_CAB_LEN_LSS_220</code> <code>DSX1_CAB_LEN_LSS_330</code> <code>DSX1_CAB_LEN_LSS_440</code> <code>DSX1_CAB_LEN_LSS_550</code> <code>DSX1_CAB_LEN_LSS_600</code> <code>DSX1_CAB_LEN_GRT_600</code>

Configuring ATM Physical and Logical Interfaces

Setting the Line Type

To set the line type, perform the following steps:

Step	Action
1	From the ATM Physical Config> prompt, enter set line-type option , where <i>option</i> is dsx1e1 for E1 CCITT Recommendation G.704 Table 4a and dsx1esf for T1 Extended Super Frame.
2	Press Return.

Displaying Current Settings for DS1 Interfaces

You can display information about current ATM physical interface configuration settings. To display current configuration settings, perform the following steps:

Step	Action
1	To display current settings for the physical interface, access the ATM Physical Config> prompt.
2	Enter list all .
3	Press Return. The desired information and the ATM Physical Config> prompt is displayed.

Example

ATM Physical Config>**list all**

```
Logical Port Numbers      : 13 - 28
Cable Length (feet)     : Less than 220
Line Type                 : DSX1E1
Loopback                  : Disabled
Payload Scramble         : Enabled
Timing                    : Local
```

Configuring ATM Physical and Logical Interfaces

Configuration Commands for DS1 Interfaces

Table 5-4: Configuration Commands for DS1 Interfaces

Command	Description
enable/disable loopback	Enables or disables loopback testing. Loopback testing is used to test the integrity of the ATM ModPHY card.
enable/disable payload-scramble	Enables or disables payload scrambling. Payload scrambling helps ensure data integrity.
list all	Displays all the settings for the ATM physical interface.
list cable-length	Displays the currently set cable length for the T1 point-to-point cable.
list line-type	Displays the operating mode: dsx1E1 or dsx1ESF.
list loopback	Displays the status of loopback testing: enabled or disabled.
list payload-scramble	Displays the status of payload scrambling: enabled or disabled.
list timing	Displays the timing set for the interface: local or loop.
list transmission-type	Displays displays the transmission type SONET or SDH used on the ATM physical interface.
set cable-length	Sets the currently set cable length for the T1 point-to-point cable. Enter the actual length of the cable used, and the software determines a length category from the following: DSX1_CAB_LEN_LSS_110 DSX1_CAB_LEN_LSS_220 DSX1_CAB_LEN_LSS_330 DSX1_CAB_LEN_LSS_440 DSX1_CAB_LEN_LSS_550 DSX1_CAB_LEN_LSS_600 DSX1_CAB_LEN_GRT_600

Configuring ATM Physical and Logical Interfaces

Command	Description
<code>set line-type</code>	Sets the operating mode. Options are: dsx1E1 for E1 CCITT Recommendation G.704 and dsx1ESF for T1 Extended Super Frame
<code>set timing</code>	Sets the source of the transmit clock. You can synchronize the devices according to either the VNswitch module's clock (local) or to the remote device's clock (loop). The default is local timing.

Configuring DS3 Interfaces for T3/E3 Links

When configuring DS3 interfaces, you can perform the following tasks:

- Set transmission timing. (See [Setting Transmission Timing](#).)
- Set payload scrambling to help ensure data integrity. (See [Enabling and Disabling Payload Scrambling](#).)
- Initiate a self-test of the ModPHY card. (See [Enabling and Disabling Loopback Mode for the ModPHY Card](#).)
- Set the line attenuation on transmission.
- Set the line type.
- Set the transmission power.
- Enable and disable PLCP (Physical Layer Convergence Protocol sublayer.)
- Display current physical interface settings.

Except where noted, the tasks for configuring DS3 interfaces are described in this section. All other tasks can be found in the [Configuring OC3 Interfaces](#) section.

Enabling and Disabling PLCP for DS3 Interfaces

T3 PHY supports transmitting cells over T3 with and without DS-3 PLCP (Physical Layer Convergence Protocol) encapsulation. The default is to have PLCP encapsulation enabled. To be compliant with the ATM Forum specification, continue to operate with PLCP enabled. If ATM Forum compliancy is not a concern, then you can disable PLCP to recover some bandwidth. Note that this value is not a MIB object and cannot be set from SNMP.

Configuring ATM Physical and Logical Interfaces

To enable or disable PLCP encapsulation, perform the following steps:

Step	Action
1	From the <code>ATM Physical Config></code> prompt, enter enable plcp or disable plcp .
2	Press Return. PLCP is enabled or disabled.

Setting Line Attenuation for DS3 Interfaces

To set the line attenuation for transmission on DS3 interfaces, perform the following steps:

Step	Action
1	From the <code>ATM Physical Config></code> prompt, enter set line-building-out option , where <i>option</i> is either high or low . In general, you should not have to change this setting. However, for T3 lines less than 225 feet long, the value should be set to low (bit = 1) to prevent the receiver from being saturated.
2	Press Return. Line attenuation is set for the interface.

Setting the Line Type for DS3 Interfaces

The line type indicates the variety of DS3 C-bit or E3 application implementing this interface. The line type setting affects the interpretation and usage of the error bits on the DS3/E3 interface. When setting the line type, both ends of the link need to be set to the same line type.

To set the line type, perform the following steps:

Step	Action
1	From the <code>ATM Physical Config></code> prompt, enter set line-type option , where <i>option</i> is a value from Table 5-5 .
2	Press Return. The line type is set.

Configuring ATM Physical and Logical Interfaces

Table 5-5: Line Type Options for DS3 Interfaces

Option	Description
dsx3-c-bit-parity	ANSI T1.107-1988 This option applies to T3 and is the default on power up for T3 lines. This line type can operate with or without PLCP encapsulation.
dsx3-clear-channel	ANSI T1.107a-1989 This option applies to T3. This line type can operate with or without PLCP encapsulation.
e3-framed-mode	CCITT Recommendation G.751 This option applies to E3 and is the default on power up for E3 lines.
e3-other-mode	CCITT Recommendation G.832.
e3-plcp-mode	CCITT Recommendation G.751 with PLCP mode encapsulation.

Setting the Transmission Power for DS3 Interfaces

To set the transmission power for a DS3 interface, perform the following steps:

Step	Action
1	From the <code>ATM Physical Config></code> prompt, enter set transmission-power option , where <i>option</i> is high or low . The factory default for T3 lines is low. For E3 lines, the default is high. It is strongly recommended that you use the factory default setting.
2	Press Return. The transmission power is set.

Displaying Current Settings for DS3 Interfaces

You can display information about current ATM physical interface configuration settings. To display current configuration settings, perform the following steps:

Step	Action
1	To display current settings for the physical interface, access the <code>ATM Physical Config></code> prompt.
2	Enter list all .

Configuring ATM Physical and Logical Interfaces

Step	Action
3	Press Return. The desired information and the ATM Physical Config> prompt is displayed.

Example

ATM Physical Config>**list all**

```
Logical Port Numbers      : 13 - 28
Line Building Out        : Low
Line Type                 : DSX3 C Bit Parity
Loopback                 : Disabled
Payload Scramble         : Enabled
PLCP                     : Enabled
Transmission Power       : Low
Timing                   : Local
```

Configuring ATM Physical and Logical Interfaces

Configuration Commands for DS3 Interfaces

Table 5-6: Configuration Commands for DS 3 Interfaces

Command	Description
enable/disable loopback	Enables or disables loopback testing. Loopback testing is used to test the integrity of the ATM ModPHY card.
enable/disable payload-scramble	Enables or disables payload scrambling. Payload scrambling helps ensure data integrity at the level of the SONET or SDH payload envelope.
enable/disable PLCP	Enables or disables PLCP encapsulation.
list all	Displays all the settings for the ATM physical interface.
list line-building-out	Displays the attenuation on transmission.
list line-type	Displays the type of DS3 C-bit or E3 application implementing this interface. The line type setting affects the interpretation and of the error bits on the DS3/E3 interface.
list loopback	Displays the status of loopback testing: enabled or disabled.
list payload-scramble	Displays the status of payload scrambling: enabled or disabled.
list plcp	Displays the status of PLCP encapsulation: enabled or disabled.
list timing	Displays the timing set for the interface: local or loop.
list transmission-power	Displays the power (high or low) on the transmitter.
set line-building-out	Sets the line attenuation (high or low) for transmission on DS3 interfaces. Specify low for short cables (less than 225 feet long). For long cables (greater than 225 feet long), specify high.

Configuring ATM Physical and Logical Interfaces

Command	Description
set <u>line-type</u>	Sets the type of DS3 C-bit or E3 application implementing this interface. The line type setting affects the interpretation and usage of the error bits on the DS3/E3 interface. Options are: dsx3-c-bit-parity dsx3-clear-channel e3-framed-mode e3-other-mode e3-plcp-mode
set <u>timing</u>	Sets the source of the transmit clock. You can synchronize the devices according to either the VNswitch module's clock (local) or to the remote device's clock (loop). The default is local timing.
set <u>transmission-power</u>	Sets the power (high or low) on the transmitter. The factory default for T3 lines is low; for E3 lines, the default is high. It is strongly recommended that you use the factory default setting.

Configuring a Logical Interface

You can perform the following tasks when configuring an ATM logical interface:

- Enable or disable the interface.
- Configure the interface as an ATM Bridge Tunnel.
- Configure the interface as a LAN Emulation Client (LEC) on an Emulated LAN (ELAN).
- Display the currently selected interface type and state.

Refer to the [Configuring an Ethernet Logical Interface](#) section earlier in this chapter for information about how to configure 10 Mb/s Ethernet interfaces on a VNswitch 900EA.

Logical Interfaces Established at Startup

[Table 5-3](#) lists the ATM logical interfaces automatically established after installation, and the conditions under which they are active. Each of these default connections can be modified. For example, you can change the default LEC configured on interface 13 to an ATM Bridge Tunnel, if desired.

Note

[Table 5-3](#) shows the ATM interface defaults for the EA module, however the default behavior is the same for the FA and XA modules. If an FA or XA module is attached to any other VNswitch, the default type of connection established is an Ethernet ATM Bridge Tunnel.

Configuring ATM Physical and Logical Interfaces

Table 5-7: ATM Interface Defaults at Startup for EA Module

Logical Interface Number	Type of Connection Established	Conditions Under Which Connection Is Established
13	ELAN	The VNswitch ATM interface is connected to a DIGITAL GIGAswitch/ATM, or any UNI 3.1-compliant switch.
27	FDDI ATM Bridge Tunnel	The VNswitch ATM interface is connected to a DIGITAL GIGAswitch/FDDI with an AGL2 Card.
28	Ethernet ATM Bridge Tunnel	The VNswitch ATM interface is connected to the ATM interface of another VNswitch 900EA.

Disabling and Enabling an ATM Logical Interface

You must disable the ATM interface before you reconfigure an ATM Bridge Tunnel or LEC. If you attempt to reconfigure the interface while it is still enabled, CLI and SNMP requests are rejected.

Configuring ATM Physical and Logical Interfaces

The procedure described in this section provides instructions about how to disable and enable the interface dynamically, through volatile memory from the `ATM/n LEC Config>` prompt. Enabling and disabling an ATM interface through volatile memory does not require that you restart the module to take effect. The setting is effective immediately and survives power outages and module restarts.

Note

You can also enable and disable the ATM interface through volatile memory or NVRAM from the Config prompt (`Config>`). Refer to the [Enabling and Disabling an Interface](#) section earlier in this chapter for information about how to do so. However, enabling and disabling the ATM interface from the `ATM/n LEC Config>` prompt allows you to do so without accessing the Config prompt, and then returning to `ATM/n LEC Config>`.

Disabling the ATM Interface

To disable the ATM interface through volatile memory, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter <code>disable</code> .
2	Press Return. The interface is disabled and the <code>ATM/n LEC Config></code> prompt is displayed.

Note

You may not be able to disable an interface if the interface is in the process of completing a self-test. If this is the case, one of two messages is displayed. The first indicates self-test is being canceled. The second indicates self-test is in progress and the interface cannot be disabled at this time. If either message is displayed, you can attempt to disable the interface at a later time.

Configuring ATM Physical and Logical Interfaces

Enabling the ATM Interface

To enable an interface through volatile memory, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter enable .
2	Press Return. The interface is enabled and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Clearing ATM Configuration Information

You can clear all ATM configuration information. When you do this, the switch restarts automatically and removes disable records associated with any ATM logical interface. To determine if disable records exist, use the Config> **list interface** command.

To clear all ATM configuration information, perform the following steps:

Step	Action
1	At the Config> prompt, enter clear atm .
2	Press Return. The following message is displayed: You are about to clear all ATM configuration information. *** WARNING *** This will invoke an automatic RESTART Are you sure you want to do this (Yes or [No]):
3	Enter yes if you want to clear the information and remove disable records associated with any ATM logical interface. The switch restarts automatically.

Displaying the Logical Interface Type and State

To display the type and state of the currently selected interface, perform the following steps: Example

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter list all .
2	Press Return. The report and the ATM/ <i>n</i> LEC Config> prompt are displayed.

Configuring ATM Physical and Logical Interfaces

Step	Action
3	To display the value for an individual parameter, you can enter list option , where <i>option</i> is one of the parameters shown in Table 5-10 .

```
ATM/14 LEC Config>list all
```

```
Interface                : 14 (Bridge Port 14)
Aging Time(seconds)      : 300
Arp Response Time(seconds) : 1
Config Mode              : Automatic
Control timeout(seconds) : 120
Elan Name(Configured)    :
ELan Name(Joined)        : Not available
Flush Timeout(seconds)   : 4
Forward Delay Time(seconds) : 15
Interface State          : Enabled
LEC Status                : Unavailable
LEC MAC Address          : 00-00-F8-4E-9D-0F
LES ATM Address           :
Mac Type                  : Ethernet
Max Frame size           : 1516
Max Retry Count           : 1
Max Unknown Frames Count : 1
Max Unknown Frame Time(seconds) : 1
Path Switch Delay(seconds) : 6
Vcc Timeout Period(seconds) : 1200
```

Configuring an ATM Bridge Tunnel Interface

ATM Bridge Tunnels are Permanent Virtual Channel (PVCs) established to provide a simple point-to-point connection between two devices through an ATM network. The link does *not* require those elements (LAN Emulation Server, LAN Emulation Configuration Server, and Broadcast and Unknown Server, for example) otherwise necessary to establish an emulated LAN.

Configuring ATM Physical and Logical Interfaces

What Type of Tunnel to Configure

You can configure either an Ethernet ATM Bridge Tunnel or an FDDI ATM Bridge Tunnel on each ATM logical interface. The type of bridge tunnel (Ethernet or FDDI) you establish should be that which minimizes the number of points at which Ethernet-to-FDDI translation is required. [Table 5-8](#) provides general guidelines for determining what type of tunnel to configure.

Note

When changing any of the LEC or BT parameters for a logical interface, you must disable the interface. After you have made the changes, reenable the interface.

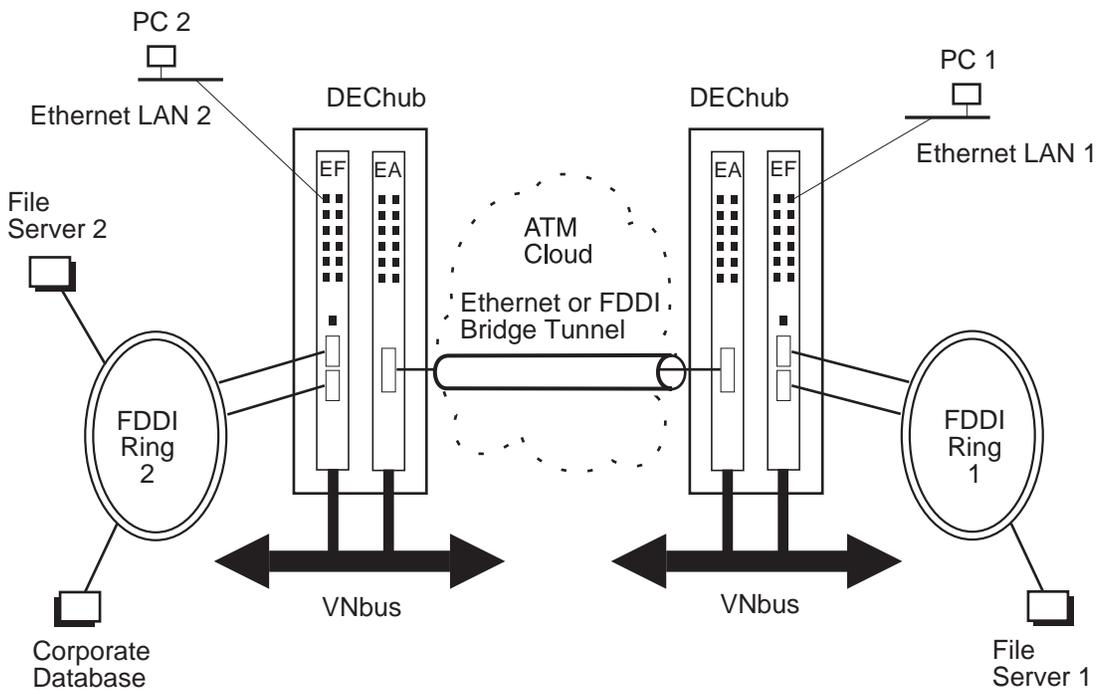
Table 5-8: Guidelines for Determining Bridge Tunnel Type

If. . .	Then. . .
The source and destination devices on either side of an ATM cloud are both on Ethernet LANs	The tunnel should be an Ethernet ATM Bridge Tunnel.
Either the source or destination device is on an FDDI ring, and the other device is on an Ethernet LAN	The tunnel should be an FDDI ATM Bridge Tunnel.
The source and destination devices on either side of an ATM cloud are both on an FDDI ring	The tunnel should be an FDDI ATM Bridge Tunnel.

Using [Figure 5-3](#) as an example, if PC 1 on Ethernet LAN 1 typically requires access to PC 2 on Ethernet LAN 2, the ATM interfaces on VNSwitch modules that connect to the ATM cloud should be configured for an Ethernet ATM Bridge Tunnel.

Again using [Figure 5-3](#) as an example, assume that PC 1 on Ethernet LAN 1 typically requires access to the Corporate Database on FDDI ring 2. In this situation, the ATM interfaces should be configured for an FDDI ATM Bridge Tunnel.

Figure 5-3: ATM Bridge Tunnels



LKG-10304-96F

How to Configure an ATM Bridge Tunnel

This section describes how to configure an ATM Bridge Tunnel on the VNswitch 900EA's ATM interface. Configuring an ATM Bridge Tunnel involves assigning a Virtual Channel Identifier to the Permanent Virtual Channel, and setting the Bridge Tunnel's MAC type.

You must also configure the device on the opposite end of the bridge tunnel. This can be, for example, another VNswitch 900EA, a DIGITAL GIGAswitch/ATM, or another vendor's ATM switch. Refer to the appropriate documentation provided by the manufacturer of that device for configuration instructions.

Each ATM logical interface is automatically set up as either a LAN emulation client (LEC), or a Bridge Tunnel at installation. Refer to [Table 5-7](#) for information about the types of ATM logical interface established at installation, and the conditions under which each is active. If the connection is a LEC, you can change the connection to a Bridge Tunnel.

Configuring ATM Physical and Logical Interfaces

To change the VNswitch 900EA ATM interface from a LEC to an ATM Bridge Tunnel, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter change bridge_tunnel .
2	Press Return. The following message is displayed: Current interface type is LEC. This command will change interface type to Bridge Tunnel Do you want to proceed? (Yes or [No]):
3	Enter y if you want to change the LEC to a Bridge Tunnel. Enter n if you do not want to change the LEC to a Bridge Tunnel.
4	Press Return. If you entered y , the following message is displayed: VPI = 0 in this release VCI(62-1022) [62]: VPI is the Virtual Path Identifier assigned to the Permanent Virtual Channel. The VPI for this release is always 0. VCI is the Virtual Channel Identifier assigned to the PVC. If you entered n , the ATM/ <i>n</i> LEC Config> prompt is displayed.
5	Enter a number from 62 through 1022. This number is the VCI you want to assign to the virtual channel.
6	Press Return. The ATM Bridge Tunnel you specified is configured and the ATM/ <i>n</i> BT Config> prompt is displayed. The Bridge Tunnel is assigned a MAC type of Ethernet, by default. Refer to the Modifying the VCI and MAC Type section for information about how to change the MAC type.

Configuring ATM Physical and Logical Interfaces

Modifying the VCI and MAC Type

You can change the VCI you configured when you first established the Bridge Tunnel. You can also change the MAC type to either an Ethernet ATM Bridge Tunnel or an FDDI ATM Bridge Tunnel. Refer to the [What Type of Tunnel to Configure](#) section for information about how to determine the type of Bridge Tunnel you should establish.

To modify the Bridge Tunnel's VCI or MAC type, perform the following steps:

Step	Action
1	At the <code>ATM/n BT Config></code> prompt, enter Enter set vc .
2	Press Return. A message similar to the following is displayed: <code>VPI = 0 in this release</code> <code>VCI(62-1022):</code> VPI is the Virtual Path Identifier assigned to the Permanent Virtual Channel. The VPI for this release is always 0. VCI is the Virtual Channel Identifier assigned to the PVC.
3	If you want to change the VCI assigned to the PVC, enter a new VCI number from 62 through 1022. If you do not want to change the VCI, reenter the VCI number currently assigned to the PVC (or just press Return).
4	To change the MAC type, use either of these commands: set mac e (if you want to establish an Ethernet Bridge Tunnel) set mac fddi (if you want to establish an FDDI Bridge Tunnel).
5	Press Return. The VCI and MAC type are set as specified, and the <code>ATM BT/n Config></code> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Displaying Current Bridge Tunnel Settings

To display configuration information about the Bridge Tunnel, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> BT Config>enter list all .
2	Press Return. The report is displayed.

Example

```
ATM/27 BT Config>list all
```

```
Interface                : 27 (Bridge Port 27)
Interface Status         : Down
Interface State          : Enabled
Mac Type                  : Ethernet
PVC (VPI,VCI)            : 0,62
LAN FCS                   : No
Quality of Service       : 3
```

Configuring a LAN Emulation Client Interface

This section describes how to configure an ATM interface serving as a LAN Emulation Client (LEC) on an Emulated LAN (ELAN). You must also configure the ATM device on the opposite end of the ELAN. This second device can be, for example, another VNswitch 900EA, another vendor's ATM switch, and servers or workstations directly attached to the ATM network. Refer to the appropriate documentation provided by the manufacturer of that device for configuration instructions.

You can perform the following tasks when configuring the ATM logical interface as an LEC:

- Assign an ELAN name to the logical interface.
- Set the LEC to identify the LES address automatically or manually.
- Set address resolution parameters.
- Set data frame control parameters.
- Set control frame and Data VCC Timeouts.
- Display current ATM logical interface configuration parameters.

Configuring ATM Physical and Logical Interfaces

Caution

It is recommended that the only tasks you perform when configuring a logical interface as a LEC are assigning an ELAN to the interface, identifying the LES address, and displaying current configuration settings. It is recommended that you perform other configuration tasks, such as setting maximum frame size or forward delay time, *only* if you fully understand the effect the change will have on the entire network.

Note

When changing any of the LEC or BT parameters for a logical interface, you must disable the interface. After you have made the changes, reenable the interface.

Assigning an ELAN Name to the Interface

You can identify the ELAN you want the LEC to join by specifying the name of the ELAN. Typically, the list of valid ELAN names is created by the network administrator on the LECS. If you do not identify the ELAN the interface is to join, the LECS directs the connection to a default ELAN compatible with the requirements of the interface.

To identify the ELAN with which you want the LEC's logical interface to join, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set elan_name .
2	Press Return. The following message is displayed: ELan Name []?
3	Enter the (new) name of the ELAN with which you want the interface to establish a connection.
4	Press Return. The specified name is set and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Identifying the LES Address

The LECS is, in part, responsible for maintaining identification and location records of ELANs that are members of an ATM network. Typically, more than one LES is configured in an ATM network, and each is responsible for the identification and location of one of the ELANs on the network. A LEC, attempting to join a particular ELAN, must first determine the address of the LES responsible for the ELAN the LEC is attempting to join.

Identifying which LES is responsible for a target ELAN can be accomplished either automatically or manually. Automatic identification is managed by the LECS based on “look-up” tables previously configured by the network administrator. Automatic identification is the default operational mode. You can choose, however, to disable automatic identification, and to identify the LES manually by specifying the LES ATM address as an LEC configuration parameter. If you choose to provide the LES address manually, the LEC begins the connection process at the Join phase.

Note

Identifying LES addresses is most easily managed on the LAN Emulation Configuration Server. Managing the address tables centrally on the LECS is easier than changing the LES address on each VNswitch module, especially when the physical configuration of a large network changes.

LEC Connection Phases

This section provides a general description of the process any LEC (including the VNswitch module) must complete to establish a connection to an ELAN. The process is divided into the five phases described in [Table 5-5](#). Refer to The ATM Forum Technical Committee’s *LAN Emulation Over ATM* specification for a detailed discussion about ATM LEC functional parameters.

Configuring ATM Physical and Logical Interfaces

Table 5-9: ELAN Connection Phases

Phase	Description
LECS Connect	The LEC establishes a connection with the LECS.
Configuration	The LEC obtains the ATM address of a LES. The LEC may also obtain other configuration parameters.
Join	The LEC establishes its connection(s) with the LES, and determines the operating parameters of the ELAN.
Initial Registration	The LEC <i>may</i> attempt to register additional LAN destinations according to the registration protocol, through which the LEC identifies the LAN destinations it requires for normal operation.
BUS Connect	The LEC establishes a connection with the BUS.

Automatic Identification of the LES Address

Automatic identification of the LES ATM address occurs during the Configuration phase, described in [Table 5-5](#). During the LECS Connect phase, the VNswitch 900EA establishes a Configuration Direct Virtual Channel Connection (VCC) with the LECS. This is the connection over which the VNswitch requests the LECS to identify the address of the LESs responsible for the target ELANs. The LEC sends a configuration request and the LECS responds with the LES address and other configuration parameters.

Automatic identification of the LES address is the default setting for the VNswitch 900EA. If, however, you previously configured the VNswitch module for manual mode and you now want to reset the module to automatic mode, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set config mode automatic .
2	Press Return. The VNswitch module is set to identify the LES address automatically through the LECS, and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Manual Identification of the LES Address

You can configure the VNswitch module so that you provide the address of the LES responsible for the target ELAN. The VNswitch 900EA uses this address during the Join phase (Table 5-5) to establish a Control Direct Virtual Channel Connection (VCC) with the LES. This is the connection over which the VNswitch module determines the operating parameters of the ELAN.

To identify the LES address manually, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set config_mode manual .
2	Enter set les_atm_addr .
3	Press Return. An address format template and the following message are displayed: Les Atm address []?
4	Enter the ATM address of the LES responsible for the ELAN to which the interface is to connect. The ATM address is a 40-character alphanumeric string divided into three segments by periods (.). The first segment is a 26-character prefix. The second is the End Station Identifier (ESI) that is equivalent to the MAC address. The third segment is a 2 character selector byte. Example: 3999990000000008002bb5f380.0000f84de70e.00 or Example: 39 99 99 00 00 00 00 08 00 2b b5 f3 80 00 00 f8 4d e7 0e 00 The address must be entered using the hexadecimal value. It can include spaces or periods.
5	Press Return. The VNswitch module is configured to access the LES identified by the specified address, and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Clearing the LES Address

You can clear the LES address by entering an address of all zeros. Perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set config_mode manual .

Configuring ATM Physical and Logical Interfaces

Step	Action
2	Enter set les_atm_addr .
3	Press Return. An address format template and the following message are displayed: <pre>Les Atm address []?</pre>
4	Enter an address of all zeros. The address must be entered as a 40-character alphanumeric string. Example: 00
5	Press Return. The VNswitch module is configured to access the LES identified by the specified address, and the <code>ATM/n LEC Config></code> prompt is displayed.
6	At the <code>ATM/n LEC Config></code> prompt, enter set config_mode automatic .
7	Press Return. The VNswitch module is set to identify the LES address automatically through the LECS, and the <code>ATM/n LEC Config></code> prompt is displayed.

Setting Address Resolution Parameters

Address resolution is the process by which a LEC associates the MAC address of a LAN destination with the ATM address of another LEC, or the BUS. This process is managed by the LAN Emulation Address Resolution Protocol (LE_ARP), which generates LE_ARP Requests for address information.

You can configure the following parameters that are related to address resolution:

- Expected ARP Response Time
- LE_ARP Request Retries
- Aging the LE_ARP cache

Setting the Expected ARP Response Time

ARP response time is the maximum amount of time a transmitting LEC expects an LE_ARP response to occur after an LE_ARP request is sent. If the length of time it takes for the transmitting LEC to receive a response equals or exceeds the expected time, the LEC initiates a retry.

To set the Expected ARP Response Time, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter set arp_response_time .

Configuring ATM Physical and Logical Interfaces

Step	Action
2	Press Return. The following message is displayed: Arp Response Time #seconds [1]?
3	Enter the number of seconds within which the transmitting LEC expects to receive an LE_ARP response after it sends an LE_ARP request. The range of acceptable values for this parameter is 1 to 30 seconds. The default is 1 second.
4	Press Return. The Expected ARP Response Time is set and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Setting the Number of LE_ARP Request Retries

If a VNswitch LEC interface sends an LE_ARP Request, and a response is not received within the Expected ARP Response Time, the LEC resends the unanswered LE_ARP Request, up to a specified number of retries. You can set the number of times the interface attempts to resend the LE_ARP Request by performing the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set max_retry_count .
2	Press Return. The following message is displayed: Maximum Retry Count [1]?
3	Enter the number of seconds within which the transmitting LEC expects to receive an LE_ARP response after it sends an LE_ARP request. The range of acceptable values is 0 (zero) to 2. The default is 1.
4	Press Return. The Retry Count is set and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Aging the LE_ARP Cache

LEC clients each maintain a table of entries that establishes a relationship between destination MAC addresses and the ATM address to which data frames for the MAC address is sent. Each entry in the table is based on address information included in LE_ARP Responses.

When the LEC receives a frame for transmission, and the destination MAC address is resident in the LE_ARP cache, the client can forward the frame based on the ATM address mapping in the table. This eliminates the need to repeatedly generate an LE_ARP request to determine the destination address. However, the MAC address may move due to a cabling change or a change in the spanning tree by a downstream bridge. This results in invalid mapping of MAC addresses to ATM addresses. To

Configuring ATM Physical and Logical Interfaces

accommodate such changes, entries in the LE_ARP cache are maintained for a limited period of time, thereby forcing generation of new LE_ARP Requests and verification of MAC address-to-ATM address mappings.

Each entry in the cache is retained for a period of time equal to either the Aging Time or Forward Delay Time. Aging time is used to age entries in the cache under normal operation. Forward delay is used to age entries in the LE_ARP cache when a topology change is in progress. Refer to The ATM Forum Technical Committee's *LAN Emulation Over ATM* specification for a detailed discussion about the conditions under which the Aging Time and Forward Delay Time are used.

To set Aging Time for entries in the LE_ARP cache, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set aging_time .
2	Press Return. The following message is displayed: Aging Time #seconds [300]?
3	Enter the number of seconds each entry in the LE_ARP cache is retained. The range of acceptable values is 10 to 300 seconds. The default is 300 seconds.
4	Press Return. The Aging Time is set and the ATM/ <i>n</i> LEC Config> prompt is displayed.

To set Forward Delay Time for entries in the LE_ARP cache, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set forward_delay_time .
2	Press Return. The following message is displayed: Forward Delay Time #seconds [15]?
3	Enter the number of seconds each entry in the LE_ARP cache is retained. The range of acceptable values is 4 to 30 seconds. The default is 15 seconds.
4	Press Return. The Forward Delay Time is set and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Setting the Maximum Unknown Frame Count

The Broadcast and Unknown Server (BUS) manages all broadcast and multicast traffic, and all initial unicast data associated with a MAC address that is unknown by the LEC. An unknown frame is a unicast frame with an unknown destination MAC address. You can set the maximum number of unknown frames that a LEC can send to the BUS within a given period of time. Setting the Maximum Number of Unknown Frames

To set the maximum number of unknown frames sent to the BUS, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter <code>set max unknown frames count</code> .
2	Press Return. The following message is displayed: Maximum Unknown Frame Count [1]?
3	Enter the maximum number of frames sent within the time specified in the Setting Unknown Frame Time section. The range of acceptable values is 1 to 10 frames. The default is 1 frame.
4	Press Return. The Maximum Unknown Frame Count is configured as specified, and the <code>ATM/n LEC Config></code> prompt is displayed.

Setting Unknown Frame Time

To set the time period within which the maximum number of unknown frames is not to be exceeded, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter <code>set max unknown frame time</code> .
2	Press Return. The following message is displayed: Maximum Unknown Frame Time #seconds [1]?
3	Enter the maximum number of seconds within which the maximum number of unknown frames is not to be exceeded. The range of acceptable values is 1 to 60 seconds. The default is 1 second.
4	Press Return. The Unknown Frame Time is configured as specified, and the <code>ATM/n LEC Config></code> prompt is displayed.

Setting Control Frame and Data VCC Timeouts

Configuring ATM Physical and Logical Interfaces

You can set the maximum amount of time a LEC waits for a response after sending a request. If this value is exceeded, the LEC issues another request.

You can also set the maximum amount of time a Data Direct VCC is allowed to be inactive (no data frame traffic is transmitted or received). If the Data Direct VCC timeout value is exceeded, the VNswitch LEC interface releases the VCC.

Setting the Control Frame Timeout

To set the Control Frame Timeout, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter <code>set control_timeout</code> .
2	Press Return. The following message is displayed: <code>Control Timeout #seconds [120]?</code>
3	Enter the maximum number of seconds request and response control frame interactions are to occur. The range of acceptable values is 10 to 300 seconds. The default is 120 seconds.
4	Press Return. The Control Timeout value is configured as specified, and the <code>ATM/n LEC Config></code> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Setting the VCC Timeout Value

To set the VCC Timeout, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set vcc_timeout .
2	Press Return. The following message is displayed: Vcc Timeout Period #seconds [1200]?
3	Enter the amount of time after which an inactive Data Direct SVC should be closed. The default is 20 minutes.
4	Press Return. The VCC Timeout value is configured as specified, and the ATM/ <i>n</i> LEC Config> prompt is displayed.

Setting the Maximum Frame Size

The Maximum Frame Size specifies the maximum data frame size the LEC can send on a Multicast Send VCC, or receive on the Multicast Send VCC or the Multicast Forward VCC. This parameter also specifies the maximum frame size for all LEC Data Direct VCCs.

ELAN services require that the maximum frame size is the same for all LECs belonging to a specific ELAN. All LECs that are members of an ELAN participate in a negotiation process to determine the maximum frame size used on the ELAN. The value set on the LEC should not exceed the Maximum Frame Size set on the LES. You can set the Maximum Frame Size to either 1516 octets, or 4544 octets. The default VNSwitch LEC Maximum Frame Size is 1516 octets. The value you enter is the size requested in the configuration message. The actual value used may be smaller, depending on the results of the negotiation process.

To set the Maximum Frame Size, perform the following steps:

Step	Action
1	At the ATM/ <i>n</i> LEC Config> prompt, enter set max_frame_size .
2	Press Return. The following message is displayed: Maximum Frame Size(1516 or 4544) [1516]?
3	Enter the maximum frame size you want to set. The values you can enter are 1516 or 4544. The default is 1516 octets.

Configuring ATM Physical and Logical Interfaces

Step	Action
4	<p>Press Return. If you are changing the frame size to a greater value than was previously set, the following message and the <code>ATM/n LEC Config></code> prompt are displayed.</p> <pre>Warning: New Maximum Frame Size is larger than current size. Restart VNswitch to use new size.</pre> <p>If you are changing the frame size to a smaller value than was previously set, the Maximum Frame Size is configured as specified, and the <code>ATM/n LEC Config></code> prompt is displayed.</p>
5	<p>If you are changing the frame size to a greater value than was previously set, restart the module according to the instructions provided in Chapter 10.</p>

Setting the Flush Timeout

LECs can transmit unicast frames to the same destination LAN address through either the BUS or a data direct VCC at different times. To avoid the possibility that frames are sent out of order, because of the dual paths, a procedure known as the Flush Message Protocol is used. The flush protocol requires that the LEC transmit an `LE_Flush_Request` on the previously used path, before switching to the second path. An `LE_Flush_Response` must be returned to the sender by the destination client. When the LEC that originated the flush request receives the response, the old path is clear of data and the originating LEC begins using the new path.

You can set the maximum amount of time that the LEC waits to receive an `LE_Flush_Response` after it sends a request. If this period of time is exceeded, the originating LEC begins using the new path. To set the Flush Timeout value, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter <code>set flush_timeout</code> .
2	Press Return. The following message is displayed: <pre>Flush Timeout #seconds [4]?</pre>
3	Enter the maximum number of seconds the LEC waits to receive an <code>LE_Flush_Response</code> after it sends a request. The range of acceptable values is 1 to 4 seconds. The default is 4 seconds.
4	Press Return. The Flush Timeout value is configured as specified, and the <code>ATM/n LEC Config></code> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Setting Path Switching Delay

You can set the period of time after which a LEC assumes a data frame it transmits is either delivered to its destination, or discarded. Flush Timeout is ignored if the Path Switching Delay value is lower than the Flush Timeout.

To set the Path Switching Delay, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter <code>set path_switching_delay</code> .
2	Press Return. The following message is displayed: <code>Path Switch Delay #seconds [6]?</code>
3	Enter the number of seconds after which an LEC assumes a data frame it transmits is either delivered to its destination, or discarded. The range of acceptable values is 1 to 8 seconds. The default is 6 seconds.
4	Press Return. The Path Switching Delay value is configured as specified, and the <code>ATM/n LEC Config></code> prompt is displayed.

Displaying Current ATM LEC Configuration Parameters

You can display specific information about various ATM LEC configuration settings. To display configuration information, perform the following steps:

Step	Action
1	At the <code>ATM/n LEC Config></code> prompt, enter <code>list command</code> , where <i>command</i> is the command you must enter to display the desired information. Refer to Table 5-10 for a list of commands, and for a description of the type of information that is displayed when you enter the command.
2	Press Return. The desired information and a new <code>ATM/n LEC Config></code> prompt is displayed.

Configuring ATM Physical and Logical Interfaces

Example

ATM/14 LEC Config>**list all**

```
Interface                : 14 (Bridge Port 14)
Aging Time(seconds)     : 300
Arp Response Time(seconds) : 1
Config Mode              : Manual
Control timeout(seconds) : 120
Elan Name(Configured)    :
ELan Name(Joined)       : Not available
Flush Timeout(seconds)   : 4
Forward Delay Time(seconds) : 15
Interface State          : Enabled
LEC Status               : Unavailable
LEC MAC Address          : 00-00-F8-4E-9D-0F
LES ATM Address          :
Mac Type                 : Ethernet
Max Frame size           : 1516
Max Retry Count          : 1
Max Unknown Frames Count : 1
Max Unknown Frame Time(seconds) : 1
Path Switch Delay(seconds) : 6
Vcc Timeout Period(seconds) : 1200
```

Configuring ATM Physical and Logical Interfaces

Table 5-10: List Command Options and Descriptions for ATM Logical Configurations

Option	Description
<u>aging-time</u>	Lists the aging time value used to limit the amount of time entries in the LE_ARP cache are retained.
<u>arp_response_time</u>	Displays the maximum amount of time, in seconds, a LEC waits before retrying a LE_ARP Request.
<u>config_mode</u>	Lists whether identification of the LES address is set to automatic or manual mode. Automatic identification is managed by the LECS based on “look-up” tables previously configured by the network administrator. Manual mode enables you to specify the LES ATM address as an LEC configuration parameter.
<u>control_timeout</u>	Displays the timeout value of LECS and LES requests.
<u>elan_name</u>	Lists the name of the ELAN with which the LEC is associated.
<u>flush_timeout</u>	Displays the maximum amount of time that the LEC waits to receive an LE_Flush_Response before starting to use the new channel.
<u>forward_delay_time</u>	Lists the forward delay value used to limit the amount of time non local entries in the LE_ARP cache are retained.
<u>interface_state</u>	Displays the whether the ATM port is management enabled or disabled.
<u>lec_status</u>	Displays the LEC connection status.
<u>lec_mac_addr</u>	Displays the local unicast MAC addresses assigned to the LEC interface.
<u>les_atm_addr</u>	Lists the ATM address of the LES when the config_mode parameter is set to manual.
<u>mac_type</u>	Lists the frame type used by the ATM LEC. The frame type is always Ethernet.

Configuring ATM Physical and Logical Interfaces

Option	Description
<u>max_frame_size</u>	Displays the maximum frame size requested for the LEC on this interface.
<u>max_retry_count</u>	Displays the number of times the LEC attempts to resend an LE_ARP Request when the Expected LE_ARP Response Time is exceeded.
<u>max_unknown_frames_count</u>	Displays the maximum number of frames with a particular unknown destination address that an LEC can send to the BUS within the time period specified by the unknown-frame-time-max parameter.
<u>max_unknown_frame_time</u>	Displays the maximum time within which the maximum number of frames with a particular unknown destination address can be sent to the BUS.
<u>path_switching_delay</u>	Displays the period of time after which an LEC assumes a data frame it transmits is either delivered to its destination, or discarded. Flush Timeout is ignored if the Path Switching Delay value is lower than the Flush Timeout.
<u>vcc_timeout</u>	Lists the maximum amount of time a Data Direct VCC is inactive (no data frame traffic is transmitted or received) before the VCC is released.
<u>all</u>	Displays information associated with all of the above commands.

Configuring ATM Physical and Logical Interfaces

Displaying ATM Physical Interface and Logical Port Information

To display a summary of the physical interface state and the states of all 16 logical ATM ports, enter **list summary** at the ATM Config> prompt.

Example

ATM Config> **list summary**

```
Phy Type           : OC3
Phy Link State     : Running
ATM Address Prefix : 39999990000000000000F84F0A80
Active UNI Version : 3.1
```

LAN Emulation interfaces

Ifc	MAC Address	Status	VSD	Frame Size	Joined ELAN Name
13	0000F84D790E	Disabled	2	1516	
14	0000F84D790F	Operational	3	1516	anna3
15	0000F84D7910	Unavailable	4	1516	
16	0000F84D7911	Unavailable	1	1516	
17	0000F84D7912	Unavailable	1	1516	
18	0000F84D7913	Unavailable	1	1516	
19	0000F84D7914	Unavailable	1	1516	
20	0000F84D7915	Unavailable	1	1516	
21	0000F84D7916	Unavailable	1	1516	
22	0000F84D7917	Unavailable	1	1516	
23	0000F84D7918	Unavailable	1	1516	
24	0000F84D7919	Unavailable	1	1516	
25	0000F84D791A	Unavailable	1	1516	
26	0000F84D791B	Unavailable	1	1516	

Bridge Tunnel interfaces

Ifc	VPI	VCI	MAC Address	MAC Type	Status	VSD
27	0	63	0000F84D791C	FDDI	Down	1
28	0	62	0000F84D791D	Ethernet	Down	1

Configuring the Address Resolution Protocol

When a network device sends an IP packet, the transmitting device must first determine the MAC address associated with the IP address in the packet. The Address Resolution Protocol (ARP) is responsible for determining the MAC-to-IP address mapping. ARP does so by broadcasting an ARP Request that includes the IP address of the destination. The destination node then responds to the request by providing its MAC address to the original node that requested the information. The original node retains the mapped addresses in an ARP cache for later use. This MAC-to-IP address learning process is typically handled automatically by ARP. Under certain conditions, however, you may need to map an IP address to its associated MAC address manually. You may need to do so, for example, if ARP is not running on the destination node.

You can perform the following tasks when configuring ARP:

- Add and change ARP cache entries manually.
- Delete a manually entered ARP entry.
- Manage the time that learned entries are retained.
- Display information about the current configuration.

Adding and Changing ARP Cache Entries Manually

On a VNSwitch module, manually entered ARP cache entries are configured on a per interface basis. A separate ARP cache is, therefore, retained for each interface. You cannot change a learned entry.

To manually map an IP address to a MAC address for an interface, or to change an existing manual entry, perform the following steps:

Step	Action
1	At the Config prompt, enter arp . Example: Config>arp
2	Press Return. The ARP configuration prompt (ARP config>) is displayed.
3	If you want to add a new ARP entry, enter add entry . If you want to change an existing ARP entry, enter change entry .
4	Press Return. The following message is displayed: Interface Number [0]?

Configuring the Address Resolution Protocol

Step	Action
5	Enter the number of the logical interface for which you want to create or change an IP-to-MAC address entry.
6	Press Return. The following message is displayed: IP Address [0.0.0.0]?
7	Enter the IP address of the device for which you want to create or change an entry.
8	Press Return. The following message is displayed: Mac Address []?
9	Enter the MAC address of the device for which you want to create or change an entry. The address must be entered using the format 08002bAABBCC. Hyphens (-) and colons (:) are not allowed.
10	Press Return. The IP-to-MAC address entry is created or modified as specified, and the ARP <code>config></code> prompt is displayed.
11	Restart the module if you want the new configuration settings to take effect.

Deleting a Manually Entered ARP Entry

You can delete ARP cache entries that were previously entered manually. You cannot delete learned entries. Learned entries are automatically aged out of the cache.

To delete an existing IP-to-MAC address ARP entry that was previously added manually, perform the following steps:

Step	Action
1	At the Config prompt, enter arp . Example: <code>Config> arp</code>
2	Press Return. The ARP configuration prompt (<code>ARP config></code>) is displayed.
3	Enter delete entry .
4	Press Return. The following message is displayed: Interface Number [0]?
5	Enter the number of the logical interface for which you want to delete an IP-to-MAC address entry.

Configuring the Address Resolution Protocol

Step	Action
6	Press Return. The following message is displayed: IP Address [0.0.0.0]?
7	Enter the IP address of the device for which you want to delete an IP-to-MAC address entry.
8	Press Return. The IP-to-MAC address entry is deleted, and the ARP <code>conf ig></code> prompt is displayed.
9	Restart the module if you want the new configuration settings to take effect.

Managing the Time That Learned Entries Are Retained

As long as an IP-to-MAC address entry is in the ARP cache, the VNswitch module can transmit data directly to their associated device without rebroadcasting ARP requests. This helps maintain network performance by reducing the need to repeatedly broadcast ARP requests for active communication between two devices.

The VNswitch module retains learned ARP cache entries for a set period of time managed by a Refresh Timer and an Auto-Refresh function. The Refresh Timer specifies the amount of time a learned entry is retained before the module attempts to reverify the presence of the device. The verification process is referred to as Auto-Refresh. Auto-Refresh is the process that attempts to verify that the device associated with a learned ARP entry still exists. The Refresh Timer and Auto-Refresh function do not affect ARP cache entries that are entered manually.

Setting the Refresh Timer

The Refresh Timer specifies the amount of time a learned entry is retained before the module attempts to verify the presence of the device. The default is 20 minutes. If the entry is no longer valid, the entry is deleted from the cache and the VNswitch module must relearn the mapping the next time it attempts to transmit data to the same IP address. If, before the Refresh Time expires, the VNswitch module detects verification of an entry through an ARP Request sent by another device, the entry is retained and the Refresh Timer is restarted. Similarly, if the Refresh Time is about to expire, the VNswitch module attempts to verify the entry by sending its own ARP Request. If the entry is still valid, the entry is retained and the Refresh Timer is restarted.

Configuring the Address Resolution Protocol

You can change the Refresh Time to affect the length of time each learned entry in the cache is retained. To set the Refresh Timer, perform the following steps:

Step	Action
1	At the Config prompt, enter arp . Example: Config>arp
2	Press Return. The ARP configuration prompt (ARP config>) is displayed.
3	Enter set refresh-timer .
4	Press Return. The following message is displayed: timeout (in minutes) [20]? The factory default is 20 minutes. If the factory default was modified, the default is the most recent setting.
5	Enter the number of minutes after which you want ARP to attempt to verify the ARP entry is still valid. The number must be a whole number in the range of 1 to 65535 minutes. A value of zero (0) disables the refresh timer.
6	Press Return. The Refresh Timer is set as specified, and the ARP config> prompt is displayed.
7	Restart the module if you want the new configuration settings to take effect.

Enabling and Disabling Auto-Refresh

Auto-Refresh is the process that attempts to verify that a device still exists when its learned entry in the ARP cache exceeds the Refresh Timer. You can enable and disable Auto-Refresh. When enabled, Auto-Refresh initiates verification shortly before the Refresh Time is exceeded. When disabled, the verification process does not occur and the learned ARP entry is deleted if no active communication is otherwise initiated between the devices before the Refresh Time is exceeded. The factory default is disabled. If the factory default was modified, the default is the most recent setting.

To enable or disable Auto-Refresh, perform the following steps:

Step	Action
1	At the Config prompt, enter arp . Example: Config>arp

Configuring the Address Resolution Protocol

Step	Action
2	Press Return. The ARP configuration prompt (<code>ARP config></code>) is displayed.
3	If you want to disable Auto-Refresh, enter disable auto-refresh . If you want to enable Auto-Refresh, enter enable auto-refresh .
4	Press Return. Auto-Refresh is disabled or enabled as specified.
5	Restart the module if you want the new configuration settings to take effect.

Displaying Information About the Current Configuration

To display information about the current ARP configuration, perform the following steps:

Step	Action
1	At the Config prompt, enter arp . Example: <code>Config>arp</code>
2	Press Return. The ARP configuration prompt (<code>ARP config></code>) is displayed.
3	Enter list option , where <i>option</i> is the command you must enter to display the desired information. Refer to Table 5-7 for a list of commands and for a description of the information that is displayed when you enter the command.
4	Press Return. The desired information and a new <code>ARP config></code> prompt is displayed.

Example

```
ARP config>list all
```

```
ARP configuration:
```

```
Refresh timeout: 8 minutes
```

```
Auto refresh: enabled
```

```
Mac address translation configuration
```

```
Ifc #   Prot #   Protocol -> Mac address
1       0       16.20.48.79 -> 112233445566
1       0       16.20.48.68 -> 112233445566
2       0       16.20.48.89 -> 223344556677
```

Configuring the Address Resolution Protocol

Table 5-11: List Command Options and Descriptions for ARP Configurations

Command	Description
<u>entry</u>	The following information is listed for each ARP cache entry added manually: <ul style="list-style-type: none">• The number of the interface with which the entry is associated• The IP address and MAC address for the entry
<u>config</u>	The following information is listed about the length or time ARP is configured to retain entries added manually: <ul style="list-style-type: none">• The Refresh Timer value in minutes• Whether Auto-Refresh is enabled or disabled
<u>all</u>	All the information described for both the entry and config options are displayed.

Chapter 6

Monitoring Network Interfaces

Overview

Introduction

This chapter provides information about how to monitor a DIGITAL VNswitch module's Ethernet, Fast Ethernet, FDDI, and ATM network interfaces.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Displaying the Interface Number and Type	6-2
Monitoring an Ethernet Interface	6-5
Monitoring a Fast Ethernet Interface	6-10
Monitoring an FDDI Interface	6-15
Monitoring an ATM Interface	6-20
Monitoring the LAN Emulation ARP Database	6-29
Displaying Interface Test Results	6-40
Clearing Interface Counters	6-41
Testing an Interface	6-42
Monitoring the Address Resolution Protocol	6-43
Monitoring ICMP Counters	6-46

Displaying the Interface Number and Type

You can display the number assigned to a logical interface and the interface type (Ethernet, Fast Ethernet, FDDI, or ATM).

To list a VNSwitch module's logical interface numbers and types, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter list all .
2	Press Return. A report that includes general information about the switch is displayed, including the number associated with each type of interface.

Example

The following is an example of the information displayed for an FDDI module. The report is divided into three sections. The first section includes switch identification information and the console baud rate. The interface number is displayed in the third

Displaying the Interface Number and Type

section of the report. The first two sections include general information about the switch and supported protocols. Refer to [Table 6-1](#) for a description of the third section of the report that includes the interface number.

Monitor>**list all**

900EF-MM, 12 Eth 1 FDDI,HW=04f1e1,RO=V1.5.5,#2278,SW=V2.0

Hostname: VN1

Console baud rate: 0

```
Num Name Protocol
0 IP DOD-IP
3 ARP Address Resolution Protocol
11 SNMP Simple Network Management Protocol
23 BRID Adaptive Source Routing Transparent Enhanced Bridge
```

```
Num Name Feature
5 RMON Remote Monitoring
6 Mirr Port Mirror
```

46 Interfaces:

Ifc	Name	MAC/Data-Link	Hardware	State
0	VNbus/0	VNbus-Encapsulation	VNbus	Up
1	Eth/1	Ethernet/IEEE 802.3	Ethernet/802.3	Up
2	Eth/2	Ethernet/IEEE 802.3	Ethernet/802.3	Up
3	Eth/3	Ethernet/IEEE 802.3	Ethernet/802.3	Up
4	Eth/4	Ethernet/IEEE 802.3	Ethernet/802.3	Down
5	Eth/5	Ethernet/IEEE 802.3	Ethernet/802.3	Down
6	Eth/6	Ethernet/IEEE 802.3	Ethernet/802.3	Down
7	Eth/7	Ethernet/IEEE 802.3	Ethernet/802.3	Down
8	Eth/8	Ethernet/IEEE 802.3	Ethernet/802.3	Down
9	Eth/9	Ethernet/IEEE 802.3	Ethernet/802.3	Down
10	Eth/10	Ethernet/IEEE 802.3	Ethernet/802.3	Down
11	Eth/11	Ethernet/IEEE 802.3	Ethernet/802.3	Down
12	Eth/12	Ethernet/IEEE 802.3	Ethernet/802.3	Down
13	FDDI/13	FDDI/IEEE 802.2	FDDI	Up
14	VLAN/14	VLAN Interface	None	Up
15	VLAN/15	VLAN Interface	None	Up
16	VLAN/16	VLAN Interface	None	Up
17	VLAN/17	VLAN Interface	None	Down
...				
45	VLAN/45	VLAN Interface	None	Down

Table 6-1: Description of the Interface Information Displayed

Information Displayed	Description
Ifc	Physical and logical interface number.
Name	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by the interface number.

Displaying the Interface Number and Type

Information Displayed	Description
MAC/Data Link	Type of MAC/data link configured for the interface.
Hardware	Interface hardware type.
State	Current state of the interface. The following states are possible: <ul style="list-style-type: none">• <code>Up</code> — The interface is connected and operational.• <code>Down</code> — The interface is not operational and failed self-test.• <code>Disabled</code> — The interface is either temporarily or permanently disabled.• <code>Testing</code> — The interface is in the process of completing a self-test.• <code>Not present</code> — Power-up (restart) diagnostics detected a fault on the interface.

Monitoring an Ethernet Interface

This section provides information about how to monitor an Ethernet interface. Monitoring includes the ability to display specific information about operational states, activity counters, and various interface configuration settings.

To monitor the interface, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter interface statistics , followed by the number of the interface you want to monitor. The Ethernet interfaces are numbered 1 through 12 on a VNswitch 900EX, 900EF, or 900EA module. The interfaces are numbered 1 through 24 on a VNswitch 900EE module.
2	Press Return. Information about operational states, activity counters, and various interface configuration settings is displayed.

Example

The following is an example of the information displayed. Refer to [Table 6-2](#) for a description of the information.

Monitoring an Ethernet Interface

Monitor>interface statistics 1

```

Ifc Ifc' Name          Self-Test      Self-Test      Maintenance
   Passed          Failed          Failed
 1 1 Eth/1             0              41799          0
 Ethernet/IEEE 802.3 MAC/data-link on Ethernet/802.3 interface
Physical address      0000F864B802   Network Speed   10 Mb/s
PROM address         0000F864B802   Duplex mode     HALF
Shared hardware errors 0
Input statistics:
  failed, frame too long      0   failed, FCS error      0
  failed, alignment error     0   failed, FIFO overrun   0
  runts                       0   packets missed         0
  internal hardware errors    0
Output statistics:
  deferred transmission       0   single collisions      0
  multiple collisions         0   total collisions       0
  excess collisions           0   failed, FIFO underrun  0
  failed, carrier sense err   0   SQE test err           0
  late collision              0   internal hardware errors 0
MAU Status:
connector            RJ45 s
type                 10BASE-T HALF DUPLEX
state                OPERATIONAL
media available state NOT AVAILABLE
media avail state exit 0
jabber state         NO JABBER
jabber state enters  0
false carriers

```

Table 6-2: Description of the Monitoring Information Displayed for Ethernet

Information Displayed	Description
Ifc	Physical and logical interface number.
Ifc'	Not applicable on VNswitch module.
Name	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
Self-Test Passed	Number of times the interface passed self-test since it was last restarted.
Self-Test Failed	Number of times the interface failed diagnostic self-test since it was last restarted.

Monitoring an Ethernet Interface

Information Displayed	Description
Maintenance Failed	Number of times the interface failed maintenance tests since it was last restarted.
Physical address	Ethernet address of the interface currently in use. This may be the PROM address or an address overwritten by another protocol.
PROM address	Permanent unique Ethernet address in the PROM for this Ethernet interface.
Shared hardware errors	Number of internal hardware errors that do not fall into any other category of counters listed in this table. These errors are related to hardware shared by multiple interfaces, but are not specific to any one interface.
Network speed	Speed of the network interface.
Duplex mode	Duplex setting for which the interface is set (half-duplex or full-duplex mode).
Input statistics: failed, frame too long	Number of times the interface receives a frame that is larger than the maximum size of an Ethernet frame (1518 bytes).
Input statistics: packets missed	Number of times the interface attempted to receive a packet, but the local packet buffer is full. This error indicates that the network has more traffic than the interface can handle.
Input statistics: internal hardware errors	Number of internal hardware errors related to packet reception that do not fall into any other category of counters listed in this table. These errors are related to a specific interface.
Input statistics: failed, FIFO over-run	Number of times the Ethernet chipset is unable to store bytes in the local packet buffer as fast as they come off the wire.
Input statistics: runts	Number of frames, shorter than 64 bytes, received by the interface since the last restart.
Output statistics: deferred transmission	Number of times the carrier sense mechanism detects line activity causing the interface to defer transmission.

Monitoring an Ethernet Interface

Information Displayed	Description
Output statistics: multiple collisions	Number of packets involved in multiple collisions before they are successfully transmitted.
Output statistics: excess collisions:	Number of packets for which transmission fails due to 16 successive collisions on a given interface. This error indicates a high volume of network traffic or hardware problems with the network.
Output statistics: failed, carrier sense err	Number of times the carrier sense condition was lost or never asserted when attempting to transmit a packet on a given interface. This error indicates a problem between the interface and its Ethernet transceiver.
Output statistics: late collision	Number of times that a collision is detected on a given interface later than 512 bit times into the transmission of a packet. This error indicates that an interface on the network failed to defer, or that the network is not properly configured.
Output statistics: internal hardware errors	Number of internal hardware errors related to packet transmission that do not fall into any other category of counters listed in this table. These errors are related to a specific interface.
Output statistics: single collisions	Number packets that have experienced exactly one collision and are then successfully transmitted.
Output statistics: total collisions	Sum of single collisions plus multiple collisions.
Output statistics: failed, FIFO underrun	Number of times a packet transmission fails due to the inability of the interface to retrieve packets from the local packet buffer fast enough to transmit them on the network.
Output statistics: SQE test error	Number of times the interface sends a packet but detects that the transceiver has no heartbeat. The packet is treated as successfully transmitted because some transceivers do not generate heartbeats.

Monitoring an Ethernet Interface

Information Displayed	Description
Output statistics: jabber	Number of times the interface attempted to transmit for an excessive period of time (between 20 and 150 milliseconds).

Monitoring a Fast Ethernet Interface

This section provides information about how to monitor a Fast Ethernet interface. Monitoring includes the ability to display specific information about operational states, activity counters, and various interface configuration settings.

To monitor the interface, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter interface statistics , followed by the number of the interface you want to monitor. The two Fast Ethernet interfaces on a VNswitch 900EX module are numbered 13 and 14.
2	Press Return. Information about operational states, activity counters, and various interface configuration settings is displayed.

Monitoring a Fast Ethernet Interface

Example

The following is an example of the information displayed. Refer to [Table 6-3](#) for a description of the information.

Monitor>interface statistics 13

```

Ifc Ifc' Name          Self-Test      Self-Test      Maintenance
                   Passed        Failed         Failed
 13 13 E100/13          1              2              0
 Ethernet/IEEE 802.3u MAC/data-link on 100 Mb/s Ethernet interface
Physical address      0000F864B80E   Network Speed      100 Mb/s
PROM address          0000F864B80E   Duplex mode         HALF
Input statistics:
  failed, frame too long      0   failed, FCS error      0
  failed, alignment error     0   failed, FIFO overrun   0
  runts                        0
Output statistics:
  deferred transmissions      3   total collisions      0
  failed, excess collisions   0   late collisions        0
MAU Status:
  connector                   RJ45 s
  type                         100BASE-TX HALF DUPLEX
  state                        OPERATIONAL
  media available state        AVAILABLE
  media avail state exits     0
  jabber state                 NO JABBER
  jabber state enters         2
  false carriers              0

  auto-negotiation state      COMPLETE
  remote signaling            NOT DETECTED
  advertised capabilities     100BASE-TX FULL DUPLEX
                              100BASE-TX HALF DUPLEX
                              10BASE-T FULL DUPLEX
                              10BASE-T HALF DUPLEX

  remote capabilities         100BASE-TX HALF DUPLEX

```

Table 6-3: Description of the Monitoring Information Displayed for Fast Ethernet

Information Displayed	Description
Ifc	Physical and logical interface number.
Ifc'	Not applicable on VNswitch module.

Monitoring a Fast Ethernet Interface

Information Displayed	Description
Name	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
CSR	The value in this field is not applicable to the VNswitch module.
Vec	The value in this field is not applicable to the VNswitch module.
Self-Test Passed	Number of times the interface passed self-test since it was last restarted.
Self-Test Failed	Number of times the interface failed diagnostic self-test since it was last restarted.
Maintenance Failed	Number of times the interface failed maintenance tests since it was last restarted.
Physical address	Ethernet address of the interface currently in use. This may be the PROM address or an address overwritten by another protocol.
PROM address	Permanent unique Ethernet address in the PROM for this Fast Ethernet interface.
Network speed	Current operating speed of the network interface.
Duplex mode	Duplex setting for which the interface is set (half-duplex or full-duplex mode).
MAU status: type	Current operational type and mode (10Base-T full-duplex mode, 10Base-T half-duplex, 100 base-TX full-duplex mode, or 100 base-TX half-duplex mode) of the Medium Attachment Unit (MAU) as determined by either the autonegotiation function, or if autonegotiation is disabled or is not implemented, the default speed and duplex mode settings.
MAU status: state	Current operational status of the MAU associated with the interface. Possible states are: Operational, Not present, Unknown, and Other.

Monitoring a Fast Ethernet Interface

Information Displayed	Description
MAU status: media available state	Status of the MAU's connection to the physical medium. Possible states are: Available, Not available, Remote fault, Unknown, and Other.
MAU status: media available state exists	Number of times the MAU is in the media available state, and then changes to another state.
MAU status: connector	Type of physical connector used by the MAU. Possible connector types are: RJ45s and fiber SC.
MAU status: jabber state	Jabbering is the state to which the MAU transitions if it transmits for an excessive period of time (26 milliseconds). While in jabber state mode, the MAU disables any further transmission onto the medium until the interface stops transmitting to the MAU for approximately 750 milliseconds. (Applicable to 10Base-T mode only) Possible states are Jabbering and Not jabbering.
MAU status: jabber state enters	Number of times the MAU detected a jabber condition. (Refer to MAU status: jabber state.)
MAU status: false carriers	Number of times the carrier sense is asserted without a J/K symbol detection. (Applicable to 100Base-TX and -FX mode only.)
Auto-negotiation state	Current status of the auto-negotiation process. Possible states are: Configuring, Complete, Disabled, and Parallel detect failed.
Remote signalling	This field indicates whether the remote device is using auto-negotiation signaling.
Advertised capabilities	This field displays the list of supported modes of operation (100Base-TX full-duplex, 100Base-TX half-duplex, and so on) that the MAU is advertising to the remote device.
Remote capabilities	This field displays the list of advertised modes of operation (100Base-TX full-duplex, 100Base-TX half-duplex, and so on) received from the remote device.

Monitoring a Fast Ethernet Interface

Information Displayed	Description
Input statistics: failed, frame too long	Number of times the interface receives a frame that is larger than the maximum size of an Ethernet frame (1518 bytes).
Input statistics: failed, alignment error	Number of times the interface receives a packet with a length that is not a multiple of eight bits.
Input statistics: runts	Number of frames, shorter than 64 bytes, received by the interface since the last restart.
Input statistics: failed, FCS error	Number of packets received on a given interface that do not pass the Frame Check Sequence (FCS) test.
Input statistics: failed, FIFO overrun	Number of times the Fast Ethernet chipset is unable to store bytes in the local packet buffer as fast as they come off the wire.
Output statistics: deferred transmission	Number of times the carrier sense mechanism detects line activity causing the interface to defer transmission.
Output statistics: failed, excess collisions	Number of packets for which transmission fails due to 16 successive collisions on a given interface. This error indicates a high volume of network traffic or hardware problems with the network.
Output statistics: total collisions	Total number of collisions a packet incurs.
Output statistics: late collisions	Number of times that a collision is detected on a given interface later than 512 bit times into the transmission of a packet. This error indicates that an interface on the network failed to defer, or that the network is not properly configured.

Monitoring an FDDI Interface

This section provides information about how to monitor the FDDI interface. Monitoring includes the ability to display specific information about operational states, activity counters, and various interface configuration settings.

To monitor the interface, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter interface statistics , followed by the number of the interface you want to monitor. The number is 13 for an FDDI interface on the VNswitch 900EF.
2	Press Return. Information about operational states, activity counters, and various interface configuration settings is displayed.

Example

The following is an example of the information displayed. Refer to [Table 6-4](#) for a description of the information.

Monitor>**interface statistics 13**

```

Ifc Ifc' Name           Self-Test      Self-Test      Maintenance
      Passed      Failed         Failed
13 13 FDDI/13           1              1              0
  FDDI/IEEE 802.2 MAC/data-link on FDDI interface
UNA: 0000F84A1A8E -> MLA: 0000F84E180E -> DNA: 0000F847F5AE
Policy: reject A-A B-B M-M                               Ring Inits: 4
Bypass: False      Purger State: purgerOff      FDX State: fdxIdle
Station: Das       ECM: In          CFM: c_wrap_b      RMT: Ring_Op
Interface 13 A/M: (PHY A, connecting)
  LEM Rejects: 0          LCT Fail Ct: 0          LEM Ct : 0
  Alarm       : 10^-8      Cutoff       : 10^-8      Estimate: 10^-15
Interface 13 B/S: (PHY B, active)
  LEM Rejects: 0          LCT Fail Ct: 0          LEM Ct : 0
  Alarm       : 10^-8      Cutoff       : 10^-8      Estimate: 10^-15
T_Neg : 7.9872ms      TVX : 2.51904ms      T_Max : 167.77216ms
T_Req : 7.9872ms      T_Notify : 10s      Ring Latency : 0.71ms
Beacon Rcv: 0        Beacon Xmt: 0        Trace Rcv: 0        Trace Xmt: 0
Frame Ct: 1394640    Frame Errors: 0      Frames Lost: 0

```

Monitoring an FDDI Interface

Table 6-4: Description of the Monitoring Information Displayed for FDDI

Information Displayed	Description
Ifc	Physical and logical interface number.
Ifc'	Not applicable on VNswitch module.
Name	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
CSR	This field is not active for the FDDI interface.
Vec	This field is not active for the FDDI interface.
Self-Test Passed	Number of times the interface passed self-test since it was last restarted.
Self-Test Failed	Number of times the interface failed diagnostic self-test since it was last restarted.
Maintenance Failed	This field is not active for the FDDI interface.
UNA:	Address of the upstream neighbor, displayed in canonical format.
MLA:	VNswitch module MAC address of the FDDI interface, displayed in canonical format.
DNA:	Address of the downstream neighbor, displayed in canonical format.
Policy:	Shows the PHY port type connection combinations not allowed (rejected) on the interface's PHY ports. For example, <code>reject A-A B-B M-M</code> , indicates a link between PHY A port on one station and PHY A port on a second station is not allowed. PHY B port-to-PHY B port and PHY M port-to-PHY M port are also not allowed.
Bypass:	A value of <code>True</code> indicates an optical bypass relay (OBR) is present. <code>False</code> indicates an OBR is not present.

Monitoring an FDDI Interface

Information Displayed	Description
Purger State:	Indicates the state of the Ring Purger. Possible states include: Purger Off, Candidate, Nonpurger, and Purger. Refer to Chapter 5 for information about how to set the Ring Purger.
FDX State:	Indicates whether full-duplex mode is active or idle on the FDDI interface. Possible states include: fdx Idle, fdx Operation, fdx Request, and fdx Confirm. A value of fdx Operation indicates the interface is set to full-duplex mode and a full-duplex circuit was established following auto-negotiation. A value of fdx Idle indicates the interface is either set to half-duplex mode, or set to full-duplex mode but a half-duplex circuit was established following auto-negotiation. The fdx Request and fdx Confirm states are transitional states. Refer to Chapter 5 for information about how to set the interface for full-duplex or half-duplex mode.
Station:	Displays the station type to which the interface is set. Possible values include DAS, SAC, or SAS. Refer to Chapter 5 for information about how to set the Station Type.
ECM:	Displays the FDDI Entity Connection Management (ECM) state. ECM controls and manages the optical bypass. Possible states include: In, Out, Trace, Leave, Path-Test, Insert, Deinsert, and Check.
CFM:	Displays the FDDI Configuration Management State (CFM) state. Possible states include: Isolated, Local-a, Local-b, Local-ab, Local-s, Wrap-a, Wrap-b, Wrap-ab, Wrap-s, C-wrap-a, C-wrap-b, C-wrap-s, and Thru.
RMT:	Identifies the Ring Management (RMT) software monitoring the state of the ring. Possible values include: Isolated, Non-Op, Ring-Op, Non-Op-Dup, Ring-Op-Dup, Directed, and Trace.
Interface <i>n</i> A/M:	Indicates the state of the PHY port. Possible states include Disabled, Connecting, Stand By, and Active.

Monitoring an FDDI Interface

Information Displayed	Description
LEM Rejects:	Displays the number of times the Link Error Monitor detects a link was rejected.
LCT Fail Ct:	Displays the number of times the link confidence test failed since the last restart.
LEM Ct:	Displays the total number of link error events that occurred since the last restart.
Alarm:	Displays the currently configured Set PHY LER-Alarm value. Refer to Chapter 5 for information about how to set the Link Error Rate Alarm.
Cutoff:	Displays the currently configured Set PHY LER-Cutoff value. Refer to Chapter 5 for information about how to set the Link Error Rate Cutoff.
Estimate:	Displays the estimated bit error rate based on link monitoring.
T_Neg:	Displays the token rotation time.
TVX:	Displays the Valid Transmission Timer setting. This value is the maximum amount of time the interface waits to receive a <i>valid</i> data frame. Refer to Chapter 5 for information about how to set the Valid Transmission Timer.
T_Max:	Displays the Maximum Token Rotation Time (tmax-lower-bound). Refer to Chapter 5 for information about how to set this value.
T_Req:	Displays the Requested Token Rotation Time (t-req). Refer to Chapter 5 for information about how to set this value.
T_Notify:	Displays how often the interface generates Neighbor Information Frames (NIF) to neighbor nodes. This parameter is not user-configurable.
Ring Latency:	Amount of time it takes for data to make one rotation around the ring.

Monitoring an FDDI Interface

Information Displayed	Description
Beacon Rcv:	Displays the total number of times the interface received its own beacon frames and those of other stations since the last restart.
Beacon Xmt:	Displays the number of times the interface entered the Beacon Transmit state since the last restart.
Trace Rcv:	Displays the number of trace frames received.
Trace Xmt:	Displays the number of trace frames transmitted.
Frame Ct:	Displays the total number of frames passed on the ring since the last restart. Note: Disable the Ring Purger on all devices on the ring to get the most accurate ring utilization frame count. If the Ring Purger is on, void frames transmitted by the MAC are also counted, artificially increasing the frame count.
Frame Errors:	Displays the total number of Cyclic Redundancy Check (CRC) errors that occur on the ring since the last restart.
Frames Lost:	Displays the total number of frames lost on the ring since the last restart.

Monitoring an ATM Interface

This section provides information about how to display the results of interface maintenance and self-tests, and the MAC data link type associated with the interface. This section also provides information about how to monitor both an ATM physical interface and an ATM logical interface. Monitoring includes the ability to display specific information about operational states, activity counters, and various interface configuration settings.

Displaying Interface Test Results and MAC Type

To display interface test results and MAC type, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter interface statistics , followed by the number of the desired interface.
2	Press Return. The test report and MAC type are displayed.

Example

Monitor>**interface statistics 14**

```
          Self-Test      Self-Test      Maintenance
Ifc Ifc' Name  Passed      Failed         Failed
14 14  ALEC/14   0              1              0
      ATM LAN Emulation C MAC/data-link on ATM interface
```

Monitoring the Physical Interface

ATM physical interface numbers vary depending on the switch model and configuration. See [Table 1](#) for a list of possible interface numbers.

To monitor a physical interface, perform the following steps:Example

Step	Action
1	At the Monitor prompt (Monitor>), enter ATM physical to display the ATM Physical> prompt.
2	Press Return. The ATM Physical> prompt is displayed. Physical ATM interface 1 maps to logical ports 13 through 28.

Monitoring an ATM Interface

Step	Action
3	Enter list option , where <i>option</i> is the command you must enter to display the desired information. Refer to Table 6-5 for a list of commands.
4	Press Return. Depending on the way the interface was set in the DIGITAL clearVISN network management product, information about either the front panel interface or the interface to the backplane is displayed.

```
ATM Physical>list mib-statistics current
Logical Port Numbers      : 13 - 28
sect status               : 0x06
sect ess                  : 869
sect sess                 : 869
sect sefs                 : 869
sect cvs                  : 869
line status               : 0x02
line ess                  : 0
line sess                 : 0
line cvs                  : 0
line uass                 : 869
path width                : sonet path sts3
path status               : 0x0C
path ess                  : 0
path sess                 : 0
path cvs                  : 0
path uass                 : 869
oocd                      : 871138
lcd_alarm                 : lcd failure
ATM Physical>
```

Monitoring an ATM Interface

Table 6-5: ATM OC3 Physical Interface List Command Options

Command	Description
<u>all</u>	Lists all physical interface information.
<u>atom3-errors</u>	Lists status and counters of the ATM SAR function.
<u>medium</u>	Provides information about the physical medium.
<u>mib-statistics</u> <i>mib-option</i> , where <i>mib-option</i> is one of the command options described in Table 6-7 .	
<u>modpmd-errors</u>	Lists status and counters of the physical medium.
<u>status</u>	Provides detailed status information on ATM links, including counters for net up count, net down count, data direct and bridge tunnel count, and the number of times the PHY went into loopback because the PHY was down.
<u>uni-version</u>	Displays information about the User Network Interface versions being used.

Table 6-6: ATM DS1 and DS3 Physical Interface List Command Options

Command	Description
<u>all</u>	Lists all physical interface information.
<u>atom3-errors</u>	Lists status and counters of the ATM SAR function.
<u>config</u>	Displays information about the physical medium
<u>mib-statistics</u> <i>mib-option</i> , where <i>mib-option</i> is one of the command options described in the Table 6-7 .	
<u>modpmd-errors</u>	Lists status and counters of the physical medium.

Monitoring an ATM Interface

Command	Description
<u>status</u>	Provides detailed status information on ATM links, including counters for net up count, net down count, data direct and bridge tunnel count, and the number of times the PHY went into loopback because the PHY was down.
<u>uni-version</u>	Displays information about the User Network Interface versions being used.

Monitoring an ATM Interface

ATM Physical Interface MIB Statistics

The ATM physical interface maintains a variety of activity counters. Each counter records the total number of occurrences of each activity per 15-minute period. Each 15-minute period is referred to as an interval. The VNswitch maintains a record of counts for the last 96 intervals which is equivalent to 24 hours. You can display counters for the current interval, for all 96 intervals (the last 24 hours), or for a selected interval.

Table 6-7: ATM Physical Interface MIB-Statistics Command Options and Descriptions

Command	Description
current	Displays counter information for the current interval (15-minute period).
total	Displays counter information for the last 24 hours (all 96 15-minute intervals).
interval	After you press Return, the following prompt is displayed: Interval [1]? Enter a number from 1 through 96, where the number specifies the interval for which you want counter information displayed.

Monitoring the Physical Interface and ATM Logical Ports

To monitor the physical interface state and the states of all 16 logical ATM ports, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter atm .
2	Press Return. The ATM> prompt is displayed.
3	Enter list summary
4	Press Return. Information about the physical interface state and the states of the logical ATM ports is displayed.

Monitoring an ATM Interface

Example

```
EA-13 ATM>list summary
Phy Type           : OC3
Phy Link State     : Running
ATM Address Prefix : 399999000000000000F84F0A80
Active UNI Version : 3.1
```

LAN Emulation interfaces

Ifc	MAC Address	Status	VSD	Frame Size	Joined ELAN Name
13	0000F84D790E	Disabled	2	1516	
14	0000F84D790F	Operational	3	1516	anna3
15	0000F84D7910	Unavailable	4	1516	
16	0000F84D7911	Unavailable	1	1516	
17	0000F84D7912	Unavailable	1	1516	
18	0000F84D7913	Unavailable	1	1516	
19	0000F84D7914	Unavailable	1	1516	
20	0000F84D7915	Unavailable	1	1516	
21	0000F84D7916	Unavailable	1	1516	
22	0000F84D7917	Unavailable	1	1516	
23	0000F84D7918	Unavailable	1	1516	
24	0000F84D7919	Unavailable	1	1516	
25	0000F84D791A	Unavailable	1	1516	
26	0000F84D791B	Unavailable	1	1516	

Bridge Tunnel interfaces

Ifc	VPI	VCI	MAC Address	MAC Type	Status	VSD
27	0	63	0000F84D791C	FDDI	Down	1
28	0	62	0000F84D791D	Ethernet	Down	1

Monitoring a Logical Interface

To monitor a logical interface, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter interface and the number of the interface you want to monitor.
2	Press Return. The ATM/n LEC> or ATM/n BT> prompt is displayed.

Monitoring an ATM Interface

Step	Action
3	Enter list option , where <i>option</i> is the command you must enter to display the desired information. Refer to Table 6-8 for a list of commands, and for a description of the type of information that is displayed when you enter the command.
4	Press Return. Information about the selected option and the ATM/ <i>n</i> LEC> or ATM/ <i>n</i> BT> prompt is displayed.

Monitoring an ATM Interface

Example

The following is an example of the information displayed when you enter the **list all** command option for interface 13, when the interface is configured as a LAN Emulation Client (LEC)

ATM/14 LEC> **list all**:

```
Interface                : 14 (Bridge Port 14)
Lec Status               : Operational
Elan Name(Configured)   :
Elan Name(Joined)       : anna3
Atm address              : 3999990000000000000F84F0A80.0000F84D790E.00
LEC Control Frames Sent : 14
LEC Control Frames Received : 0
LEC Illegal Control Frames Received : 0
LE Arp Request Sent     : 3
LE Arp Request Received : 4
LE Arp Response Sent    : 2
LE Arp Response Received : 3
LECS Atm address        : 3999990000000000000F84F0A80.0000F84F0A90.00
```

Ifc	VPI	VCI	Type	S/PVC	ATM Address
13	0	5	Signaling	PVC	
13	0	16	ILMI	PVC	
14	0	107	LE Control	SVC	3999990000000000000F84F0A80.0000F84F0A9E.00
14	0	108	LE Multicast	SVC	3999990000000000000F84F0A80.0000F84F0A9F.00
14	0	130	LE Control	SVC	3999990000000000000F84F0A80.0000F84F0A9E.00
14	0	131	LE Multicast	SVC	3999990000000000000F84F0A80.0000F84F0A9F.00
14	0	134	LE DataDirect	SVC	3999990000000000000F84F0A80.0000F88E230E.00
14	0	135	LE DataDirect	SVC	3999990000000000000F84F0A80.0000F88E230E.00

If the interface is configured as an ATM bridge tunnel, the **list all** command displays information such as the following:

ATM/28 BT>**list all**

```
Interface                : 28 (Bridge Port 28)
Interface Status         : Down
```

Ifc	VPI	VCI	Type	S/PVC	ATM Address
28	0	62		PVC	

Monitoring an ATM Interface

Table 6-8: ATM Logical Interface List Command Options

Command	Description
<u>atm_address</u>	Displays the ATM address for this logical VNswitch port.
<u>counters</u>	Displays Statistics on the following types of frames: LEC control frames sent LEC control frames received LEC illegal frames received
<u>database</u>	Displays entries in the LE-ARP database associated with this LEC port.
<u>elan_name</u>	Displays the emulated LAN identification.
<u>le_arp_counters</u>	Displays LAN Emulation ARP information for: LE Arp Request Sent LE Arp Request Received LE Arp Response Sent LE Arp Response Received
<u>status</u>	Displays LAN emulation client (LEC) or bridge tunnel (BT) status (enabled or disabled)
<u>vc</u>	Displays all virtual circuits on this interface
<u>all</u>	Displays all applicable information from this table

Monitoring the LAN Emulation ARP Database

To monitor the LAN Emulation ARP Database, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter <code>atm</code> .
2	Press Return. The <code>ATM></code> prompt is displayed.
3	Enter <code>list database learp option</code> , where <i>option</i> is the command you must enter to display the desired information. Refer to Table 6-9 for a list of commands, and for a description of the type of information that is displayed when you enter the command.
4	Press Return. Information about the selected option and the <code>ATM></code> prompt is displayed.

Table 6-9: ATM Physical Interface List Database LEARP Command Options and Descriptions

Command	Description
<code>addr</code>	Displays the entry for the specified MAC address.
<code>interface</code>	Displays the entries associated with the specified interface number.
<code>range</code>	Displays the entries whose MAC address falls within the specified range of addresses.

Monitoring an ATM Interface

Example

```
ATM>list database learp addr 00-00-F8-8E-23-00
```

MAC Address	Ifc	VPI/VCI	Type	Rmt	ATM Address
00-00-F8-8E-23-00	14	0/135	LeArp	Y	39999900000000000000F84F0A80.0000F88E230E.00

Example

```
ATM>list database learp interface 14
```

MAC Address	Ifc	VPI/VCI	Type	Rmt	ATM Address
00-00-F8-8E-23-00	14	0/135	LeArp	Y	39999900000000000000F84F0A80.0000F88E230E.00

Example

```
ATM>list database learp range
```

```
First MAC address [00-00-00-00-00-00]?
```

```
Last MAC address [FF-FF-FF-FF-FF-FF]?
```

MAC Address	Ifc	VPI/VCI	Type	Rmt	ATM Address
00-00-F8-8E-23-00	14	0/135	LeArp	Y	39999900000000000000F84F0A80.0000F88E230E.00

Monitoring Packet Statistics and Error Counts

You can monitor the following information about one or more interfaces:

- Packet buffers
- Error counts
- Input and output queues
- Packet statistics

Displaying Packet Buffer Data

You can display a report that displays packet buffer data for one or all interfaces. The buffer size for each interface is built dynamically and is different for each type of interface. Global system buffers are of one size and are equal to the maximum buffer size of all interface buffers.

Note

The information provided by the packet buffer report reflects activity occurring on the VNswitch module's management processor only. It does not include data about buffers in the VNswitch module's forwarding subsystem.

To display the report, perform the following steps:

Step	Action
1	If you want to view packet buffer data for all interfaces, enter buffer at the Monitor prompt (<code>Monitor></code>). If you want to view packet buffer data for a selected interface, enter buffer interface# at the Monitor prompt (<code>Monitor></code>), where interface# is the number of the interface for which you want to display packet buffer data.
2	Press Return. The packet buffer report is displayed, and the Monitor prompt is displayed.

Monitoring Packet Statistics and Error Counts

Example

The following is an example of the information displayed for a single (FDDI) interface on a VNswitch 900EF. If you chose to display data for all interfaces, a similar report is displayed, containing the same information for each interface. Refer to [Table 6-10](#) for a description of the information.

Monitor>**buffer 13**

Nt	Interface	Input Buffers				Buffer sizes				Bytes	
		Req	Alloc	Low	Curr	Hdr	Wrap	Data	Trail	Total	Alloc
13	FDDI/13	16	16	4	16	28	42	4478	4	4552	72832

Table 6-10: Description of the Buffer Information Displayed

Information Displayed	Description
Nt	Physical and logical interface number.
Interface	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
Input Buffers Req	Number of buffers requested.
Input Buffers Alloc	Number of buffers allocated.
Input Buffers Low	Low water mark (for flow control).
Input Buffers Curr	Current number of buffers on this interface. The value is 0 if the interface is down. The packet is eligible for flow control if the value of Curr is below Low when a packet is received. Refer to the Input and Output Queues section for conditional information.
Buffer Sizes Hdr	Sum of the maximum hardware, MAC, and data link headers.
Buffer Sizes Wrap	Allowance given for MAC, LLC, or network layer headers due to protocol wrapping.
Buffer Sizes Data	Maximum data link layer packet size.
Buffer Sizes Trail	Sum of the largest MAC and hardware trailers.
Buffer Sizes Total	Overall size of each packet buffer.

Monitoring Packet Statistics and Error Counts

Information Displayed	Description
Buffer Sizes Bytes Alloc	Amount of buffer memory for the interface. This value is determined by multiplying the values of <code>Alloc X Total</code> .

Monitoring Packet Statistics and Error Counts

Packet Error Statistics

You can display a report that displays packet error statistics for one or all interfaces. To display the report, perform the following steps:

Step	Action
1	If you want to view error statistics for all interfaces, enter error at the Monitor prompt (Monitor>). If you want to view error statistics for a selected interface, enter error interface# at the Monitor prompt (Monitor>), where interface# is the number of the interface for which you want to display error statistics.
2	Press Return. The error statistics report is displayed, and the Monitor prompt is displayed.

Example

The following is an example of the information displayed for a single (FDDI) interface. If you chose to display data for all interfaces, a similar report is displayed, containing the same information for each interface. Refer to [Table 6-11](#) for a description of the information.

```
Monitor>error 13
      Input      Input      Input      Input      Output      Output
Ifc Name Discards Errors  Unk Proto  Flow Drop  Discards Errors
13  FDDI/13      0      0      0      0      1      0
```

Monitoring Packet Statistics and Error Counts

Table 6-11: Description of the Packet Error Information Displayed

Information Displayed	Description
Ifc	Physical and logical interface number.
Name	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
Input Discards	Number of packets received unsuccessfully.
Input Errors	Number of packets that were defective at the data link.
Input Unk Proto	Number of packets received for an unknown protocol.
Input Flow Drop	Number of packets received that are flow controlled on input.
Output Discards	Number of packets the switch chose to discard rather than transmit due to flow control.
Output Errors	Number of output errors including, for example, attempts to send data over a network that is down or that went down during transmission.

Note

The sum of the Output Discards is not the same as Input Flow Drops over all networks. Output Discards may indicate locally originated packets.

Monitoring Packet Statistics and Error Counts

Input and Output Queues

You can display a report that displays information about the length of input and output queues for one or all interfaces.

Note

The information provided by the input and output queue report reflects activity occurring on the VNswitch module's management processor only. It does not include data about queues in the VNswitch module's forwarding subsystem.

To display the input and output queue report, perform the following steps:

Step	Action
1	If you want to view input and output queues for all interfaces, enter queue at the Monitor prompt (Monitor>). If you want to view input and output queues for a selected interface, enter queue interface# at the Monitor prompt (Monitor>), where interface# is the number of the interface for which you want to display queues.
2	Press Return. The queue report is displayed, and the Monitor prompt is displayed.

Example

The following is an example of the information displayed for a single (FDDI) interface. If you chose to display data for all interfaces, a similar report is displayed, containing the same information for each interface. Refer to [Table 6-12](#) for a description of the information.

Monitor>**queue 13**

Nt	Interface	Input Queue			Output Queue	
		Alloc	Low	Curr	Fair	Curr
13	FDDI/13	16	4	16	8	0

Table 6-12: Description of the Input and Output Queue Information Displayed

Information Displayed	Description
Ifc	Physical and logical interface number.
Interface	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
Input Queue Alloc	Number of buffers allocated to the interface.
Input Queue Low	Low water mark for flow control on this interface.
Input Queue Curr	Current number of buffers on this interface. The value is 0 if the interface is disabled or down.
Output Queue Fair	Fair level for the length of the output queue on the interface.
Output Queue Curr	Number of packets currently waiting to be transmitted on the interface. For locally originated packets, the eligibility discard depends on the global low water mark value you can monitor by displaying the VNswitch module's memory as described in Chapter 4 .

Flow Control and Dropped Buffers

The VNswitch module attempts to keep at least the `Low` value packets available for receiving over an interface. If a packet is received and the value of `Curr` is less than `Low`, then the packet is subject to flow control. If a buffer subject to flow control is to be queued on the interface and the `Curr` value is greater than `Fair`, the buffer is dropped instead of queued. The dropped buffer is displayed in the `Output Discards` column of the packet error statistics report. (Refer to the [Packet Error Statistics](#) section.) It also generates ELS event GW.036 or GW.057.

Interpreting Current Buffer Values During Packet Forwarding

Due to the scheduling algorithms of the VNswitch module, the dynamic values of `Curr` (especially `Input Queue Curr`) may not be fully representative of typical values. The Monitor code runs when the input queues are drained. Therefore, `Input Queue Curr` is nonzero when those packets are waiting on slow transmit queues.

Monitoring Packet Statistics and Error Counts

Note

The information provided by the input and output queue report reflects activity occurring on the VNswitch module's management processor only. It does not include data about queues in the VNswitch module's forwarding subsystem.

Packet Statistics

You can display a report that displays packet information for one or all interfaces. To display the report, perform the following steps:

Step	Action
1	If you want to view packet statistics for all interfaces, enter statistics at the Monitor prompt (Monitor>). If you want to view packet statistics for a selected interface, enter statistics interface# at the Monitor prompt (Monitor>), where interface# is the number of the interface for which you want to display packet statistics.
2	Press Return. The queue report is displayed, and the Monitor prompt is displayed.

Example

The following is an example of the information displayed for a single (FDDI) interface. If you chose to display data for all interfaces, a similar report is displayed, containing the same information for each interface. Refer to [Table 6-13](#) for a description of the information.

```
Monitor>statistics 1
```

```
Monitor>statistics 1
          Unicast      Multicast      Unicast      Multicast
Ifc Name  Pkts Rcv      Pkts Rcv      Pkts Trans   Pkts Trans
 1 Eth/1      2             0             2            0
```

Monitoring Packet Statistics and Error Counts

Table 6-13: Description of the Packet Statistics Displayed

Information Displayed	Description
Ifc	Physical and logical interface number.
Name	Interface type (Ethernet, Fast Ethernet, FDDI, ATM LEC, ATM FDDI Bridge Tunnel, and so on), followed by a number indicating the instance within this type. Instance numbering starts at zero.
Unicast Pkts Rcv	Number of nonmulticast, nonbroadcast, specifically addressed packets at the MAC layer.
Multicast Pkts Rcv	Number of multicast or broadcast packets received.
Packets Trans	Number of packets transmitted.
Multicast Pkts Trans	Number of multicast or broadcast packets transmitted.

Displaying Interface Test Results

You can display a report that includes the results of interface maintenance tests and self-tests for all interfaces. Refer to the [Testing an Interface](#) section for additional information about interface tests.

To display interface test results for all interfaces, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter interface statistics .
2	Press Return. The test report is displayed.

Example

Monitor> **interface statistics**

Ifc	Ifc'	Name	Self-Test Passed	Self-Test Failed	Maintenance Failed
0	0	VNbus/0	1	2	0
1	1	Eth/1	1	5	0
2	2	Eth/2	1	4	0
3	3	Eth/3	1	2	0
4	4	Eth/4	0	62750	0
5	5	Eth/5	0	62750	0
6	6	Eth/6	1	57052	1
7	7	Eth/7	0	62750	0
8	8	Eth/8	0	62750	0
9	9	Eth/9	0	62750	0
10	10	Eth/10	0	62750	0
11	11	Eth/11	0	62750	0
12	12	Eth/12	0	62750	0
13	13	FDDI/13	1	1	0
14		VLAN/14			
.....VLAN's have no hardware, selftest or maintenance.					
45		VLAN/45			

Clearing Interface Counters

You can clear (reset to zero) error, packet, and other counters associated with a specific interface, or all interfaces. You may find the ability to do so useful if you want to track recent changes in one or more counters. The counters cleared are those discussed in the [Monitoring an Ethernet Interface](#), [Monitoring a Fast Ethernet Interface](#), [Monitoring an FDDI Interface](#), [Monitoring an ATM Interface](#), and [Monitoring the LAN Emulation ARP Database](#) sections. For a complete list of interface and bridge port counters, see [Appendix C](#).

To clear interface counters, perform the following steps:

Step	Action
1	Access the Monitor prompt (Monitor>).
2	If you want to clear counters for all interfaces on the module, enter clear . If you want to clear counters for a specific interface on the module, enter clear interface# , where <i>interface#</i> is the number of the interface for which you want to clear counters.
3	Press Return. The following message is displayed: Clear network statistics? (Yes or [No]):
4	Enter y if you want to clear the interface counters. Enter n if you do not want to clear the interface counters.
5	Press Return. The counters are cleared or not cleared as specified, and the Monitor prompt is displayed.

Testing an Interface

Ethernet, Fast Ethernet, FDDI, and ATM interfaces perform a self-test when they are first powered up. With the exception of ATM, these interfaces periodically perform maintenance tests while they are up. Self-tests typically provide a more thorough verification process than maintenance tests. ATM interfaces perform self-tests only during power-up and do not perform maintenance tests. Doing so would disable all ATM logical interfaces, not just the interface you specify.

To test a physical interface, use the **enable interface *n*** and **disable interface *n*** commands at the Config prompt as explained in the [Enabling and Disabling an Interface](#) section of Chapter 5.

On interfaces where no cable is attached, the "Self-Test Failed" counter on the `Monitor> interface` screen continues to increment. To get more detailed information on the state of an interface, use the `Monitor> interface statistics n` command.

Monitoring the Address Resolution Protocol

This section provides information about how to monitor the Address Resolution Protocol (ARP). Monitoring includes the ability to perform the following tasks:

- List learned and manual entries in the ARP cache.
- List the interfaces registered with ARP.
- List the protocols with addresses registered with ARP.
- Display ARP operational statistics, including such information as input packet overflows.
- Delete (clear) all learned entries from the ARP cache.

To monitor the interface, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter <code>arp</code> .
2	Press Return. The <code>ARP></code> prompt is displayed.
3	Enter one of the command options shown in Table 6-14 to display the desired information.
4	Press Return. The desired information and the <code>ARP></code> prompt is displayed.

Table 6-14: ARP Monitor Command Options

Command Option	Description
<code>clear interface</code> , where <i>interface</i> is the number of the TCP/IP Host Services logical interface. (Refer to Chapter 1 for a description of the Host Services logical interface.) To determine the correct interface number, first use the <code>hardware</code> command, below. The interface number you enter is the number adjacent to the TCP/IP Host Services logical interface name <code>VLAN/n</code> .	Deletes all learned and manually entered entries from the ARP cache associated with the specified interface. This command option can be used to force the deletion of bad transactions.

Command Option	Description
<p><u>d</u>ump <i>interface</i> [<i>protocol</i>], where <i>interface</i> is the number of the TCP/IP Host Services logical interface. (Refer to Chapter 1 for a description of the Host Services logical interface.) To determine the correct interface number, first use the hardware command, below. The interface number you enter is the number adjacent to the TCP/IP Host Services logical interface name VLAN/<i>n</i>. The <i>protocol</i> parameter specifies the short name or number of the protocol for which you want to display information. Use the protocol option to determine the protocol's name or number. The protocol name and number are optional if ARP is in use by only one protocol on the specified interface.</p>	<p>Displays the ARP cache for the specified interface/protocol combination. If more than one protocol is on the interface, the protocol number must also be entered. This displays the address-to-protocol mappings.</p>
<u>h</u>ardware	<p>Displays the interfaces registered with ARP. The interface is identified by number and type. The command option also displays the interface's hardware address space (AS) and local hardware address.</p>
<u>p</u>rotocol	<p>Displays the protocols for each interface that have addresses registered with ARP. The information displayed includes the interface number and type, the protocol name, the protocol address space (AS) in hexadecimal, and the associated local protocol address.</p>

Monitoring the Address Resolution Protocol

Command Option	Description
statistics	<p>Displays the number of ARP input packet overflows by interface, and the following ARP cache information for each protocol per interface:</p> <p>Max — All time maximum length hash chain. Cur — Current maximum length hash chain. Cnt — Number of currently active entries. Alloc — Number of entries created. Refresh:Tot — Number of refresh requests sent for the network interface and protocol. Failure — Number of Auto-Refresh attempt failures due to unavailability of internal resources. This count is not related to whether or not an entry was refreshed. TMOs: Refresh — Number of entries deleted due to a timeout of the Refresh Timer.</p>

Monitoring ICMP Counters

Internet Control Message Protocol (ICMP) is a part of IP that handles error and control messages, including an echo request/reply function to test whether a destination is reachable and responding. ICMP provides error, status and administrative messages that are incorporated into the data field of an IP packet. To display the list of ICMP counters, perform the following steps:

Step	Action
1	At the <code>IP></code> prompt, enter: <code>icmp-counters</code>
2	Press Return. A list of ICMP counters are displayed.

Example

`IP> icmp-counters`

ICMP counters	Receive	Transmit
Total number of messages	15	17
Number of errors	0	0
Destination Unreachable	0	0
Time Exceeded	0	2
Parameter Problem	0	0
Source Quench	0	0
Redirect	0	0
Echo request	13	0
Echo reply	0	13
Timestamp request	1	0
Timestamp reply	0	1
Address Mask request	1	0
Address Mask reply	0	1

Monitoring ICMP Counters

Field	Description
Total number of messages	Total number of ICMP messages sent and received.
Number of errors	Generic errors, such as buffer allocation errors, detected by the router.
Destination Unreachable	Used by the router or the destination host and is invoked if a router encounters problems reaching the destination network specified in the IP destination address. Also, the destination host can invoke this if an identified higher-level protocol is not available on the host or if a specified port is not available.
Time Exceeded	Initiated by the router when the time-to-live value becomes zero or if a timer expires during reassembly of a fragmented datagram.
Parameter Problem	Initiated by a host or the router if it encounters problems processing any part of an IP header.
Source Quench	Flow and congestion control that is used if the router has insufficient buffer space for queuing incoming datagrams.
Redirect	Invoked by the router and sent to the source host and is used to provide routing management information.
Echo request	A PING message sent to any IP address to determine the state of the internet or a network segment.
Echo reply	A PING message sent by a host in response to an echo request.
Timestamp request	This is used by the router and hosts to determine the delay incurred when delivering packets through a network.
Timestamp reply	Timestamp values sent by a host in response to a timestamp request.
Address Mask request	Used by a host to obtain a subnet mask used on the host's network. The requesting host can send the request directly to a router or broadcast it.
Address Mask reply	A reply from an address mask agent host or any authoritative originator of the address mask on the network containing the network's subnet mask.

Monitoring ICMP Counters

Chapter 7

Configuring the Transparent Bridge

Overview

Introduction

This chapter provides information about how to configure the DIGITAL VNswitch module transparent bridge and how to display information about the current configuration. The parameters you configure, as described in this chapter, are stored in nonvolatile RAM (NVRAM). Information stored in NVRAM is retained in memory even if power to the module is interrupted or the switch is reset.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Accessing and Exiting the Bridge Configuration Prompt	7-2
Setting and Enabling Switch-wide Rate Limiting	7-3
Configuring Permanent Address Filters	7-5
Configuring Protocol Filters	7-10
Configuring the Spanning Tree Protocol	7-21
Forwarding Using Only Manually Created Address Filters	7-30
Enabling and Disabling a Bridge Port	7-32
Bridging Ethernet and FDDI Networks	7-33
Auto-Testing of Ports Inactive for Extended Periods	7-39
Setting the Time That Unused Addresses Are Retained	7-40
Displaying Current Bridge Configuration Parameters	7-41
Duplicate MAC Addresses on Separate VSDs	7-45

Accessing and Exiting the Bridge Configuration Prompt

You must access the bridge configuration prompt to configure and display network parameters. The prompt is accessed from the Config prompt (`Config>`).

Accessing the Bridge Configuration Prompt

To access the Bridge configuration prompt, perform the following steps:

Step	Action
1	At the Config prompt, enter bridge . The default protocol selection is bridge . Example: <code>Config> bridge</code>
2	Press Return. The Bridge configuration prompt (<code>BRIDGE config></code>) is displayed.

Exiting the Bridge Configuration Prompt

You exit the Bridge config prompt to return to the Config prompt.

Example

To return to the Config prompt (`Config>`) from the `BRIDGE config>` prompt, enter **exit** and then press Return.

Setting and Enabling Switch-wide Rate Limiting

This section describes how to configure rate limiting on a switch-wide basis. To configure switch-wide rate limiting, perform the following tasks:

Task	Description
1	Set the maximum number of packets per second (pps).
2	Enable or disable rate limiting.

Rate limiting is used to minimize the effects of broadcast storms. A broadcast storm is typically caused when a host system responds to multicast packets that are circulating continuously on the network, or when it tries to respond to another system that never replies. The generation of such traffic at an uncontrolled rate can severely affect the available bandwidth on a network, perhaps making communications impossible.

You can limit broadcast storms to that segment of the network from which the packets are generated. You do so by setting the maximum number of multicast packets per second (pps) that the VNswitch module is to forward. The rate (pps) you set is applied only if rate limiting is enabled. If the maximum number of packets per second is reached, the switch forwards packets at the specified maximum rate, thereby limiting the effect of the broadcast storm on the other side of the switch. Switch-wide Rate Limiting is disabled by default.

Setting Maximum Frames Per Second

You can set the maximum number of frames per second either independent of enabling rate limiting, or while enabling rate limiting. To set the maximum number of frames per second independent of enabling rate limiting, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set rate-limit .
2	Press Return. The following message is displayed: Rate Limit Value (frames/sec) [400]?
3	Enter the maximum number of packets per second the VNswitch module is to forward destined for a particular MAC address, or that are of a particular protocol type. The default is 400 frames per second.
4	Press Return. The maximum number of frames per second is set and the <code>BRIDGE config></code> prompt is displayed.

Setting and Enabling Switch-wide Rate Limiting

Enabling and Disabling Rate Limiting

You enable and disable rate limiting on both a switch-wide basis and on a per-filter (MAC address or protocol) basis. This section describes how to enable and disable switch-wide rate limiting. Refer to the [Creating and Modifying a Permanent Address Filter](#) section for information about how to set rate limiting for multicast address filters. Refer to the [Creating and Modifying Protocol Filters](#) section for information about how to set rate limiting for protocol filters.

If rate limiting is disabled on the switch, it is disabled for all address and protocol filters, even if rate limiting for the individual filter is enabled. If rate limiting is enabled on the switch, it is enabled for a specific filter only if rate limiting is also enabled for the filter. Disabled is the default for switch-wide rate limiting.

Enabling Switch-wide Rate Limiting

To enable switch-wide rate limiting, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>enable rate-limit</code> .
2	Press Return. The following message is displayed: <code>Rate limit value [400]?</code>
3	Enter the maximum number of packets per second the VNswitch module is to forward destined for a particular MAC address, or that are of a particular protocol type. The factory default is 400 frames per second. If the factory default was modified, the default is the most recent setting.
4	Press Return. Rate limiting is enabled and the specified rate is applied on a switch-wide basis. The <code>BRIDGE config></code> prompt is again displayed.

Disabling Switch-wide Rate Limiting

To disable switch-wide rate limiting, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>disable rate-limit</code> .
2	Press Return. Rate limiting is disabled on a switch-wide basis.

Configuring Permanent Address Filters

This section discusses how to add, modify, and delete MAC address filters. The process of filtering network traffic is used to reduce the amount of unnecessary traffic over specific segments of a network, thereby maximizing network capacity and performance. It can also be used to restrict the distribution of sensitive information to specific locations. You can configure the VNswitch module to selectively filter or forward packets it receives, based on their MAC addresses and protocol types.

MAC address filters are stored in NVRAM as permanent entries. (Refer to [Chapter 1](#) for information about NVRAM.) Permanent address entries are retained in memory even if power to the module is interrupted or the switch is reset. Permanent entries are not affected by address Aging Time, and can exist concurrently with dynamic entries having the same address. (Refer to [Chapter 8](#) for information about configuring a MAC address as a static entry.)

You configure a filter by specifying a set of ports that are allowed for a given MAC address. The address can be an individual, multicast, or broadcast address. If a packet has a source or destination MAC address that is listed in the filter, the packet is allowed on the specified set of ports. If the address is not listed in the filter, the packet is dropped (filtered).

Source Address Filtering

A packet received on an input port is dropped (filtered) if the source address is listed in a filter and the input port is not one of the allowed ports. In addition, the source address is not learned.

If the input port is one of the allowed ports, the packet is a candidate for forwarding, based on the behavior of destination address filtering and protocol filtering. (Refer to the [Configuring Protocol Filters](#) section for information about protocol filtering.)

Destination Address Filtering

A packet about to be placed on an output port is forwarded if the destination address is listed in a filter and the output port is one of the allowed ports. If the address was not previously learned, the packet is flooded to the subset of the allowed ports that are in the forwarding state.

Configuring Permanent Address Filters

Creating and Modifying a Permanent Address Filter

To create or modify a MAC address filter, you must perform the following tasks:

Task	Description
1	Identify the address to be filtered, and the ports on which the address is allowed.
2	Enable or disable multicast rate limiting (multicast addresses only).

Identifying the Address and Allowed Ports

The addresses you specify when creating a filter are added to the switch's permanent database (NVRAM). To identify the allowed ports for an address, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set address . You can, alternatively, enter set address mac_address , where <i>mac_address</i> is the 12-digit MAC address for which you want to create or modify a filter on the module. If you enter the command using this syntax, go to step 4.
2	Press Return. The following message is displayed: Address (in 12-digit hex) []?
3	Enter the 12-digit MAC address for which you want to create or modify a filter. You can enter the address with or without hyphens separating the octets. For example, you can enter 11-22-33-44-55-66 or 112233445566 .
4	Press Return. The following message is displayed: Enter allowed ports ("None" or "All") []? If you are modifying the forwarding ports for an existing address, the previously configured ports are the default.

Configuring Permanent Address Filters

Step	Action
5	<p>Enter a list of the bridge ports to which you want the packet with the specified MAC address forwarded. (The packet is filtered from those ports not entered.) Enter None if you want the packet filtered from all ports. Enter all if you want the packet allowed on all ports. You can list the forwarding ports individually by separating each with a comma (for example, 0,1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 0-5). You can also combine a list of individual entries with a range of entries (for example, 0-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.</p>
6	<p>Press Return.</p> <p>If the address is a unicast address, the address is set and the <code>BRIDGE config></code> prompt is displayed.</p> <p>If the address is a multicast address, the following message is displayed:</p> <pre>Enable Multicast Rate Limiting (Yes or No)? [No]:</pre> <p>Refer to the Enabling and Disabling Multicast Rate Limiting section.</p> <p>If any of the specified ports do not exist, the address is set for forwarding on the port but the port setting has no effect. A message similar to the following is displayed:</p> <pre>Warning, ports 9, 10 do not exist</pre>

Enabling and Disabling Multicast Rate Limiting

Multicast rate limiting is used to minimize the effects of multicast storms. You can restrict multicast storms, consisting of packets that have a specific destination MAC address, to that segment of the network from which the packets are generated. You do so by setting the maximum number of those packets the VNswitch module is to forward per second, and by enabling rate limiting.

Rate limiting is enabled on both a per-address basis and a switch-wide basis. This section provides instructions about how to enable or disable rate limiting on a per MAC address basis. Refer to the [Setting and Enabling Switch-wide Rate Limiting](#) section for information about setting the maximum number of packets per second and enabling or disabling rate limiting on a switch-wide basis.

Configuring Permanent Address Filters

To enable or disable multicast rate limiting for the address specified in the [Identifying the Address and Allowed Ports](#) section, perform the following steps:

Step	Action
1	Enter Yes to enable Multicast Rate Limiting. Enter No to disable Multicast Rate Limiting. No is the default unless it was previously enabled for the address.
2	Press Return. The selected filter is applied to the module and rate limiting is enabled or disabled, as specified.

Deleting Permanent Address Filters

This section describes how to delete MAC address filters that you, or another switch administrator, previously entered manually. (Refer to the [Creating and Modifying a Permanent Address Filter](#) section.) You can delete address filters one at a time, or you can delete all MAC address filters previously added to the switch's permanent database (NVRAM). Deleting an address filter removes the address from the forwarding database as well as its associated multicast rate limiting status.

Note

Registered unicast and reserved addresses cannot be deleted.

Deleting a Single Permanent Address Filter

To delete one MAC address, its filter, and its associated rate limiting status, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter delete address . You can, alternatively, enter delete address mac-address , where <i>mac-address</i> is the 12-digit MAC address you want to delete. If you enter the command using this syntax, go to step 4.
2	Press Return. The following message is displayed: <code>Address (in 12-digit hex) []?</code>
3	Enter the 12-digit MAC address you want to delete. You can enter the address with or without hyphens separating the octets. For example, you can enter 11-22-33-44-55-66 or 112233445566 .

Configuring Permanent Address Filters

Step	Action
4	Press Return. The specified address is deleted. If you entered an address that does not exist, the following message is displayed: No entry found for this address

Deleting All Permanent Address Filters

To delete all MAC addresses and associated filters and rate limiting status, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter delete address all .
2	Press Return. All MAC addresses are deleted.

Configuring Protocol Filters

Protocol filtering enables you to configure a module so that it selectively filters traffic based on the data's MAC frame format. The VNswitch 900 family of modules supports protocol filtering for the following frame types:

- Ethernet-II (IEEE 802.3)
- IEEE 802.3 Subnetwork Access Protocol (SNAP)
- IEEE 802.3 Destination Service Access Point (DSAP)

Each frame type can conform to one of a number of different protocols. Common protocols are listed in [Table 7-1](#) through [Table 7-3](#). When you create a protocol filter, you can configure the module to receive or forward frames on one or more selected bridge ports, based on the frame's protocol. For example, you can configure one or more ports on the module so that Ethernet frames, conforming to the AppleTalk Phase 1 protocol, are received or forwarded only on those ports.

A received packet is dropped (filtered) if there is a filter for the protocol and the input port is not one of the allowed ports. Also, the packet can only be output on a port that is one of the allowed ports.

You can also configure the module to receive or forward *all* frames of a particular type from one or more bridge ports, using protocols not already filtered as described above. You do so by configuring a default protocol filter. For example, if you configure the module so that Ethernet frames conforming to the AppleTalk Phase 1 protocol are discarded, you can then use a default filter to discard (or receive/forward) all Ethernet frames of any other protocol type such as IP, ARP, DECnet, and so on. Refer to the [Creating and Modifying Default Protocol Filters](#) section for information about how to create a default protocol filter.

Note

VNswitch firmware allows you to configure protocol filters based on both the encapsulation and the protocol type. That is, to configure a protocol filter for a given set of ports, the user chooses the encapsulation, the protocol type, and the list of the ports to which the filter applies.

When configuring protocol filters, keep the following in mind. Forwarding of a packet from a LAN segment with one type of encapsulation to another LAN segment with a different encapsulation type requires translation. The translation of the packet takes place after the filter-forwarding decision is made. Therefore, if an Ethernet IP packet is forwarded to the FDDI port, an Ethernet IP filter needs to be set for the FDDI port for filtering to occur, even though the transmitted packet has a SNAP encapsulation.

Configuring Protocol Filters

To prevent an error in the protocol filter configuration for a given set of ports, the best approach is to set protocol filters for all encapsulation types of the protocol to be filtered on each set of ports. This, typically, can be done without any side effects. If this approach interferes with other considerations, base the configuration on filter encapsulation and choose with caution.

Creating and Modifying Protocol Filters

To create or modify a protocol filter, perform the following tasks:

Task	Description
1	Identify the specific protocol to be filtered and the allowed ports on which frames of the specified protocol type can be received and forwarded.
2	Enable or disable protocol rate limiting.

Identifying the Protocol and Allowed Ports

To specify the protocol to be filtered and the allowed ports on which the frames are received or forwarded, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set protocol-filter .
2	Press Return. The following message is displayed: <code>Protocol type [DSAP]?</code>
3	Enter one of the following types of frames for which you want to specify a protocol: <ul style="list-style-type: none">• ether (Ethernet-II)• snap• dsap DSAP is the default.
4	Press Return. If you entered ether , the following message is displayed: <code>Protocol Type in hex (5DD-FFFF) [0800]?</code> If you entered snap , the following message is displayed: <code>Address (in 10-digit hex) [00-00-00-08-00]?</code> If you entered dsap , the following message is displayed: <code>Protocol Type in hex (0-FF) [1]?</code>

Configuring Protocol Filters

Step	Action
5	<p>Enter the hexadecimal value associated with the protocol you want to filter from receiving ports.</p> <p>Refer to Table 7-1 for the hexadecimal value of common Ethernet protocols you can filter.</p> <p>Refer to Table 7-2 for the hexadecimal value of common SNAP protocols you can filter.</p> <p>Refer to Table 7-3 for the hexadecimal value of common DSAP protocols you can filter.</p>
6	<p>Press Return. The following message is displayed:</p> <pre>Enter allowed ports ("None" or "All") [0-254]?</pre> <p>The default is all ports if you are creating a new filter. The previously configured ports are the default if you are modifying an existing protocol filter.</p>
7	<p>Enter a list of the bridge ports on which you want frames to be received or forwarded based on the specified protocol. (The frames are filtered from those ports not entered.) Enter None if you want the frames filtered from all ports. Enter all if you want the frames allowed on all ports. You can list the ports individually by separating each with a comma (for example, 0,1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 0-5). You can also combine a list of individual entries with a range of entries (for example, 0-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.</p>
8	<p>Press Return. The following message is displayed:</p> <pre>Enable Multicast Rate Limiting (Yes or No)? [Yes]:</pre> <p>The default is No if you are creating a new filter. The default is the previously configured setting if you are modifying an existing filter.</p> <p>Refer to the Enabling and Disabling Multicast Rate Limiting section.</p> <p>If the specified ports do not exist, the frame is set for forwarding on the port but is not operational until the port is added. A message similar to the following is displayed in addition to the rate limiting prompt:</p> <pre>Warning, ports 9, 10 do not exist</pre>

Table 7-1: Hexadecimal Values for Common Ethernet-II (IEEE 802.3) Protocols

Protocol	Hexadecimal Value
IP	0800
ARP	0806
CHAOS	0804
DECnet MOP Dump/Load	6000
DECnet MOP Remote Console	6002
DECnet	6003
DEC LAT	6004
DEC LAVC	6007
XNS	0600
Maintenance Packet Type	7030
Apollo Domain	8019 (Ethernet)
Novell NetWare IPX	8137 (Ethernet)
AppleTalk Phase 1	809B
Apple ARP Phase 1	80F3
Loopback assistance	9000

Table 7-2: Hexadecimal Values for Common SNAP OUI/IP Protocols

Protocol	Ten-Digit Hexadecimal Value
AppleTalk Phase 2	08-00-07-80-9B
Apple ARP Phase 2	00-00-00-80-F3
Proprietary AppleTalk Phase 1 for FDDI	00-00-93-00-02
Proprietary AppleTalk ARP Phase 1 for FDDI	00-00-93-00-03

Configuring Protocol Filters

Table 7-3: Hexadecimal Values for Common DSAP Protocols

Protocol	Hexadecimal Value
Banyan SAP	BC (used for only 802.5)
Novell IPX SAP	E0 (used for only 802.5)
NetBIOS SAP	F0
ISO Connectionless Internet	FE

Enabling and Disabling Multicast Rate Limiting

Multicast rate limiting is used to minimize the effects of broadcast storms. You can limit broadcast storms, consisting of frames of a particular protocol type, to that segment of the network from which the frames are generated. You do so by setting the maximum number of those frames the VNswitch module is to receive per second, and by enabling rate limiting.

Rate limiting is enabled on both a protocol basis and a switch-wide basis. This section provides instructions about how to enable or disable rate limiting on a protocol basis. Refer to the [Setting and Enabling Switch-wide Rate Limiting](#) section for information about setting the maximum number of packets per second (pps) and enabling/disabling rate limiting on a switch-wide basis.

To enable or disable rate limiting for the protocol specified in the [Identifying the Protocol and Allowed Ports](#) section, perform the following steps:

Step	Action
1	Enter Yes to enable Multicast Rate Limiting. Enter No to disable Multicast Rate Limiting. No is the default.
2	Press Return. The selected filter is applied to the module and rate limiting is enabled or disabled, as specified.

Deleting Protocol Filters

You can delete previously configured protocol filters from the VNswitch module. You can delete filters one at a time by protocol type, or you can delete all protocol filters by frame type or protocol type.

Deleting a protocol filter may increase the amount of traffic over specific segments of a network and may, therefore, reduce network capacity and performance.

Deleting a Single Filter by Protocol Type

To delete a single filter by protocol type, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>delete protocol-filter</code> .
2	Press Return. The following message is displayed: <code>Protocol type [DSAP]?</code>
3	Enter one of the following types of frames for which you want to delete a protocol filter: <ul style="list-style-type: none"> • <code>ether</code> (Ethernet-II) • <code>snap</code> • <code>dsap</code> DSAP is the default.
4	Press Return. If you entered <code>ether</code> , the following message is displayed: <code>Protocol Type in hex (5DD-FFFF or ALL) [800]?</code> If you entered <code>snap</code> , the following message is displayed: <code>Address (in 10-digit hex or ALL) [00-00-00-08-00]?</code> If you entered <code>dsap</code> , the following message is displayed: <code>Protocol Type in hex (0-FF or ALL) [1]?</code>
5	Enter the hexadecimal value associated with the filter you want to delete. Refer to Table 7-1 for the hexadecimal value of common Ethernet protocols you can filter. Refer to Table 7-2 for the hexadecimal value of common SNAP protocols you can filter. Refer to Table 7-3 for the hexadecimal value of common DSAP protocols you can filter.
6	Press Return. The protocol filter is deleted and the <code>BRIDGE config></code> prompt is displayed.

Configuring Protocol Filters

Deleting All Protocol Filters of a Particular Frame Type

To delete all protocol filters associated with a particular frame type, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>delete protocol-filter</code> .
2	Press Return. The following message is displayed: <code>Protocol type [DSAP]?</code> The default is DSAP, ETHER, or SAP if the respective filters of that protocol are present.
3	Enter one of the following frame types for which you want to delete all associated protocol filters: <ul style="list-style-type: none">• <code>ether</code> (Ethernet-II)• <code>snap</code>• <code>dsap</code> The default is DSAP.
4	Press Return. If you entered <code>ether</code> , the following message is displayed: <code>Protocol Type in hex (5DD-FFFF or ALL) [800]?</code> If you entered <code>snap</code> , the following message is displayed: <code>Address (in 10-digit hex or ALL) [00-00-00-08-00]?</code> If you entered <code>dsap</code> , the following message is displayed: <code>Protocol Type in hex (0-FF or ALL) [1]?</code>
5	Enter a protocol type or enter <code>all</code> and go to Step 6.
6	Press Return. A message similar to the following is displayed: <code>Modify default protocol filters for DSAP (Yes or No)? [Yes]:</code>
7	If you want to modify the default protocol filters associated with all the protocol filters you are about to delete, enter <code>Yes</code> . If you do not want to modify the default protocol filters associated with all the protocol filters you are about to delete, enter <code>No</code> .
8	Press Return. If you entered <code>Yes</code> , go to step 9. If you entered <code>No</code> , go to step 11.

Configuring Protocol Filters

Step	Action
9	The following message is displayed: Allowed port numbers ("None" or "All") [0-254]? The previously configured ports are the default if you are modifying an existing default protocol filter.
10	Enter a list of the bridge output ports from which you want frames forwarded based on the specified protocol. (The frames are filtered from those ports not entered.) Enter None if you want the frames filtered from all ports. Enter all if you want the frames forwarded on all ports. You can list the forwarding ports individually by separating each with a comma (for example, 0,1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 0-5). You can also combine a list of individual entries with a range of entries (for example, 0-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.
11	Press Return. All filters associated with the specified frame type are modified and the Bridge config prompt is displayed.

Deleting All Protocol Filters For All Frame Types

To delete all protocol filters associated with all frame types, perform the following steps:

Step	Action
1	At the BRIDGE config> prompt, enter delete protocol-filter all .
2	Press Return. The following message is displayed: Modify default protocol filters for DSAP (Yes or No)? [Yes]:
3	If you want to modify the default protocol filters associated with all the DSAP (Ether or SNAP) protocol filters you are about to delete, enter Yes . If you do not want to modify the default protocol filters associated with all the DSAP (Ether or SNAP) protocol filters you are about to delete, enter No .
4	Press Return. If you entered Yes , go to step 5. If you entered No , go to step 7.

Configuring Protocol Filters

Step	Action
5	The following message is displayed: Allowed port numbers ("None" or "All") [0-254]? The previously configured ports are the default if you are modifying an existing default protocol filter.
6	Enter a list of the bridge output ports from which you want frames forwarded based on the specified protocol. (The frames are filtered from those ports not entered.) Enter None if you want the frames filtered from all ports. Enter all if you want the frames forwarded on all ports. You can list the forwarding ports individually by separating each with a comma (for example, 0,1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 0-5). You can also combine a list of individual entries with a range of entries (for example, 0-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.
7	Press Return. When you have modified all the protocol filters for a given frame type, you can modify protocol filters for the next frame type (dsap, ether, or snap). A message similar to the following is displayed: Modify default protocol filters for ETHER (Yes or No)? [Yes]: If you want to modify additional protocol filters, enter yes and go to step 3. If there are no additional default protocol filters to modify, go to step 8.
8	All protocol filters associated with all frame types are deleted and the Bridge config prompt is displayed.

Creating and Modifying Default Protocol Filters

Default protocol filters forward (or discard) all frames of a particular type (Ethernet, SNAP, or DSAP) from one or more bridge ports. They act on only those frames that conform to protocols not already specified as described in the [Identifying the Protocol and Allowed Ports](#) section.

To create or modify a default protocol filter, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set default-protocol-filter .
2	Press Return. The following message is displayed: Protocol type [DSAP]?

Configuring Protocol Filters

Step	Action
3	<p>Enter one of the following types of frames for which you want to specify a default protocol:</p> <ul style="list-style-type: none">• ether (Ethernet-II)• snap• dsap <p>DSAP is the default.</p>
4	<p>Press Return.</p> <p>Allowed port numbers ("None" or "All") [0-254]?</p> <p>The default is all ports if you are creating a new filter. The previously configured ports are the default if you are modifying an existing filter.</p>
5	<p>Enter a list of the bridge output ports from which you want frames forwarded based on the specified protocol. (The frames are filtered from those ports not entered.) Enter None if you want the frames filtered from all ports. Enter all if you want the frames forwarded on all ports. You can list the forwarding ports individually by separating each with a comma (for example, 0,1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 0-5). You can also combine a list of individual entries with a range of entries (for example, 0-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.</p>
6	<p>Press Return. The default protocol filter is created or modified and the <code>BRIDGE config></code> prompt is displayed.</p>

Example

You can configure a module so that Ethernet frames conforming to the AppleTalk Phase 1 protocol are forwarded. You do so by creating a discrete protocol filter. Refer to the [Creating and Modifying Protocol Filters](#) section for information about how to do so. You can then use a default filter to discard all other Ethernet frames of any other protocol type such as IP, ARP, DECnet, and so on. Refer to [Table 7-1](#) for a list of common Ethernet protocol types that can be filtered in this way.

Configuring Protocol Filters

Deleting Default Protocol Filters

You can delete previously configured default protocol filters from the VNswitch module. When you delete a default protocol filter, all protocols associated with the specified frame type are then forwarded on all ports.

Deleting a default protocol filter may increase the amount of traffic over specific segments of a network and may, therefore, reduce network capacity and performance.

To delete a default protocol filter, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>delete default-protocol-filter</code> .
2	Press Return. The following message is displayed: <code>Protocol type [DSAP]?</code>
3	Enter one of the following frame types for which you want to delete all default protocol filters: <ul style="list-style-type: none">• <code>ether</code> (Ethernet-II)• <code>snap</code>• <code>dsap</code> DSAP is the default.
4	Press Return. The default protocol filter is deleted and the Bridge config prompt is displayed.

Configuring the Spanning Tree Protocol

The Spanning Tree Protocol (STP) is used to detect and break circular traffic patterns, or loops. Loops can cause exponential traffic replication resulting in severe network congestion. The VNswitch module implements the IEEE 802.1D Spanning Tree Protocol. The VNswitch module does not interoperate with DEC LANbridge 100 STP.

STP eliminates loops by selecting one of the redundant bridges as the primary or root bridge, and maintaining secondary bridges as backups. The root bridge is selected, in part, based on a priority value you can assign to each bridge. STP then selects a root port on each secondary bridge. The root port is selected based on the relative priority of each port, and which port provides access to the root bridge at the least cost. The bridge then enables all root ports for forwarding, and disables other ports to prevent loops.

The Spanning Tree Protocol algorithm runs automatically for each VSD you create. Therefore, multiple spanning trees are supported on the bridge. The STP parameters you set, as described in this section, affect all instances of STP. You cannot configure different values for each instance of STP. Refer to [Chapter 9](#) for information about Virtual LANs (VLANs).

You can enable and disable STP on a port. You can also set certain parameters that influence selection of the root bridge, the root port, and the frequency with which changes in network topology are detected. These options are discussed in the following sections.

Note

The concepts presented in this section are simplified for the purpose of providing a general introduction. Refer to the *Digital VNswitch 900 Series Technical Overview* for a more detailed discussion of STP concepts.

Configuring the Spanning Tree Protocol

Influencing Selection of the Root Bridge

The root bridge is selected based on a value that is the combination of a bridge's address and a number you can assign that is the bridge's priority. A lower number for bridge priority makes it likely that the bridge is selected as the root.

To set the bridge's priority, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set bridge priority .
2	Press Return. The following message is displayed: <code>Bridge-Priority [32768]?</code>
3	Enter a number from 0 to 65535 for the bridge priority. The lower the number, the more likely the bridge is selected as the root. The factory default bridge priority is 32768. If the factory default was modified, the default is the most recent setting.
4	Press Return. The number entered is set as the bridge's priority and the <code>BRIDGE config></code> prompt is displayed.

Influencing Selection of the Root Port

The root port on a bridge is selected based on which port provides access to the root bridge at the lowest Path Cost and the relative Priority of each port.

You can choose either to set the Path Cost for a port manually through the CLI or allow STP to do so automatically. If the spanning tree sets the Path Cost, the value it assigns is based on the line speed associated with the port. The higher the line speed, the lower the assigned cost. However, you can override this value by manually entering a value that weights the Path Cost in favor of a particular port. The lower a port's Path Cost, the more likely it is to be selected as the root port. You may want to set a high Path Cost for a port if, for example, the LAN to which the port is connected has a low bandwidth.

Setting port priorities also affects which port is selected as the bridge's root port when a bridge has two ports connected in a loop. The port having the lowest assigned number is more likely to be selected as the root port. For example, if port 3 is assigned a Priority of 1, and port 6 is assigned a Priority of 2, port 3 is likely to be selected as the root port by STP. If more than one port has the same Priority, the port that has the lowest port number is used.

Configuring the Spanning Tree Protocol

To set the Path Cost and Priority for one or more ports on a switch, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set stp port .
2	Press Return. The following message is displayed: <code>Port Number [0]?</code>
3	Enter the number of the port for which you want to set Path Cost and Priority. The default is port 0.
4	Press Return. The following message is displayed: <code>Port Path-cost (0 for default) [0]?</code>
5	If you want STP to calculate the Path Cost for the port, enter the default value of 0. If you want to enter your own value, enter a number from 1 through 65535. The lower the value entered, the more likely the port will be selected as the root port.
6	Press Return. The following message is displayed: <code>Port Priority [128]?</code>
7	Enter a number from 0 through 255. The lower the number entered, the more likely the port will be selected as the root port.
8	Press Return. The port's Path Cost and Priority are set and the <code>BRIDGE config></code> prompt is displayed.

Detecting Changes in Network Topology

All bridges periodically transmit Hello Messages (Configuration Bridge Protocol Data Units — BPDUs) on their subset of designated ports. These messages, which include such information as the Path Cost of the port from which the message is transmitted, are initially used to determine the network topology and set up a spanning tree. When the spanning tree is established, Hello Messages are then used to detect changes in network topology by comparing the latest BPDU received on a port with the best BPDU previously received. (The best BPDU is, in part, based on Path Cost.) If a change in topology is detected, STP recomputes the spanning tree.

Configuring the Spanning Tree Protocol

You can fine-tune the ability of STP to detect changes in topology by performing the following tasks:

- Adjusting the frequency with which bridges transmit Hello Messages
- Fine-tuning the ability to detect a failed bridge or link
- Avoiding loops

Caution

It is recommended that you alter the STP default settings associated with these tasks *only* if you fully understand the effect the change will have on the bridge network.

Adjusting Hello Message Frequency

You can configure the frequency, in seconds, with which bridges transmit Hello Messages to check for changes in network topology. The parameter used to configure this value is called the Bridge Hello Time. The greater the frequency (the lower the number of seconds), the sooner changes in topology are detected. The lower the frequency (the greater the number of seconds), the lower the overhead associated with detecting changes.

Refer to the [Setting Bridge Hello Time, Bridge Max Age, and Bridge Forward Delay](#) section for information about how to configure the Bridge Hello Time.

Detecting a Failed Bridge or Link

You can also fine-tune STP's ability to detect a change in topology that is the result of a failed bridge or link. This is accomplished by setting or modifying the Bridge Max Age. If the time since a bridge last received a Hello Message on a port exceeds the Bridge Max Age setting, perhaps the result of a failed link, the bridge recalculates the root, path cost, and root port.

Note

Although you can configure each bridge in a spanning tree with a different value for Bridge Max Age, the value configured on the root bridge is used by all bridges in the spanning tree. The values configured on non-root bridges are used only if a non-root bridge subsequently becomes root.

Configuring the Spanning Tree Protocol

The lower the Bridge Max Age value, the earlier a failed bridge or link may be detected. However, if the Bridge Max Age time is exceeded due to a normal drop in network activity and not because of a link or bridge failure, the result may be failure to compute a correct spanning tree. This may cause forwarding loops and severe network congestion.

Refer to the [Setting Bridge Hello Time, Bridge Max Age, and Bridge Forward Delay](#) section for information about how to configure the Bridge Max Age.

Loop Avoidance

The Bridge Forward Delay parameter is used to prevent temporary forwarding loops between bridges. If a temporary forwarding loop occurs, it can cause severe network congestion. The Bridge Forward Delay must be at least twice the maximum amount of time it takes for data to traverse the network.

Refer to the [Setting Bridge Hello Time, Bridge Max Age, and Bridge Forward Delay](#) section for information about how to configure the Bridge Forward Delay.

Setting Bridge Hello Time, Bridge Max Age, and Bridge Forward Delay

There is a functional relationship among Bridge Hello Time, Bridge Max Age, and Bridge Forward Delay. That relationship requires that the values set for the three parameters conform to the following algorithm:

$$(2) \times (\text{Bridge Forward Delay} - 1 \text{ second}) \geq \text{Bridge Max Age}$$

and

$$\text{Bridge Max Age} \geq (2) \times (\text{Bridge Hello Time} + 1 \text{ second})$$

To set or modify Bridge Hello Time, Bridge Max Age, and Bridge Forward Delay, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set stp bridge .
2	Press Return. The following message is displayed: Bridge-Max-Age [20]?
3	Enter a value for the maximum amount of time a bridge port waits to receive a Hello Message before the bridge recalculates the root, path cost, and root port. The value can be from 6 through 40. The default is 20.
4	Press Return. The following message is displayed: Bridge-Hello-Time [2]?

Configuring the Spanning Tree Protocol

Step	Action
5	Enter the frequency with which you want the bridge port to transmit Hello Messages. The value can be from 1 through 10 seconds. The default is 2 seconds.
6	Press Return. The following message is displayed: Bridge-Forward-Delay [15]?
7	Enter the number of seconds you want a bridge to wait before allowing disabled bridge ports to transition to the forwarding state. The value can be from 4 through 30 seconds. The default is 15 seconds.
8	Press Return. The values you entered are set and the BRIDGE config> prompt is displayed. If the entered parameters do not conform to the algorithm described at the beginning of the Setting Bridge Hello Time, Bridge Max Age, and Bridge Forward Delay section, an error message is displayed.

Enabling and Disabling STP

You can disable and reenable STP on a specific port, if the port currently exists. STP is automatically enabled (the default) on all bridge ports at installation.

Caution

DIGITAL recommends you keep STP enabled on all ports. Disabling STP may result in loops and severe network congestion.

Disabling STP

To disable STP on a specific port, perform the following steps. Disabling STP on a port causes the port to transition to the forwarding state immediately.

Step	Action
1	At the BRIDGE config> prompt, enter disable stp .
2	Press Return. The following message is displayed: Port Number [0]?
3	Enter the number of the bridge port for which you want to disable STP.

Configuring the Spanning Tree Protocol

Step	Action
4	Press Return. STP is disabled on the specified port and the BRIDGE config> prompt is displayed.

Enabling STP

To enable STP on a specific port, perform the following steps:

Step	Action
1	At the BRIDGE config> prompt, enter enable stp .
2	Press Return. The following message is displayed: Port Number [0]?
3	Enter the number of the bridge port for which you want to enable STP.
4	Press Return. STP is enabled on the specified port and the BRIDGE config> prompt is displayed.

Setting a Port to Start Forwarding Immediately (Fast Start)

You can use the Fast Start feature to enable ports in the down state to begin forwarding immediately once they are in the up state.

A port transitions from the down to up state in the following cases:

- After a power-up or restart
- After a bridge port is connected to a network
- If a bridge port is terminated when a directly-connected workstation on the port is turned on, such as would be the case for a private ethernet

When the down port comes up, it skips the learning and listening states of the spanning tree protocol and begins forwarding immediately. The STP algorithm is executed on a port with Fast Start enabled exactly as it is on a port with Fast Start disabled, except during the initial down to up transition.

Because the Fast Start feature skips the listening and learning stages until after the port has come up, it has two side effects:

- A small amount of flooding may occur initially, until addresses are learned
- A short time of looping (limited to the hello-time, typically 1 or 2 seconds) may occur until the BPDU is received.

Configuring the Spanning Tree Protocol

Enabling Fast Start

To enable Fast Start, enter the following command at the BRIDGE Config> prompt:
enable fast start *n*, where *n* is the port number.

Disabling Fast Start

To disable the fast start feature on a port, enter the following command at the BRIDGE Config> prompt: **disable fast start *n***, where *n* is the port number.

Displaying the Status of the Fast Start Feature on a Port

You can view the status of the Fast Start feature for a port in several ways:

Using the List Port Command

From either the bridge configuration (BRIDGE config>) prompt or the bridge monitor prompt (BRIDGE>), use the list port command.

```
BRIDGE config> list port 3

Port Id (dec)      : 128:03, (hex): 80-03
Port State        : Enabled
STP Participation: Enabled
STP Fast Start    : Enabled
Manual Mode       : Disabled
Assoc Interface   : 3 (Eth/3  )
Path Cost         : 100
```

Using the List Bridge Command

From either the bridge configuration (BRIDGE config>) prompt or the bridge monitor prompt (BRIDGE>), use the list bridge command.

```
BRIDGE config> list bridge

Transparent Bridge Configuration

Bridge:           Enabled      Bridge Behavior: STB
Rate Limiting:    Disabled     Rate Limit Value: 400
Raw 802.3 IPX Trans: Disabled  No Frame Interval: 300

Bridge Database Information

Ageing Time: 37  Entries in Perm Database: 3
Number of Rehashes: 0

Bridge Priority: 32768/0x8000

Port Information

Number of ports: 29
```

Configuring the Spanning Tree Protocol

Port	Interface	State	Behaviour	STP	Manual Mode
0	0	Enabled	STB Only	Enabled	Disabled
1	1	Enabled	STB Only	Enabled	Disabled
2	2	Enabled	STB Only	Enabled	Disabled
3	3	Enabled	STB Only	Ena/Fast	Disabled
4	4	Enabled	STB Only	Enabled	Disabled
5	5	Enabled	STB Only	Enabled	Disabled
6	6	Enabled	STB Only	Enabled	Disabled
7	7	Enabled	STB Only	Enabled	Disabled
8	8	Enabled	STB Only	Enabled	Disabled
9	9	Enabled	STB Only	Ena/Fast	Disabled
10	10	Enabled	STB Only	Enabled	Disabled
11	11	Enabled	STB Only	Enabled	Disabled
12	12	Enabled	STB Only	Enabled	Disabled
13	13	Enabled	STB Only	Enabled	Disabled
14	14	Enabled	STB Only	Enabled	Disabled
15	15	Enabled	STB Only	Enabled	Disabled
16	16	Enabled	STB Only	Enabled	Disabled
17	17	Enabled	STB Only	Ena/Fast	Disabled
18	18	Enabled	STB Only	Ena/Fast	Disabled
19	19	Enabled	STB Only	Enabled	Disabled
20	20	Enabled	STB Only	Enabled	Disabled
21	21	Enabled	STB Only	Enabled	Disabled
22	22	Enabled	STB Only	Enabled	Disabled
23	23	Enabled	STB Only	Enabled	Disabled
24	24	Enabled	STB Only	Enabled	Disabled
25	25	Enabled	STB Only	Enabled	Disabled
26	26	Enabled	STB Only	Enabled	Disabled
27	28	Enabled	STB Only	Enabled	Disabled
28	28	Enabled	STB Only	Enabled	Disabled

Forwarding Using Only Manually Created Address Filters

You can configure the switch to forward frames using only manually created address filters. Doing so allows you to strictly control traffic through selected ports and can provide an additional level of security to the desired subnetworks.

For example, you may want to operate a port securely so that only one system is allowed to source and sink traffic on that port. To do so, you enable manual mode on the port. You then create an address filter entry (permanent or static) for that system's MAC address, specifying only that port in the filter's allowed port map. In addition, if you want multicast or broadcast frames to be forwarded to that port, you must create an address filter for each such address, specifying at least that port in the allowed port map.

You configure the use of only manually created filters by enabling manual mode on selected ports. When manual mode is enabled, the port's forwarding table that contains learned addresses is purged, and new addresses are not learned. However, you can still configure new MAC address filters manually. (Refer to the [Configuring Permanent Address Filters](#) section for information about creating address filters manually.) The addresses and locations of other network devices are relearned if manual mode is later disabled.

Enabling Manual Mode

To enable manual mode, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter enable manual-mode .
2	Press Return. The following message is displayed: <code>Port Number [0]?</code> The value displayed for port number is based on the ports that are currently enabled or disabled.
3	Enter the number of the port on which you want to enable manual mode (disable bridge learning).
4	Press Return. Manual mode is enabled on the specified port and the <code>BRIDGE config></code> prompt is displayed.
5	Repeat steps 1 through 4 for each port on which you want to enable manual mode.

Forwarding Using Only Manually Created Address Filters

Disabling Manual Mode

To disable manual mode (enable bridge learning), perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>disable manual-mode</code> .
2	Press Return. The following message is displayed: <code>Port numbers [0-254]?</code> The value displayed for port number is based on the ports that are currently enabled or disabled. For example, because you do not need to disable ports that are already disabled, the enable command displays only the ports that are enabled.
3	Enter the number of the port for which you want to disable manual mode (enable bridge learning).
4	Press Return. Manual mode is disabled on the specified port and the <code>BRIDGE config></code> prompt is displayed.
5	Repeat steps 1 through 4 for each port on which you want to disable manual mode.

Enabling and Disabling a Bridge Port

You can enable or disable selected bridge ports, if the ports are already added. You may want to disable a bridge port to help isolate network problems, for example. Disabling a port turns off bridging services on the port but does not disable its associated logical interface. Although a disabled port does not forward frames, it continues to receive certain STP packets. All bridge ports are enabled (the default) at startup.

Enabling a Port

To enable a bridge port, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter enable port .
2	Press Return. The following message is displayed: <code>Port Number [0]?</code>
3	Enter the number of the bridge port you want to enable.
4	Press Return. The selected port is enabled and the <code>BRIDGE config></code> prompt is displayed.

Disabling a Port

To disable a bridge port, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter disable port .
2	Press Return. The following message is displayed: <code>Port Number [0]?</code>
3	Enter the number of the bridge port you want to disable.
4	Press Return. The selected port is disabled and the <code>BRIDGE config></code> prompt is displayed.

Bridging Ethernet and FDDI Networks

When bridging an Ethernet LAN to an FDDI backbone, frames must be repackaged to resolve differences in frame formats between the two networks. VNswitch 900EF modules repackage the frames automatically using standard encapsulation and translation. They also implement IP fragmentation automatically. However, you may need to enable or disable IPX translation on a bridge if the bridge is to handle datagrams transmitted by network nodes using IPX protocols.

IP Fragmentation

IP fragmentation is used to “cut up” (fragment) large data packets into frame sizes that can be handled by Ethernet or Fast Ethernet LAN segments, or an Ethernet or Fast Ethernet ATM Emulated LAN or Bridge Tunnel. This may occur, for example, if IP datagrams traverse multiple LAN segments to reach their destination. If the first LAN segment the datagram enters is an FDDI or ATM segment, larger IP datagrams are accommodated by FDDI and ATM frames. If subsequent hops include an Ethernet or Fast Ethernet LAN segment, however, the FDDI or ATM frames may then be too large for the smaller Ethernet frames. IP fragmentation is used to resolve this incompatibility. Destination devices reassemble the fragmented packets when the packets are received. You cannot disable IP fragmentation.

Condition Under Which Packets Are Discarded

IP packets cannot be fragmented if their Don't Fragment Bit (DF Bit) is set. If a packet's DF Bit is set, the packet is discarded and an ICMP message is sent notifying the source node of the condition.

Note

An ICMP message is sent only if an IP address has been configured.

Enabling and Disabling IPX Translation

Internetwork Packet Exchange (IPX) is a peer-to-peer networking protocol for Novell NetWare. You may need to enable IPX translation on a VNswitch module under the following circumstances:

- Some nodes on your network are set to generate and receive IPX datagrams in Raw Ethernet frame format
- Some nodes in the same NetWare network number as above are directly connected to an FDDI LAN

Bridging Ethernet and FDDI Networks

Translation is necessary under the above conditions because the FDDI-attached NetWare node cannot recognize raw frames, while the Ethernet-attached nodes are configured to recognize only raw frames. For all these nodes to communicate, the VNswitch modules can be configured to translate between raw format (on the Ethernet LAN) and SNAP format (on the FDDI LAN).

Note

IPX translation incurs extra overhead when forwarding frames. If possible, reconfigure the Ethernet-attached nodes in this scenario to use SNAP encapsulation so that IPX translation is not needed. Note also that if all the nodes in the network number use raw encapsulation, then IPX translation is not needed since Raw frames can be sent over the FDDI network.

Translation Rules

To determine whether IPX translation should be enabled or disabled on a VNswitch 900EF module, you must understand the translation rules listed in [Table 7-4](#). The examples in the table refer to [Figure 7-1](#).

Table 7-4: IPX Translation Rules

If . . .	And . . .	Then . . .
The switch receives a Raw Ethernet frame from an Ethernet LAN (switch 1)	IPX translation is <i>enabled</i> on the switch	The Raw Ethernet frames are translated and are output to the FDDI backbone as SNAP frames.
The switch receives a Raw Ethernet frame from an Ethernet LAN (switch 1)	IPX translation is <i>disabled</i> on the switch	The Raw Ethernet frames are <i>not</i> translated and are output to the FDDI backbone as Raw Ethernet frames.
The switch receives a SNAP frame from an FDDI backbone (switch 2)	IPX translation is <i>enabled</i> on the switch	The SNAP frames are translated and are output to the Ethernet LAN (LAN B) as Raw Ethernet frames.

Bridging Ethernet and FDDI Networks

If. . .	And. . .	Then. . .
The switch receives a SNAP frame from an FDDI backbone (switch 2)	IPX translation is <i>disabled</i> on the switch	The SNAP frames are <i>not</i> translated to the Raw Ethernet frame type and are output to the Ethernet LAN (LAN B) as SNAP frames.

Bridging Ethernet and FDDI Networks

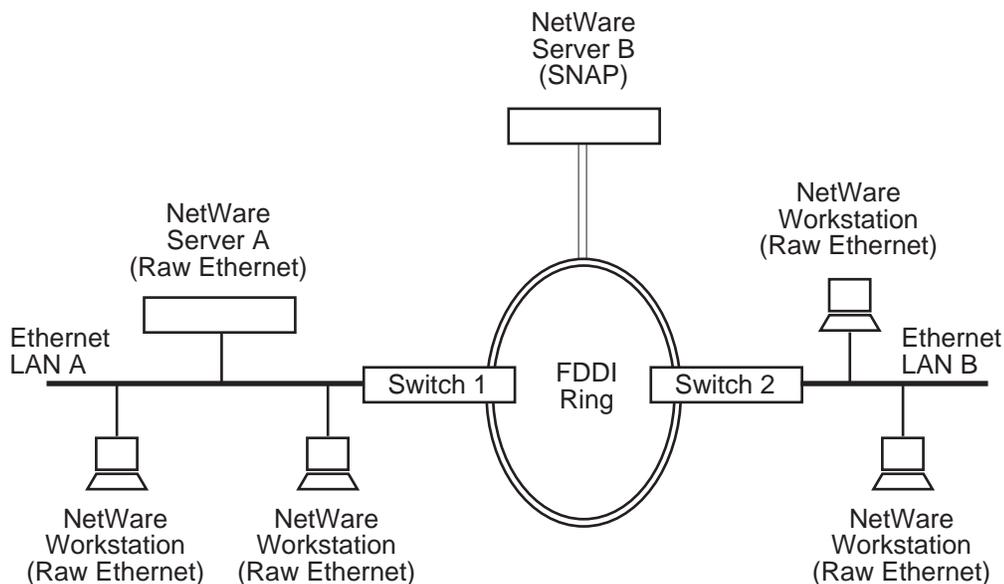
Examples

The following examples show how you can use the translation rules to accommodate several network configurations. [Figure 7-1](#) is used to illustrate the examples.

- NetWare server A on LAN A must exchange data with the Novell workstations on LAN B. Because the nodes on both LAN A and B generate and receive only Raw Ethernet frames, IPX translation does not need to be enabled on either switch 1 or switch 2.
- NetWare server B, directly connected to the FDDI backbone, must exchange data with the Novell workstations on LAN B. Because the server is set to handle SNAP frames, it can place data on the backbone without modification. However, in order for the workstations on LAN B to receive the SNAP frames, the frames must be translated back to Raw Ethernet format by switch 2. IPX translation must, therefore, be enabled on switch 2.

Conversely, in order for the workstations on LAN B to place their Raw Ethernet frames on the FDDI backbone, switch 2 must have IPX translation enabled. The switch can then translate the frames between raw format and SNAP format.

Figure 7-1: VNswitch 900EF Modules Requiring IPX Translation



LKG-10264-96F

NetWare Ethernet Frame Type Options

The NetWare operating system can be set to encapsulate IPX datagrams into one of several Ethernet frame types, two of which are Raw Ethernet and SNAP. What frame type you can set on the NetWare node varies, depending on the type of network to which the node is connected:

If . . .	Then . . .
The NetWare node is directly connected to an Ethernet LAN	The NetWare operating system can be configured to encapsulate IPX datagrams into any one of the following four Ethernet frame types: <ul style="list-style-type: none"> • Ethernet-II • Raw Ethernet • Subnetwork Access Protocol (SNAP) • Destination Service Access Point (DSAP)
The NetWare node is directly connected to an FDDI backbone	The NetWare operating system can be configured to encapsulate IPX datagrams into any one of the following two Ethernet frame types: <ul style="list-style-type: none"> • Subnetwork Access Protocol (SNAP) • Destination Service Access Point (DSAP)

Note

DIGITAL recommends that IPX network nodes, connected directly to Ethernet LANs, use the Ethernet-II frame format for full connectivity of IPX stations across Ethernet networks and FDDI bridged networks. If Ethernet-II is used by all IPX network nodes connected to an Ethernet LAN, IPX translation does not need to be enabled on any VNswitch 900EF module.

Bridging Ethernet and FDDI Networks

Enabling IPX Translation

To enable IPX translation, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>enable ipx-translation</code> .
2	Press Return. IPX translation is enabled on the VNswitch module.

Disabling IPX Translation

To disable IPX translation, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>disable ipx-translation</code> .
2	Press Return. IPX translation is disabled on the VNswitch module.

Auto-Testing of Ports Inactive for Extended Periods

A failure can sometimes cause a port to lose the ability to receive frames although it can still transmit frames. STP is unable to detect this condition. Such a failure can result in a forwarding loop and severe network congestion.

You can set the maximum amount of time a port operates without receiving a frame, although it may be transmitting frames. This setting, referred to as the No Frame Interval, is expressed in seconds and is applied to all ports on a switch. If this time is exceeded, the module tests the port for correct operation. If the self-test determines a problem exists, the port and STP are disabled. (Refer to the [Displaying Current Bridge Configuration Parameters](#) section for information about displaying a change in port and STP status. Refer to [Chapter 5](#) for information about displaying a change in the status of the logical interface.) All three conditions are logged to the Event Logging System. If no problem is found during the self-test, the No Frame Interval clock is restarted.

To set or modify the No Frame Interval, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter set no-frame-interval [<i>interval in seconds</i>], where <i>interval in seconds</i> is the maximum number of seconds the port is to remain inactive (no frames are received by, or transmitted from, the port).
2	Enter the maximum number of seconds (from 0 to 1000) the port is to remain inactive before a test is run. The default is 300 seconds.
3	Press Return. The No Frame Interval you entered is set for all ports on the switch.

Setting the Time That Unused Addresses Are Retained

A port's forwarding database stores Dynamic (learned) MAC addresses for a predetermined period of time, referred to as the Aging Time. If the switch does not receive a datagram from the device associated with that address within the period specified by the Aging Time, the address is "unlearned" by the port. If the switch subsequently receives data with that destination address, the data is forwarded through all ports, except the one on which it was received. This process is referred to as "flooding."

You can set or modify the Aging Time for all ports on a switch. To do so, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>set age</code> .
2	Press Return. The following message is displayed: <code>Seconds [300] ?</code>
3	Enter the Aging Time in seconds (from 10 through 1,000,000). The default is 300 seconds.
4	Press Return. The Aging Time is set.

Displaying Current Bridge Configuration Parameters

You can display specific information about various bridge configuration settings. To display configuration information, perform the following steps:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <i>list configuration</i> , where <i>configuration</i> is the command you must enter to display the desired information. Refer to Table 7-5 for a list of commands and for a description of the information that is displayed when you enter the command.
2	Press Return. The desired information and a new <code>BRIDGE config></code> prompt is displayed.

Example

BRIDGE config>list bridge

```
Transparent Bridge Configuration
Bridge: Enabled Bridge Behavior: STB
Rate Limiting: Enabled Rate Limit value: 400
Raw 802.3 IPX Trans: Disabled No Frame Interval: 300
Bridge Database Information
Ageing Time: 300 Entries in Perm Database: 0
Number of Reshases: 0
Spanning Tree Protocol Information
Bridge Priority: 1500/0x5DC
Port Information
Number of ports: 14
Port Interface State Behaviour STP Manual Mode
0 0 Enabled STB Only Enabled Disabled
1 1 Enabled STB Only Enabled Disabled
2 2 Enabled STB Only Enabled Disabled
3 3 Enabled STB Only Enabled Disabled
4 4 Enabled STB Only Enabled Disabled
5 5 Enabled STB Only Enabled Disabled
6 6 Enabled STB Only Enabled Disabled
7 7 Enabled STB Only Enabled Disabled
8 8 Disabled STB Only Enabled Disabled
9 9 Enabled STB Only Enabled Disabled
10 10 Enabled STB Only Enabled Disabled
11 11 Enabled STB Only Enabled Disabled
12 12 Enabled STB Only Enabled Disabled
13 13 Enabled STB Only Enabled Disabled
```

Table 7-5: List Command Options and Descriptions

Command	Description
address	The following information is listed for any permanent or static address you specify: <ul style="list-style-type: none">• Whether Rate Limiting is enabled for a specific MAC address• The address type (permanent, static, dynamic, and so on) for the MAC address• The allowed (forwarding) ports for the address

Displaying Current Bridge Configuration Parameters

Command	Description
<u>bridge</u>	<p>The following information is listed about the transparent bridge:</p> <ul style="list-style-type: none">• Bridge state (enabled or disabled)• Bridge behavior (Spanning Tree Bridge only)• Rate limiting state and value• IPX translation state• No Frame Interval• Ageing Time• Number of entries in permanent database• Number of rehashes (the number of times the module's programmable randomizer function was reprogrammed to more efficiently access the address and protocol databases used by the lookup engine)• Bridge priority and STP standard• Number of ports, and a list of port numbers and their associated interface numbers• Port states and whether STP and Manual Mode is enabled on each port
<u>default-protocol-filter</u>	<p>A list of default protocol filters in use and a list of their forwarding ports</p>
<u>port</u>	<p>The following information is listed for the port you specify:</p> <ul style="list-style-type: none">• ID and state of the port• Whether STP is enabled• Type of functionality supported by the port (Transparent Bridging only)• Whether Manual Mode is enabled• The number of the logical interface associated with the port, and its type (Ethernet, FDDI, and so on)• Path Cost of the port

Displaying Current Bridge Configuration Parameters

Command	Description
<u>protocol-filter</u>	<p>The following information is listed about protocol filters:</p> <ul style="list-style-type: none">• List of protocol filter classes in use• List of protocol filter classes in use, including type• Whether Rate Limiting is enabled• The ports on which the protocol is forwarded
<u>range</u>	<p>The following information is listed for each address within the range of permanent or static MAC addresses you specify:</p> <ul style="list-style-type: none">• Whether Rate Limiting is enabled• The address type (permanent, static, dynamic, and so on) for the range of MAC addresses• The allowed ports for the addresses
<u>stp</u>	<p>The following information is listed about STP:</p> <ul style="list-style-type: none">• STP's Bridge Max Age time• STP's Bridge Hello Time• STP's Bridge Forward Delay

Duplicate MAC Addresses on Separate VSDs

In most networks, the same MAC address is not expected to appear as a Source Address (SA) on more than one VLAN Secure Domain (VSD). Exceptions do exist, for example:

- A DECnet router can be attached to multiple VSDs to perform routing between those VSDs. DECnet routers force a phase IV-style derived MAC address on all the router's interfaces. This MAC address then appears as a duplicate on each VSD where the router has an interface.
- Sun systems with multiple interfaces use the same MAC address on all interfaces. However, you can configure these systems to use a unique MAC address on each interface.
- Any address can appear transiently as a duplicate if the address moves from one VSD to another. A VNswitch module sees this transient duplicate address when it receives packets from such an SA over the VNbus.

Before Version 3.0, the VNswitch learned the SA on the VSD where it first appeared and forwarded packets with that address as a Destination Address (DA). Packets containing the same address as a DA received on other VSDs were flooded. This presented a problem for VSDs containing LANE ATM ports because flooded packets are transmitted on LANE ATM ports via the BUS circuit with the rate limited to 2 packets per second.

With this release of the firmware, the VNswitch learns MAC addresses separately in each created VSD; that is, each VSD represents a separate learning domain. As a result, the VNswitch correctly forwards packets with the same DA received on multiple VSDs and flooding is not required. LANE ATM ports transmit these packets with the ATM line rate.

DIGITAL still recommends that you manually configure duplicate MAC addresses because it may help to avoid reachability problems in cases where the VNswitch has complex protocol filtering configuration and the duplicate address is learned via the VNbus port in multiple VSDs, or both. When you manually configure a duplicate address, you can specify the set of permitted ports on which the address is allowed to be received and transmitted.

Duplicate MAC Addresses on Separate VSDs

Configuring Duplicate MAC Addresses

Permanent Duplicate MAC Addresses

You configure permanent duplicate MAC addresses from the `BRIDGE config>` prompt. To set a permanent duplicate MAC address, perform the following tasks:

Step	Action
1	At the <code>BRIDGE config></code> prompt, enter <code>set duplicate-address <i>address</i></code> where <i>address</i> is the 12-digit MAC address. You can enter the address with or without hyphens separating the octets. For example, you can enter 11-22-33-44-55-66 or 112233445566 .
2	Press Return. The following message is displayed: <code>Enter allowed ports, "None", or "All" []?</code>
3	Enter the ports on which this address is reachable.

Static Duplicate MAC Addresses

You configure a static duplicate MAC address from the `BRIDGE>` prompt. To set a static duplicate MAC address, perform the following tasks:

Step	Action
1	At the <code>BRIDGE></code> prompt, enter <code>set duplicate-address <i>address</i></code> where <i>address</i> is the 12-digit MAC address. You can enter the address with or without hyphens separating the octets. For example, you can enter 11-22-33-44-55-66 or 112233445566 .
2	Press Return. The following message is displayed: <code>Enter allowed ports, "None", or "All" []?</code>
3	Enter the ports on which this address is reachable.

Duplicate MAC Addresses on Separate VSDs

Displaying Duplicate MAC Addresses

To display all user-configured duplicate MAC addresses, perform the following tasks:

Step	Action
1	At the BRIDGE> prompt, enter list database duplicate
2	Press Return. All duplicate MAC addresses are displayed.

Note

This command lists only static duplicate addresses that have been manually configured; it does not list learned duplicate addresses.

Example

```
BRIDGE>list database duplicate
```

```
MAC Address      Multi   Rate      Last  
Cast*   Limit  Entry Type  Port   Port   Seen   Learned  Allowed  
Port(s)  
08-00-09-00-4A-E8  N/A    Stat Dupl                All
```

Duplicate MAC Addresses on Separate VSDs

Deleting A Duplicate MAC Address

Permanent Duplicate MAC Addresses

To delete a permanent duplicate MAC address, perform the following tasks:

Step	Action
1	At the BRIDGE config> prompt, enter <u>delete address address</u> where <i>address</i> is the 12-digit MAC address. You can enter the address with or without hyphens separating the octets. For example, you can enter 11-22-33-44-55-66 or 112233445566 .
2	Press Return. The duplicate MAC address is deleted.

Static Duplicate MAC Addresses

To delete a static duplicate MAC address, perform the following steps:

Step	Action
1	At the BRIDGE > prompt, enter <u>delete address</u> where <i>address</i> is the 12-digit MAC address. You can enter the address with or without hyphens separating the octets. For example, you can enter 11-22-33-44-55-66 or 112233445566 .
2	Press Return. The duplicate MAC address is deleted.

Chapter 8

Monitoring the Transparent Bridge

Overview

Introduction

This chapter provides information about how to monitor the DIGITAL VNswitch module transparent bridge and how to configure static addresses. Monitoring includes the ability to display specific information about operational states, activity counters, and various bridge configuration settings. Static addresses you configure are stored in volatile RAM and are lost when the switch is powered down or restarted. The addresses, therefore, must be relearned when power is again applied to the switch.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Accessing and Exiting the Bridge Monitor Prompt	8-2
Monitoring the Bridge	8-3
Configuring a MAC Address As a Static Entry	8-18

Accessing and Exiting the Bridge Monitor Prompt

You must access the Bridge monitor prompt to monitor bridge parameters and to configure static MAC addresses. The prompt is accessed from the Monitor prompt (`Monitor>`).

Accessing the Bridge Monitor Prompt

To access the Bridge monitor prompt, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter bridge . Example: <code>Monitor>bridge</code>
2	Press Return. The Bridge monitor prompt (<code>BRIDGE></code>) is displayed.

Exiting the Bridge Monitor Prompt

You exit the Bridge monitor prompt to return to the Monitor prompt.

Example

To return to the Monitor prompt (`Monitor>`) from the `BRIDGE>` prompt, enter **exit** and then press Return.

Monitoring the Bridge

You can monitor the bridge by displaying such information as operational states, activity counters, and various configuration settings associated with the following bridge software components:

- General bridging operation
- Bridge ports
- MAC address database
- Protocol filters
- Spanning Tree Protocol

General Bridging Operation

You can display the following types of information about the bridge:

- Operational states and port profile
- Aging parameters

To display a profile of the bridge's ports, aging parameters, and information about the operational state of various bridge functions such as IP fragmentation and rate limiting, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter list bridge .
2	Press Return. The desired information and a new BRIDGE> prompt is displayed.

Monitoring the Bridge

Example

BRIDGE>**list bridge**

```
Bridge state:           Enabled
Multicast Rate Limiting: Disabled (400 pps)
No Frame Interval:     300 seconds
Raw 802.3 IPX Trans:   Disabled
Bridge type:           STB
Ageing time:           300 seconds
Number of rehashes:    0
Number of static entries: 0
Number of dynamic entries: 2308
Number of ports:       14
```

Port	Interface	Operational		MAC Address	Maximum	Flags
		State			PDU	
0	VNbus/0	Up		00-00-F8-6F-51-01	4491	MD
1	Eth/1	Down		00-00-F8-6F-51-02		MD
2	Eth/2	Down		00-00-F8-6F-51-03		MD
3	Eth/3	Down		00-00-F8-6F-51-04		MD
4	Eth/4	Down		00-00-F8-6F-51-05		MD
5	Eth/5	Down		00-00-F8-6F-51-06		MD
6	Eth/6	Down		00-00-F8-6F-51-07		MD
7	Eth/7	Down		00-00-F8-6F-51-08		MD
8	Eth/8	Down		00-00-F8-6F-51-09		MD
9	Eth/9	Down		00-00-F8-6F-51-0A		MD
10	Eth/10	Down		00-00-F8-6F-51-0B		MD
11	Eth/11	Down		00-00-F8-6F-51-0C		MD
12	Eth/12	Down		00-00-F8-6F-51-0D		MD
13	FDDI/13	Down		00-00-F8-6F-51-0E		MD

Flags: ME = Manual Mode Enabled, MD = Manual Mode Disabled

Port Activity Counters

To display information about port activity, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter list counters ports , where <i>ports</i> is the option you must enter to display the desired information. Refer to Table 8-1 for a list of options and for descriptions of the information that is displayed when you enter the option.
2	Press Return. The desired information and a new BRIDGE> prompt is displayed.

Table 8-1: List Counters Command Options and Descriptions

Command Option	Description
<u>a</u>ll-ports	Lists port activity counters for all ports.
<u>p</u>ort <i>port-number</i>	Lists port activity counters for a specific port, where <i>port-number</i> is the number of the port for which you want to display information.
<u>s</u>ummary	Displays the sum of activity across all ports for each counter.

ExampleBRIDGE>**list counters summary**

```

Port restarts:                               1
Total frames received by interfaces:         748388
IP frames fragmented:                        0
IP frames not fragmented:                   0
Frames submitted to bridging:                748388
Frames submitted to routing:                 154831
Frames with unknown dest address:           0
Frames causing learning transactions:        10078
Dropped, source address filtering:           14
Dropped, dest address filtering:             96473
Dropped, protocol filtering:                 0
Dropped, address rate limiting:              0
Dropped, protocol rate limiting:            0
Dropped, no buffer available:                0
Dropped, input queue overflow:               0
Dropped, source or dest port blocked:        0
Dropped, terminating queue overflow:         14
Dropped, fragmentation queue overflow:       0
Dropped, translate flood queue overflow:     0
Dropped, translation failure:                0
Frames sent by bridging:                     102077
Dropped, transmit queue overflow:            0
Dropped, transmit error:                    0
Dropped, too big to send on port:           0

```

Monitoring the Bridge

Bridge Ports

You can display the following information for one or all bridge ports:

- Port ID
- Whether STP and Manual Mode are enabled
- Name and number of the interface associated with the port
- State of the port
- For the VNbus port, state information is also provided for each instance of STP running over the VNbus through the VNbus port

If the port is a member of multiple VLAN Secure Domains (VSDs), the state of each VSD is listed, followed by the number and name of the VSD. The VNbus port is an example of such a port.

To display information about one or all bridge ports, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter list port port-number , where <i>port-number</i> is either the number of the port for which you want to display information or all for information about all ports.
2	Press Return. The desired information and a new BRIDGE> prompt is displayed.

Examples

The first example shows the types of information displayed for any port, other than the VNbus port. The second example shows the types of information displayed for the VNbus port, including state information for each instance of STP running over the VNbus and through the VNbus port. Each instance of STP indicates that one or more VSDs, each of which might be physically located on a different module in the same hub, are linked into a single, larger VSD. The link is established over the VNbus on the DEChub 900. Refer to [Chapter 9](#) for additional information about VSDs linked over the VNbus.

```
BRIDGE>list port 13
Port Id (dec)      : 128:13, (hex): 80-0D
Port State        : Enabled
STP Participation: Enabled
Manual Mode       : Disabled
Assoc Interface   : 13 (FDDI/13)
```

Monitoring the Bridge

BRIDGE>**list port 0**

Port Id (dec) : 128: 0, (hex): 80-00
STP Participation: Enabled
Manual Mode : Disabled
Assoc Interface : 0 (VNbus/0)

Port state per STP instance:

VSD#	Name	State
1	DEFAULT	Forwarding
2	TEST	Forwarding

Monitoring the Bridge

MAC Address Database

You can display the following information for each MAC address in the filtering database:

- Whether Rate Limiting is enabled for an address
- Whether the address is a multicast address (an asterisk indicates it is a multicast address)
- Address entry type (dynamic, reserved, registered, and so on)
- Last port on which the address was seen (dynamic, unicast static, and unicast permanent addresses only)
- Port on which the address was learned (dynamic, unicast static, and unicast permanent addresses only)
- Ports on which the address is allowed (permanent and static addresses only)

You can display this information selectively for addresses of a particular entry type (permanent, static, registered, and so on), for a range of addresses, or on a per-port basis.

To display information about selected portions of the database, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter list database <i>selected-information</i> , where <i>selected-information</i> is the option you must enter to display the desired information. Refer to Table 8-2 for a list of options and for descriptions of the type of information that is displayed when you enter the option.
2	Press Return. The desired information and a new BRIDGE> prompt is displayed.

Table 8-2: List Database Command Options and Descriptions

Command Option	Description
<u>all</u>	Lists information for all addresses associated with all ports.
<u>dynamic</u>	Lists information for all dynamic entries in the database.
<u>local</u>	Lists information for all registered entries in the database.
<u>permanent</u>	Lists information for all permanent entries in the database.
<u>port</u>	Lists information for all addresses associated with the specified port.
<u>range</u>	Lists information for all addresses within a specified range.
<u>static</u>	Lists information for all static entries in the database.
<u>atm</u>	Lists information for all the LEC learned entries in the database.

Monitoring the Bridge

Example

```
BRIDGE>list database local
```

MAC Address	Multi Cast*	Rate Limit	Entry Type	Last Seen Port	Learned Port	Allowed Port(s)
00-00-F8-6F-51-00		N/A	Registered			All
00-00-F8-6F-51-01		N/A	Registered			0
00-00-F8-6F-51-02		N/A	Registered			1
00-00-F8-6F-51-03		N/A	Registered			2
00-00-F8-6F-51-04		N/A	Registered			3
00-00-F8-6F-51-05		N/A	Registered			4
00-00-F8-6F-51-06		N/A	Registered			5
00-00-F8-6F-51-07		N/A	Registered			6
00-00-F8-6F-51-08		N/A	Registered			7
00-00-F8-6F-51-09		N/A	Registered			8
00-00-F8-6F-51-0A		N/A	Registered			9
00-00-F8-6F-51-0B		N/A	Registered			10
00-00-F8-6F-51-0C		N/A	Registered			11
00-00-F8-6F-51-0D		N/A	Registered			12
00-00-F8-6F-51-0E		N/A	Registered			13
00-00-F8-6F-51-20		N/A	Registered			All
00-00-F8-6F-51-22		N/A	Registered			All
01-80-C2-00-00-00*	Disabled		Registered			All
FF-FF-FF-FF-FF-FF*	Disabled		Registered			All

Protocol Filters

You can display the types of protocol filters and the default protocol filter applied to one or more ports. Protocol filtering is supported for the following frame types:

- Ethernet-II
- Subnetwork Access Protocol (SNAP)
- Destination Service Access Point (DSAP)

Displaying Protocol Filters

To display the protocol filters applied to one or more ports, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter list protocol-filter frame-type , where <i>frame-type</i> is the option you must enter to display the desired information. Refer to Table 8-3 for a list of options and for descriptions of the information that is displayed when you enter the command.

Step	Action
2	Press Return. The desired information and a new BRIDGE> prompt is displayed.

Table 8-3: List Protocol Filter Command Options and Descriptions

Command Option	Description
all	Lists all the protocol filters in use, and the ports on which the protocols are filtered. The filters are listed by frame type (Ethernet-II, SNAP, and SNAP) and hexadecimal value. Refer to Table 8-4 , Table 8-5 , and Table 8-6 for a list of hexadecimal values for common protocols.
ethertype protocol hexadecimal value	Lists the ports on which the specified protocol is filtered, where <i>protocol hexadecimal value</i> is the hexadecimal value of the protocol in which you are interested. Refer to Table 8-4 for a list of hexadecimal values for common Ethernet-II protocols. Enter 0 , the default, to display all Ethernet-II protocols that are filtered.
dsap protocol hexadecimal value	Lists the ports on which the specified protocol is filtered, where <i>protocol hexadecimal value</i> is the hexadecimal value of the protocol in which you are interested. Refer to Table 8-5 for a list of hexadecimal values for common DSAP (also referred to as SAP) protocols. Enter 100 , the default, to display all SAP (DSAP) protocols that are filtered.
snap protocol hexadecimal value	Lists the ports on which the specified protocol is filtered, where <i>protocol hexadecimal value</i> is the 10-digit hexadecimal value of the protocol in which you are interested. Refer to Table 8-6 for a list of hexadecimal values for common SNAP protocols. Enter 00-00-00-00-00 , the default, to display all SNAP protocols that are filtered.

Example

Monitoring the Bridge

```
BRIDGE>list protocol-filter all
```

Destination SAP	Rate Limit	Port(s)
01	Disabled	2-3
Ethernet type	Rate Limit	Port(s)
0800	Enabled	4-8
No SNAP filters configured		

Table 8-4: Hexadecimal Values for Common Ethernet-II Protocols

Protocol	Hexadecimal Value
IP	0800
ARP	0806
CHAOS	0804
DECnet MOP Dump/Load	6000
DECnet MOP Remote Console	6002
DECnet	6003
DEC LAT	6004
DEC LAVC	6007
XNS	0600
Maintenance Packet Type	7030
Apollo Domain	8019 (Ethernet)
Novell NetWare IPX	8137 (Ethernet)
AppleTalk Phase 1	809B
AppleTalk ARP Phase 1	80F3
Loopback assistance	9000

Table 8-5: Hexadecimal Values for Common DSAP Protocols

Protocol	Hexadecimal Value
Banyan SAP	BC (used for only 802.5)
Novell IPX SAP	E0 (used for only 802.5)
NetBIOS SAP	F0
ISO Connectionless Internet	FE

Monitoring the Bridge

Table 8-6: Hexadecimal Values for Common SNAP OUI/IP Protocols

Protocol	Ten-Digit Hexadecimal Value
AppleTalk Phase 2	08-00-07-80-9B
AppleTalk ARP Phase 2	00-00-00-80-F3
Proprietary AppleTalk Phase 1 for FDDI	00-00-93-00-02
Proprietary AppleTalk ARP Phase 1 for FDDI	00-00-93-00-03

Displaying Default Protocol Filters

You can also display default protocol filters used to filter *all* frames of a particular type from one or more output bridge ports, using protocols not already filtered as described above. For example, the module can be configured so that Ethernet frames conforming to the AppleTalk Phase 1 protocol are discarded, while a default filter is used to discard all Ethernet frames of any other protocol type such as IP, ARP, DECnet, and so on.

To display the default protocol filters applied to one or more ports, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter list default-protocol-filter .
2	Press Return. The desired information and a new BRIDGE> prompt is displayed.

Example

```
BRIDGE>list default protocol-filter
Protocol      Allowed Ports
DSAP         0-254
ETHER        0-254
SNAP         0-254
```

Spanning Tree Protocol

You can display the following information about STP:

- Configuration parameters such as the Bridge Hello Time, port Priorities and Costs, the number of VSDs on a port, and whether STP is enabled on a port
- Activity data such as the number of times the network topology has changed and the number of BPDUs sent and received
- Whether STP is active on a given port
- Information about the designated root and designated bridge for each port

To display information about the Spanning Tree Protocol, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter list stp option , where <i>option</i> is the option you must enter to display the desired information. Refer to Table 8-7 for a list of options and for descriptions of the type of information that is displayed when you enter the option.
2	Press Return. The desired information and a new BRIDGE> prompt is displayed.

Example

BRIDGE>**list stp counters 1**

```
VSD 3 Math-dept
Time since topology change (seconds):      248262
Topology changes:                          0
BPDUs received:                            3308
BPDUs sent:                                0
Port  Interface  BPDUs received  BPDU input overflow  Forward transitions
  0   VNbus/0    3308           0                     1
 10   Eth/10     0              0                     0
```

Table 8-7: List STP Command Options and Descriptions

Command Option	Description
<u>configuration</u>	<p>Lists the following STP parameters configured on the switch:</p> <ul style="list-style-type: none"> • Bridge Maximum Age • Bridge Hello Time • Bridge Forward Delay • Hold Time (This value is fixed at one second and is not configurable, as required by IEEE standard 802.1D.) • Interface type and number, Priority, Cost, Administrative State, and number of VSDs on each port
<u>counters</u> <i>vsd#-or-name</i>	<p>Lists the following information about STP activity, where <i>vsd#-or-name</i> is the number or name of the VSD for which you want the information displayed. If no number or name is specified, information for all VSDs is displayed.</p> <ul style="list-style-type: none"> • Time, in seconds, since the current topology change within the VSD was detected. This value is zero when there is no topology change in effect. • Number of changes in network topology within the VSD since the VSD was created • Number of BPDUs received and sent by the VSD • Number of BPDUs received, BPDU input overflow, and forward transitions for each port in the VSD

Command Option	Description
<code>state vsd#-or-name</code>	<p data-bbox="732 590 1224 762">Lists the following information about STP activity, where <i>vsd#-or-name</i> is the number or name of the VSD for which you want the information displayed. If no number or name is specified, information for all VSDs is displayed.</p> <ul data-bbox="732 779 1224 1486" style="list-style-type: none"> <li data-bbox="732 779 1224 856">• ID (a combination of the bridge priority and address) for the instance of STP in the VSD <li data-bbox="732 873 1224 951">• ID (a combination of the bridge priority and address) of the root bridge in the VSD's spanning tree <li data-bbox="732 968 1224 1045">• Path Cost of the root port in the VSD's spanning tree <li data-bbox="732 1062 1224 1140">• Which port is the root port on the VSD's spanning tree <li data-bbox="732 1157 1224 1234">• Current Bridge Max Age, Hello Time, and Forward delay dictated by the root bridge and used by all bridges in the VSD's spanning tree <li data-bbox="732 1251 1224 1329">• Whether the VSD's spanning tree has detected (True or False) a topology change <li data-bbox="732 1346 1224 1423">• Whether the root bridge has confirmed (True or False) a topology change in the VSD's spanning tree <li data-bbox="732 1440 1224 1486">• State (Forwarding, Blocking, Listening, Learning, or Down) of each port in the VSD
<code>tree vsd#-or-name</code>	<p data-bbox="732 1518 1224 1690">Lists the following information for each port in the VSD's spanning tree, where <i>vsd#-or-name</i> is the number or name of the VSD for which you want the information displayed. If no number or name is specified, information for all VSDs is displayed.</p> <ul data-bbox="732 1707 1224 1885" style="list-style-type: none"> <li data-bbox="732 1707 1224 1736">• Associated interface number <li data-bbox="732 1745 1224 1774">• Designated root <li data-bbox="732 1782 1224 1812">• Designated cost <li data-bbox="732 1820 1224 1850">• Designated bridge <li data-bbox="732 1858 1224 1885">• Designated port

Configuring a MAC Address As a Static Entry

This section discusses how to add MAC address filters as static entries. This section also describes how to modify and delete static entries. Because static addresses are stored in volatile RAM and are lost when the switch is powered down or restarted, you may find the ability to add static entries to be most useful when trying to isolate network problems. (Refer to [Chapter 1](#) for information about volatile RAM.) Static entries are not affected by address Aging Time. (Refer to [Chapter 7](#) for information about address aging.)

The process of filtering network traffic is used to reduce the amount of unnecessary traffic over specific segments of a network, thereby maximizing network capacity and performance. It can also be used to restrict the distribution of sensitive information to specific locations. You can configure the VNSwitch module to selectively filter or forward packets it receives, based on their MAC addresses and protocol types.

You configure a filter by specifying a set of ports that are allowed for a given MAC address. The address can be an individual, multicast, or broadcast address. If a packet has a source or destination MAC address that is listed in the filter, the packet is allowed on the specified set of ports. If the address is not listed in the filter, the packet is dropped (filtered).

Source Address Filtering

A packet received on an input port is dropped (filtered) if the source address is listed in a filter and the input port is not one of the allowed ports. In addition, the source address is not learned.

If the input port is one of the allowed ports, the packet is a candidate for forwarding, based on the behavior of destination address filtering and protocol filtering. (Refer to [Chapter 7](#) for information about protocol filtering.)

Destination Address Filtering

A packet about to be placed on an output port is forwarded if the destination address is listed in a filter and the output port is one of the allowed ports. If the address was not previously learned, the packet is flooded to the subset of the allowed ports that are in the forwarding state.

Configuring a MAC Address As a Static Entry

Creating and Modifying a Static MAC Address Filter

To create or modify a static MAC address filter, perform the following tasks:

Task	Description
1	Identify the address to be filtered and the ports on which the address is allowed.
2	Enable or disable multicast rate limiting.

Identifying the Address and Allowed Ports

The addresses you specify when creating a filter are added to the switch's static database (volatile RAM). To identify the allowed ports for an address, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter set static-address . You can, alternatively, enter set static-address mac_address , where <i>mac_address</i> is the 12-digit MAC address for which you want to create or modify a filter on the module. If you enter the command using this syntax, and go to step 4.
2	Press Return. The following message is displayed: Address (in 12-digit hex) []?
3	Enter the 12-digit MAC address for which you want to create or modify a static filter. You can enter the address with or without hyphens separating the octets. For example, you can enter 11-22-33-44-55-66 or 112233445566 .
4	Press Return. The following message is displayed: Enter allowed ports, ("None", or "All") []? If you are modifying the forwarding ports for an existing address, the previously configured ports are the default.

Configuring a MAC Address As a Static Entry

Step	Action
5	<p>Enter a list of the bridge ports to which you want the packet with the specified MAC address forwarded. (The packet is filtered from those ports not entered.) Enter None if you want the packet filtered from all ports. Enter all if you want the packet allowed on all ports. You can list the forwarding ports individually by separating each with a comma (for example, 0,1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 0-5). You can also combine a list of individual entries with a range of entries (for example, 0-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.</p>
6	<p>Press Return.</p> <p>If the address is a unicast address, the address is set and the <code>BRIDGE></code> prompt is displayed.</p> <p>If the address is a multicast address, the following message is displayed:</p> <pre>Enable Multicast Rate Limiting (Yes or No)? [No]:</pre> <p>Refer to the Enabling and Disabling Multicast Rate Limiting section.</p> <p>If the specified ports do not exist, the address is set for forwarding on the port but the port setting has no effect. A message similar to the following is displayed in addition to the Bridge config or rate limiting prompts, whichever is appropriate:</p> <pre>Warning, ports 9, 10 do not exist</pre>

Enabling and Disabling Multicast Rate Limiting

Multicast rate limiting is used to minimize the effects of multicast storms. You can restrict multicast storms, consisting of packets that have a specific destination MAC address, to that segment of the network from which the packets are generated. You do so by setting the maximum number of those packets the VNswitch module is to forward per second, and by enabling rate limiting.

Rate limiting is enabled on both a per-address basis and a switchwide basis. This section provides instructions about how to enable or disable rate limiting on a per static MAC address basis. Refer to [Chapter 7](#) for information about setting the maximum number of packets per second and enabling or disabling rate limiting on a switchwide basis.

Configuring a MAC Address As a Static Entry

To enable or disable multicast rate limiting for the static address specified in the [Identifying the Address and Allowed Ports](#) section, perform the following steps:

Step	Action
1	Enter Yes to enable Multicast Rate Limiting. Enter No to disable Multicast Rate Limiting. No is the default unless it was previously enabled for the address.
2	Press Return. The selected filter is applied to the module and rate limiting is enabled or disabled, as specified.

Deleting Static MAC Address Filters

This section describes how to delete static MAC addresses that you, or another switch administrator, previously entered manually. (Refer to the [Creating and Modifying a Static MAC Address Filter](#) section.) Deleting an address filter removes the address from the forwarding database as well as its associated multicast rate limiting status.

Note

Permanent, registered, and reserved addresses cannot be deleted.

To delete a static MAC address, its filter, and its associated rate limiting records, perform the following steps:

Step	Action
1	At the BRIDGE> prompt, enter delete mac_address . You can, alternatively, enter delete address , where <i>address</i> is the 12-digit MAC address you want to delete. If you enter the command using this syntax, and go to step 4.
2	Press Return. The following message is displayed: Address (in 12-digit hex) []?
3	Enter the 12-digit MAC address you want to delete.
4	Press Return. The specified address is deleted. If you entered an address that does not exist, the following message is displayed: No entry found for this address.

Chapter 9

Configuring Virtual LANs

Overview

Introduction

A VLAN is a group of bridge ports logically linked to define a LAN on one or more hubs. This network configuration scheme enables you to configure a set of devices so they logically appear to be on the same LAN segment, although they may be physically on different segments. You can create a maximum of 32 VLANs per DIGITAL VNswitch module.

Assume, for example, that two of five members in a bank's loan department are located in different areas of a building. Further, the two individuals are using terminals that are physically connected to an ATM emulated LAN, while the other three workers' terminals are physically connected to a segment in an Ethernet LAN. You can configure the switch software so that all five terminals are logically on the same LAN, thereby minimizing traffic from the loan department over those segments of the network outside the loan department's VLAN. The five terminals can be physically connected to either the same VNswitch module (900AE, in this case), or to two or more VNswitch modules on the same or different hubs. If any individual is later assigned to another department, the switch software can be dynamically reconfigured so the user's terminal is no longer associated with the loan department's VLAN, but is associated with a different VLAN, if desired.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
VLAN Secure Domains	9-3
Accessing and Exiting the VSD Configuration Prompt	9-5
Creating VSDs	9-6
Modifying VSDs	9-12

Topic	Page
Deleting VSDs	9-15
Displaying Information About VSDs	9-16
Configuring the VNbus to Support VLANs	9-18
Assigning a VNswitch IP End Node to a VSD	9-20

VLAN Secure Domains

A VLAN Secure Domain (VSD) is a logical set of one or more VLANs that operates with one spanning tree. The resulting configuration is a set of distinct bridge ports isolated from other ports on the same switch by blocking all unicast and multicast traffic between VSDs. The VNswitch 900 series of modules presently supports one VLAN per VSD, but the VSD concept provides for expanded support of multiple VLANs within a single VSD.

VNbus

The Virtual Network bus (VNbus) is a set of three dynamically configured buses on the backplane of the MultiSwitch 900. Each operates at 400 Mb/s for a total available capacity of 1.2 Gb/s. The VNbus is the medium used to pass Virtual LAN traffic between modules, when you configure a VLAN across two or more modules in the same hub.

VNbus Tags

A VNbus tag is a unique number used to link two or more VSDs, each of which is physically located on a different module in the same hub. Assigning the same VNbus tag number to the VSDs effectively links the individual VSDs (on different modules) into a single, larger VSD. The link is established over the VNbus on the MultiSwitch 900.

A VSD remains local to the module on which its ports reside, if no VNbus tag is assigned.

Default VSD

All bridge ports on a VNswitch module are, by default, members of a default VSD. The default VSD is numbered VSD 1 and is assigned the name "DEFAULT." The VNbus tag number for the default VSD is 65. The number, name, and VNbus tag cannot be changed.

When more than one VNswitch modules is resident in a MultiSwitch 900, all ports on all VNswitch modules are, by default, members of the same (default) VSD. Ports that are members of the default VSD operate as a traditional bridge without VLANs, except that the bridge spans the VNbus. When you create a new VSD as described in the [Creating VSDs](#) section, the ports you assign as members of the new VSD are removed from the default VSD. Conversely, when a VSD is deleted, all ports that were members of the VSD automatically become members of the default VSD.

VLAN Secure Domains

Spanning Tree Protocol Support

The Spanning Tree Protocol (STP) algorithm runs automatically for each VSD you create. Therefore, multiple spanning trees are supported on the bridge. STP inhibits loops in redundant bridges by assigning a redundant bridge as a backup.

Accessing and Exiting the VSD Configuration Prompt

You must access the VSD configuration prompt to create and manage VSDs and to display configuration information about VSDs.

Accessing the VSD Configuration Prompt

To access the VSD configuration prompt, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter vlan s.
2	Press Return. The <code>VSD Config></code> prompt is displayed.

Exiting the VSD Configuration Prompt

You exit the VSD prompt to return to the Bridge configuration prompt. For example, to return to the Bridge configuration prompt (`Bridge Config>`) from the `VSD Config>` prompt, enter **exit** and then press Return.

Creating VSDs

How you create a VSD varies, depending on which of the following port groupings is planned:

- The VSD is composed of a port group located on a single module.
- The VSD is composed of multiple port groups, each of which is located on a different module in the same hub.
- The VSD includes ATM Emulated LANs or Bridge Tunnels.

You use the **create vsd** command to create a VSD, followed by one or more of the command options shown in [Table 9-1](#).

Table 9-1: Create VSD Command Options

Command Option	Description
ports <i>port-list</i>	<p>The <i>port-list</i> is a list of the bridge ports you want to assign to the VSD. You can list the ports individually by separating each with a comma (for example, 1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 1-5). You can also combine a list of individual entries with a range of entries (for example, 1-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.</p> <p>Assigning a port-list is optional.</p>
name <i>name</i>	<p>The <i>name</i> is the name you want to assign to the VSD (Math_dept, for example). A name can be composed of up to 32 ASCII-printable characters, except a question mark (?), must not include spaces, and must contain at least one alphabetic character. The name assigned must be unique. You cannot create another VSD on the same module using the same name. Assigning a name to a VSD is optional.</p>

Creating VSDs

Command Option	Description
tag <i>vnbus-tag</i>	A VNbus tag is a unique number used to link two or more VSDs, each of which is physically located on a different module in the same hub. Assigning the same VNbus tag number to the VSDs effectively links the individual VSDs (on different modules) into a single, larger VSD. The variable <i>vnbus-tag</i> can be any number from 66 through 128. To connect VSDs on different hub modules, assign the same VNbus tag number to each. For example, if you want to make VSD 3 on a VNswitch 900AE module and VSD 2 on a VNswitch 900EE module members of the same VSD, assign the same VNbus tag number (66, for example) to both VSDs.

Creating VSDs

VSDs Within a Single Module

You create a VSD within a single module by identifying those ports that are members of the VSD. The switch then assigns a number to each VSD you create. The default VSD always retains the VSD number of 1, even if all ports are reassigned to newly created VSDs. You can also, optionally, assign a name to the VSD. Assigning a name (for example, **Math_dept**) to the VSD provides a way of more easily recognizing functional groupings of ports, rather than simply using the system-assigned VSD number.

To create a VSD within a single module, perform the following steps:

Step	Action
1	Access the VSD configuration (VSD Config>) prompt. (Refer to the Accessing the VSD Configuration Prompt section for instructions.)
2	If you want to create a VSD without assigning it a name, enter create vsd ports port-list . If you want to create a VSD and assign it a name, enter create vsd ports port-list name name . Refer to Table 9-1 for a description of the options you can enter.
3	Press Return. The VSD you configured is created, including any name you assigned. A message is displayed showing the number assigned by the switch to the VSD. You can view the list of assigned VSD numbers at any time. Refer to the Displaying Information About VSDs section for instructions about how to do so.

VSDs Across Modules in a Single Hub

You can configure multiple VSDs located on different modules in the same hub, so that they are members of the same VSD. You do so by assigning a number, called a VNbus tag, to each VSD that is to be a member of the larger VSD. The logical VSD connection that results from the assigned VNbus tag occurs over the hub's 400 Mb/s VNbus.

Assume, for example, that a VNswitch 900AE module and a VNswitch 900EE module are both resident on a single MultiSwitch 900. The 900AE module includes VSD 1 named **Geology_dept** and VSD 2 named **Math_dept**. The 900EE module includes VSD 1 named **Math_dept** and VSD 2 named **Chemistry_dept**. To make VSD 2 (**Math_dept**) on the 900AE module and VSD 1 (**Math_dept**) on the 900EE module members of the same VSD, you assign the same VNbus tag number (66, for example) to both.

Creating VSDs

To create a VSD across modules in a single hub, perform the following steps:

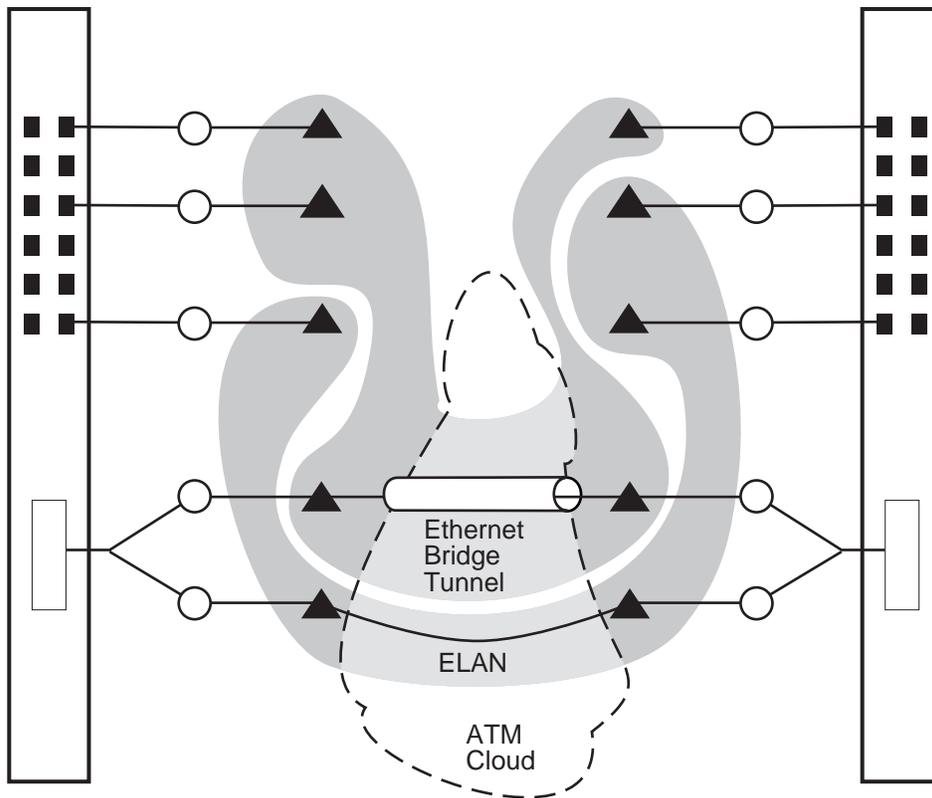
Step	Action
1	Access the VSD configuration (<code>VSD Config></code>) prompt.
2	If you want to create a VSD without assigning it a name, enter <code>create vsd ports port-list tag vnbus-tag</code> . If you want to create a VSD and assign it a name, enter <code>create vsd ports port-list name name tag vnbus-tag</code> . Refer to Table 9-1 for a description of the options you can enter.
3	Press Return. The VSD you configured is created, including any name and VNbus tag you assigned. A message is displayed showing the number assigned by the switch to the VSD on the module. You can view the list of assigned VSD numbers at any time. Refer to the Displaying Information About VSDs section for instructions about how to do so.
4	If you want to link another group of ports located on a different module to the VSD created in step 3, log in to the next module and repeat steps 2 through 4. If you do not want to link another group of ports to the VSD, exit from the VSD configuration prompt.
5	Configure the VNbus to support VLANs. Refer to the Configuring the VNbus to Support VLANs section for information about how to do so.

VSDs Across ATM Emulated LANs and Bridge Tunnels

You can configure VSDs across ATM ports located on different modules in different hubs. You do so by providing a logical connection from the bridge port associated with an ATM LEC interface or ATM bridge tunnel that is a member of one module's VSD, to the bridge port associated with an ATM LEC interface or ATM bridge tunnel that is a member of a second module's VSD.

Creating VSDs

Figure 9-1: ATM ELAN and Bridge Tunnel VSDs



▲ = Bridge Port

○ = Logical Interface

LKG-10307-96F

To create a VSD across ATM ports located on different modules in different hubs, you must first configure either an ATM LEC logical interface or an ATM bridge tunnel logical interface. (Refer to [Chapter 5](#) for information about how to do so.) You then create a VSD according to the instructions in the [VSDs Within a Single Module](#) section earlier in this chapter, including the LEC bridge port number or bridge tunnel port number that is to be included in the VSD.

Note

If you also want the VSD to span across modules in the same hub, you can add the **tag vnbus-tag** option to the **create vsd** command. For example, **create vsd ports port-list tag vnbus-tag**. Refer to the [VSDs Across Modules in a Single Hub](#) section earlier in this chapter.

Assume, for example, that two VNswitch 900AE modules are resident on different MultiSwitch 900s. An ELAN is configured from ATM LEC interface number 16 (port 17) on one module to ATM LEC interface number 19 (port 20) on the second module. (Refer to [Chapter 5](#) for information about configuring ATM LEC and bridge tunnel logical interfaces.) The first module includes VSD 3 named **Math_dept**. The second module includes VSD 2, also named **Math_dept**. To make VSD 3 (**Math_dept**) on the first module and VSD 2 (**Math_dept**) on the second module members of the same VSD, ATM LEC port number 17 is added to VSD 3 on the first module, and ATM LEC port number 20 is added to VSD 2 on the second module.

Reserving VSDs

You reserve a VLAN Secure Domain by creating a VSD without assigning it ports. You may want to do so if, for example, certain work groups or portions of the network are not ready to go on line, while others are. You can later specify what ports are members of the VSD when the work group or network segments are ready.

To reserve a VLAN Secure Domain, create a VSD according to the instructions in this section, but do not use the **ports port-list** or **tag vnbus-tag** options to assign ports or tags. The VSD you configure is created, including any name you assigned. A message is displayed showing the number assigned by the switch to the VSD.

You can view a list of assigned VSD numbers, names, and tags at any time, and can later assign ports and tags to the VSD. Refer to the [Displaying Information About VSDs](#) section for instructions about how to view VSD numbers, names, and tags. Refer to the [Modifying VSDs](#) section for instructions about how to add ports to reserved VSDs.

Modifying VSDs

You can modify a VSD's name, list of assigned ports, or VNbus tag. You may want to change assigned ports or a VNbus tag if, for example:

- Nodes logically associated with one VSD are moved and require new port assignments
- Some members of a work group or department associated with one VSD are moved and should be reassigned to another VSD

To modify a VSD's name, list of ports, or VNbus tag, perform the following steps:

Step	Action
1	Access the VSD configuration (VSD <code>Config></code>) prompt. (Refer to the Accessing and Exiting the VSD Configuration Prompt section for instructions.)
2	<p>If you want to modify a VSD's name, enter: <code>modify ysd number name new-name</code></p> <p>If you want to modify the ports assigned to a VSD, enter: <code>modify ysd number ports new-port-list</code></p> <p>If you want to modify the VNbus tag assigned to a VSD, enter: <code>modify ysd number tag new-vnbus-tag</code></p> <p>If you want to modify the VSD using more than one command option (for example the name, port list, and VNbus tag), enter: <code>modify ysd number name new-name ports new-port-list tag new-vnbus-tag</code></p> <p>Refer to Table 9-2 for a description of the options you can enter.</p>
3	Press Return. The VSD is modified according to the changes you entered. You can view the list of assigned VSD numbers at any time. Refer to the Displaying Information About VSDs section for instructions about how to do so.

Table 9-2: Modify VSD Command Options

Command Option	Description
<i>ports new-port-list</i>	<p>The <i>new-port-list</i> is the new list of bridge ports you want to assign to the VSD. You can list the ports individually by separating each with a comma (for example, 1,2,3,4,5) or you can enter a range of ports by separating the first and last port numbers in the range with a hyphen (for example, 1-5). You can also combine a list of individual entries with a range of entries (for example, 1-4,6,7). Spaces are not permitted between port numbers and the comma or hyphen.</p>
<i>name new-name</i>	<p>The <i>new-name</i> is the new name you want to assign to the VSD (changing Math_dept to Advanced_Math_dept, for example). A name can be composed of up to 32 ASCII-printable characters, except a question mark (?), must not include spaces, and must contain at least one alphabetic character. The name assigned must be unique. You cannot create another VSD on the same module using the same name. Assigning a name to a VSD is optional.</p>

Modifying VSDs

Command Option	Description
<code>tag new-vnbus-tag</code>	<p>A VNbus tag is a unique number used to link two or more VSDs, each of which is physically located on a different module in the same hub. Assigning the same VNbus tag number to the VSDs effectively links the individual VSDs (on different modules) into a single, larger VSD. The variable <i>vnbus-tag</i> can be any number from 66 through 128. To connect VSDs on different hub modules, assign the same VNbus tag number to each.</p> <p>Modifying the VNbus tag for a VSD can be used to change the membership of the ports to that of another VSD that spans more than one module on the same hub. Assume, for example, that VSD 3 on a VNswitch 900AE module and VSD 2 on a VNswitch 900EE module are members of the same VSD because they are both assigned the VNbus tag 66. The 900AE module has a second VSD (VSD 4) that is assigned a VNbus tag of 67. You can change the membership of VSD 2 on the 900EE module by modifying its VNbus tag to 67, so that it is now a member of the same VSD as VSD 4 on the 900AE module.</p>

Deleting VSDs

You can delete VSDs one at a time, or you can delete all VSDs on a particular VNswitch module at the same time. When a VSD is deleted, all ports that were members of the VSD automatically become members of the default VSD. You cannot delete the default VSD.

To delete one or all VSDs, perform the following steps:

Step	Action
1	Access the VSD configuration (VSD Config>) prompt. (Refer to the Accessing and Exiting the VSD Configuration Prompt section or instructions.)
2	If you want to delete a single VSD, enter delete vsd number or delete vsd name , where <i>number</i> is the number of the VSD you want to delete and <i>name</i> is the name of the VSD you want to delete. (Refer to the Displaying Information About VSDs section for instructions about how to list the numbers and names assigned to existing VSDs. If you want to delete all VSDs on the module, enter delete all .)
3	Press Return. The VSDs you entered are deleted and the ports that are members of the deleted VSD become members of the default VSD.

Displaying Information About VSDs

You can display the following configuration information for either a specific VSD or for all VSDs:

- VSD number
- VSD name
- List of ports that are members of a VSD on a single module
- VNbus tag associated with the VSD
- Warning messages

To display VSD configuration information, perform the following steps:

Step	Action
1	Access the VSD configuration (VSD Config>) prompt. (Refer to the Accessing and Exiting the VSD Configuration Prompt section for instructions.)
2	If you want to display information about a single VSD, enter list vsd number or list vsd name , where <i>number</i> is the number of the VSD for which you want to display information and <i>name</i> is the name of the VSD for which you want to display information. If you want to display information about all VSDs on the module, enter list all .
3	Press Return. Information about the specified VSDs is displayed.

Examples

```
VSD>list vsd 3
```

```
VSD Number:      3
VSD Name:        Math-Dept
VNbus Tag:       98
```

```
Port    Assoc Interface #/Name
  5      5/Eth/5
 12     12/Eth/12
 14     14/ATMLEC/14
```

Displaying Information About VSDs

VSD Config>**list all**

VSD Name	Ports	VNbus Tag	Ifc
1 DEFAULT	4,5,7-12	65	14
2 vsd1	1,2,6,13		15
3 vsd2	3		16

Warning Messages

The following warning messages are also displayed under certain conditions when you use the **list** command:

Warning Message	Description
TCP/IP-Host's affiliated VSD has no ports	This message is displayed when the VSD associated with TCP/IP Host Services currently has no ports in it. When the VSD associated with IP Host has no ports in it, there is no remote access to TCP/IP Host Services. Refer to the Assigning a VNswitch IP End Node to a VSD section for additional information about the relationship between VSDs and TCP/IP Host Services. Refer to the Creating VSDs section for information about how to add ports to an existing VSD.
VNbus unavailable -- VSDs do not span the hub	<p>This message is displayed when the VNswitch module is installed in a MultiSwitch 900 and the module is not connected to the VNbus. The message is a reminder that VLANs are not currently spanning the hub even if VNbus tags are assigned to the VSDs.</p> <p>One reason the VNbus may not be connected to the module is that the module's VNbus interface may be disabled. Refer to Chapter 5 for information about how to enable and disable an interface.</p>
Module not in Hub -- VNbus tags not used	This message is displayed when the VNswitch module is installed in a DEChub ONE docking station. The message is a reminder that VNbus tags are not used in a docking station and that the values in the VNbus column can be ignored. The VNbus tags in the VNbus column are only applicable if the module is installed in a MultiSwitch 900.

Configuring the VNbus to Support VLANs

If you plan to create a VLAN across two modules in the same hub, you may have to configure the VNbus to establish a LAN segment on the bus. You may also have to enable the VNbus port on those modules where the port is disabled. Because of certain default settings and automated configuration functions, you need to configure the VNbus only under the following conditions:

If. . .	Then. . .
The MultiSwitch 900 Auto-LAN feature is disabled	You must establish a LAN segment on the VNbus manually.
The VNbus port is disabled by you or another switch administrator	You must enable the VNbus port.

Establishing a LAN Segment on the VNbus

If you plan to create a VLAN across two modules in the same hub, an internal LAN segment over the VNbus must first be created. A LAN segment is established after any two modules are installed on the MultiSwitch 900, and can be created either automatically by the hub's Auto-LAN feature, or manually using the DIGITAL clearVISN MultiChassis Manager. Refer to the appropriate MultiSwitch 900 documentation for information about how to enable or disable the Auto-LAN feature. Refer to the appropriate MultiChassis Manager documentation for information about how to create an internal LAN manually.

Note

The VNbus uses only a portion of the MultiSwitch 900 backplane's flexible channels. The remaining channels on the backplane are used when creating additional LAN segments in the hub over which modules can exchange data. For example, two Ethernet or two FDDI modules can be connected over the backplane instead of over an external LAN. Such an internal LAN configuration can be reconfigured through software only, without changing hardware connections as would be required if the connection is through an external LAN. This capability is provided through the optional MultiChassis Manager, a component of the DIGITAL clearVISN network management product. Refer to the appropriate MultiChassis Manager documentation for information about creating and deleting LAN segments over the hub backplane, and how to create and delete connections to the segments.

Disabling and Enabling a VNbus Port

VNswitch modules connect to the VNbus through the VNbus port on each module's back panel. You can disable and enable the VNbus port from your terminal console. You may want to do so if, for example, you need to temporarily inhibit traffic flow between two segments of a VLAN separated by two modules. The VNbus tag used to create a VLAN across multiple modules remains unaltered if you disable a VNbus port. Therefore, when you are ready to resume traffic flow across a VLAN between the modules, you do not need to reapply the VNbus tags, you need only enable the VNbus port.

The default setting for the VNbus port is enabled. Therefore, when you create a VLAN across modules in a hub, you do not need to enable the VNbus port unless it is first disabled by either you or another switch administrator.

Disabling the Port

To disable the VNbus port, perform the following steps:

Step	Action
1	At the <code>Bridge Config></code> prompt, enter disable port 0 .
2	Press Return. The module's VNbus port is disabled.

Enabling the Port

To enable the VNbus port, perform the following steps:

Step	Action
1	At the <code>Bridge Config></code> prompt, enter enable port 0 .
2	Press Return. The module's VNbus port is enabled.

Assigning a VNswitch IP End Node to a VSD

VNswitch modules can serve as IP end nodes in a network. As an IP end node, a module supports several IP protocols including, for example, Telnet, SNMP, TFTP, and Ping. Telnet and SNMP provide two vehicles for remotely managing modules. TFTP is used to backup and restore switch configurations and to load new software images. Ping is used to diagnose network problems. Refer to [Chapter 12](#) for information about using Telnet and SNMP. Refer to [Chapter 10](#) for information about backing up and restoring switch configurations.

When the VNswitch module serves as an IP end node, you must perform the following tasks before any of these IP protocols (Telnet, SNMP, and so on) can be used:

- Assign an IP address and subnet mask to the module
- Select the VSD on which you want a VNswitch IP end node to reside

Assume, for example, each of three VSDs represents a different IP subnet. You must select the VSD on which the IP end node is to reside. If you select VSD 1, then you must assign an IP address and subnet mask to the IP end node that is valid for VSD 1's subnet. IP nodes directly attached to VSD 1's subnet can initiate an ARP request for, and connect to, the module directly. IP nodes connected to other subnets (VSD 2, for example) must connect to the module through a router.

You do not need to assign an IP address to the VNswitch module or assign the address to a VSD if you do not want to use any of the IP end node protocols' features (remote management, network backup and restoration, and so on).

Restrictions

Assigning an IP host to a VSD affects only in-band management (IP packets received over Ethernet, Fast Ethernet, FDDI, or ATM network interfaces). IP packets received through the OBM port of a MultiSwitch 900 or a DEChub ONE are not affected. Packets received through an OBM port are not subject to VSD or bridge filtering.

Assigning an IP Address and Subnet Mask

Refer to [Chapter 12](#) for information about how to assign the module's TCP/IP host address.

Assigning a VNswitch IP End Node to a VSD

Selecting the VSD on Which the Host Resides

To associate an IP Host address with a specific VSD, perform the following steps:

CAUTION

You should assign the VNswitch module's IP host address and associate the host address with a particular VSD one after the other. If you change the IP address or the address-to-VSD association without also changing the other and then restart the module, you may lose inband connectivity to the module. This may occur because the IP host address and the VSD's subnet do not match. If you encounter this problem, you may be able to correct the mismatch only by using a local console (inband management), or by resetting the module to factory defaults.

Step	Action
1	Access the VSD configuration (VSD Config>) prompt. (Refer to the Accessing and Exiting the VSD Configuration Prompt section for instructions.)
2	Enter set vsd ip-host .
3	Press Return. The following message is displayed: VSD Number or Name (1): []?
4	Enter the number or name of the VSD to which you want to assign the module's IP address. The factory default is 1.
5	Press Return. The module's IP address is associated with the specified VSD.
6	Verify that the IP address assigned to the IP host matches the subnet of the VSD. If it does not match, you should change the IP host address before completing the next step.
7	Restart the module for the setting to take effect. Refer to Chapter 10 for information about how to do so.

Chapter 10

Performing Routine Maintenance

Overview

Introduction

This chapter describes DIGITAL VNswitch module maintenance procedures that you may need to perform periodically, and those that you should perform regularly.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Accessing and Exiting the Boot Config Prompt	10-2
Restarting the Switch	10-3
Upgrading and Reinstalling Switch Software	10-4
Backing Up and Restoring the Switch	10-13
Checking Available RAM	10-19
Capturing Restart or Crash Messages and Diagnostic Data	10-20
Displaying All Boot Config Settings	10-34

Accessing and Exiting the Boot Config Prompt

You must access the Boot config prompt (`Boot config>`) to configure and manage many of the parameters discussed in this chapter. The prompt is accessed from the Config prompt (`Config>`).

Accessing the Boot Config Prompt

To access the Boot config prompt, perform the following steps:

-
- 1 At the Config prompt, enter **boot**.
 - 2 Press Return. The following message is displayed, followed by the Boot config prompt:

```
TFTP Boot/dump configuration
Boot config>
```
-

Exiting the Boot Config Prompt

You exit the Boot config prompt to return to the Config prompt.

Example

To return to the Config prompt (`Config>`) from the `Boot config>` prompt, enter **exit** and then press Return.

Restarting the Switch

Certain switch configuration tasks require that you restart the switch for new configuration settings to take effect. The switch is also started when power is applied.

Restarting the switch causes reinitialization of the switch software, using the same executable and configuration it is currently running. The executable (also known as the boot or image file) and configuration parameters are stored in the switch's nonvolatile flash memory. Restarting the switch also clears information stored in volatile memory, including all bridge table entries not saved to NVRAM, and drops any packets in the bridge. Recovery of the lost packets is the responsibility of the sending and receiving nodes.

Note

Restarting a switch from a remote console terminates the console's Telnet session.

How to Restart the Switch

To restart the switch, perform the following steps:

Step	Action
1	At the Main prompt (Main>), enter restart .
2	Press Return. The following prompt is displayed: <code>Are you sure you want to restart the system? (Yes or [No]):</code> No is the default.
3	If you want to restart the switch, enter Yes and press Return. The switch is restarted. If you do not want to restart the switch, press Return. The Main prompt is again displayed.

Note

You can also restart the switch from the logon menu. Refer to [Chapter 2](#) for information about how to do so.

Upgrading and Reinstalling Switch Software

This section discusses the following topics:

- Installing the software
- Configuring installation file locations
- Viewing installation file locations
- Modifying installation file locations
- Canceling the installation procedure

An upgrade replaces the current version of software with a newer version. A reinstallation reloads the switch with a copy of the current version of the software. You may need to reinstall the software if, for example, you have a problem upgrading to a newer version and need to reinstall the earlier version until the problem is resolved.

The steps you complete to perform an upgrade or reinstallation are the same. Therefore, for the purpose of this discussion, the term installation is used to mean either an upgrade or a reinstallation.

Installing the Software

The source software you use for an installation must be stored on a network host or server. You may want to preconfigure the network location of the software prior to performing an upgrade or reinstallation. (Refer to the [Configuring Installation File Locations](#) section for instructions.) Although preconfiguring the network location is not required, it reduces the number of steps you perform when installing the software.

The steps you perform to install the software vary, depending on whether you preconfigured the network location of the software to be installed.

Using a Preconfigured Network File Location

To upgrade or reinstall the switch software, perform the following steps:

Step	Action
1	At the Main prompt (Main>), enter reload .
2	Press Return. The following prompt is displayed: Are you sure you want to reload the system (Yes or No)? No is the default.

Upgrading and Reinstalling Switch Software

Step	Action
3	<p>If you want to reload the switch software, enter Yes and press Return. The switch performs a restart, taking itself off the network, reinitializes itself, and reconnects to the network in IP Host-Only Mode. The switch software is copied from the preconfigured location and stored in NVRAM/flash memory. The following sequence of messages is displayed:</p> <pre>System restart... Copyright 1995-1996 Digital Equipment Corp. MOS Operator Control IP Host-Only Mode-Upgrading Operational Image... Network FLASH Upgrade Proceeding...</pre> <p>The installation is complete when the VNbus Traffic and VNbus Status LEDs flash alternating green and yellow for about 10 seconds, and then the LED pairs remain lit (either green or yellow) for another 10 seconds. This occurs about 45 seconds after the following message is displayed:</p> <pre>TFTP transfer complete: Status: OK.</pre> <p>After the installation or upgrade is successfully completed, hardware diagnostics are run. This takes about 1 minute.</p> <p>If the network location of the software to be installed is not preconfigured, the following message and the Main prompt is displayed:</p> <pre>Aborted, No boot entries defined. Configure a boot entry or use LOAD REMOTE.</pre> <p>If you do not want to reload the switch software, press Return. The Main prompt is again displayed.</p> <p>Note: Refer to Table 10-1 for a full description of the LED lighting sequences during an upgrade.</p>
4	<p>If you are performing an upgrade, update the software version number according to the instructions in the Configuring Installation File Locations section.</p>

Using an Unconfigured Network File Location

To upgrade or reinstall the switch software, perform the following steps:

Step	Action
1	At the <code>Boot config></code> prompt, enter load remote .
2	Press Return. The following message is displayed: <code>Remote Host Address [0.0.0.0]?</code>

Upgrading and Reinstalling Switch Software

Step	Action
3	Enter the IP address of the host or server on which the software to be installed is located. The default address is 0 . 0 . 0 . 0.
4	Press Return. The following message is displayed: Remote Pathname []?
5	Enter the path and file name that identifies where on the host or server the software is located. Example: /usr/tftp/dvnee206.bin
6	Press Return. The following message is displayed: First Hop Address [0.0.0.0]?
7	Enter the IP address of the first hop router that routes to other networks. The first hop router and its address are needed if the remote host on which the software is located is not on a directly connected network. The default address is 0 . 0 . 0 . 0.
8	Press Return. The following message is displayed: TFTP Timeout Value [10]?
9	Enter the desired TFTP timeout value. TFTP is the protocol the switch uses to download the installation file. If the TFTP request is lost, TFTP retransmits the request several times. The timeout value affects the amount of time that passes before the next retry. The default timeout value of 10 seconds is recommended. A value greater than 10 seconds should be used if the path over which the installation is to occur is typically slow.
10	Press Return. The following message is displayed: Are you sure you want to reload the system (Yes or No)? No is the default.

Upgrading and Reinstalling Switch Software

Step	Action
11	<p>If you want to reload the switch software, enter Yes and press Return. The switch performs a restart, taking itself off the network, reinitializes itself, and reconnects to the network in IP Host-Only Mode. The switch software is copied and stored in NVRAM/flash memory. If the module is in a DEChub ONE docking station, the following sequence of messages is displayed:</p> <pre>System restart... Copyright 1995-1996 Digital Equipment Corp. MOS Operator Control IP Host-Only Mode-Upgrading Operational Image... Network FLASH Upgrade Proceeding...</pre> <p>The installation is complete when the VNbus Traffic and VNbus Status LEDs flash alternating green and yellow for about 10 seconds, and then the LED pairs remain lit (either green or yellow) for another 10 seconds. This occurs about 45 seconds after the following message is displayed:</p> <pre>TFTP transfer complete: Status: OK.</pre> <p>After the installation or upgrade is successfully completed, hardware diagnostics are run. This takes about 1 minute.</p> <p>If you are upgrading a boot block, the yellow and green lights will remain lit until you power-cycle the board.</p> <p>If you do not want to reload the switch software, press Return. The Main prompt is again displayed.</p> <p>Note: Refer to Table 10-1 for a full description of the LED lighting sequences during an upgrade.</p>
12	<p>If you are performing an upgrade, update the software version number according to the instructions in the Configuring Installation File Locations section.</p>

LED Lighting Sequence During Load and Reload

The module's LEDs provide information regarding the progress of an installation or upgrade. If the TFTP file transfer is successful, a CRC of the image received at the module is performed. [Table 10-1](#) describes the LED lighting sequences and the status of the upgrade they reflect.

Table 10-1: LED Lighting Sequence During Load/Reload

LED Lighting Sequence	Upgrade Status
<ul style="list-style-type: none"> VNbus Traffic LED (upper): green VNbus Status (lower): yellow 	CRC is in progress.
<ul style="list-style-type: none"> VNbus Traffic LED (upper): yellow VNbus Status (lower): green 	CRC detects an error.
<ul style="list-style-type: none"> VNbus Traffic LED (upper): green VNbus Status (lower): green 	The image is being written to flash, following a successful CRC.
<ul style="list-style-type: none"> VNbus Traffic LED (upper) and VNbus Status LED (lower) both flash alternating green and yellow for about 10 seconds, and then remain solidly lit (either green or yellow) for another 10 seconds 	The process of writing the image to flash is complete.

Configuring Installation File Locations

To predefine the location of the file to be used for an upgrade or reinstallation, you specify the name of the installation file to be installed and the location of the server on which the file resides. You do not need to specify the VNswitch module interface for transparent bridging. The upgrade operation defaults to the IP-HST address or the appropriate IP address if routing is enabled.

Specifying the File Name and Server Location

To specify the name of the file to be installed and the location of the server on which the file resides, perform the following steps:

Step	Action
1	At the Boot config prompt, enter add boot-entry .
2	Press Return. The following message is displayed: remote host [0.0.0.0]?

Upgrading and Reinstalling Switch Software

Step	Action
3	Enter the IP address of the remote host on which the installation file resides. The default address is 0.0.0.0.
4	Press Return. The following message is displayed: <code>via gateway (0.0.0.0 if none) [0.0.0.0]?</code>
5	Enter the IP address of the first hop router, if any. The default address is 0.0.0.0. The default should be used if there is no first hop router.
6	Press Return. The following message is displayed: <code>timeout in seconds [10]?</code>
7	Enter the desired TFTP timeout value. TFTP is the protocol the switch uses to download the installation file. If the TFTP request is lost, TFTP retransmits the request several times. The timeout value affects the amount of time that passes before the next retry. The default timeout value of 10 seconds is recommended. A value greater than 10 seconds should be used if the path over which the installation is to occur is typically slow.
8	Press Return. The following message is displayed: <code>File name []?</code>
9	Enter a path and file name for the location on the remote server where the installation file is located.
10	Press Return. The specified values are set and the Boot config prompt is displayed.

Displaying File Names and Server Locations

Displaying:installation file locationsTo view a list of the files to be installed and the location of the server on which the files reside, perform the following steps:

Step	Action
1	At the Boot config prompt, enter list boot-entries .
2	Press Return. A reference number is displayed for each configured interface, followed by a colon. The reference number is then followed by the location of each file to be used for installation, including the path and name of the file on the remote system, the IP address of the remote host on which the installation file resides, the IP address of the first hop router (if any), and the TFTP retransmission timer value.

Upgrading and Reinstalling Switch Software

Example

```
Boot config>list boot-entries
```

```
Boot files:
```

```
1: "/usr/bt/dvnee154.bin" on 192.9.1.2 through 0.0.0.0 for 3 secs
2: "/usr/bt/dvnee154.bin" on 192.9.2.2 through 192.9.1.4 for 3 secs
```

Modifying Installation File Locations

Modifying predefined installation file locations used for an upgrade or reinstallation includes the ability to change and delete the currently configured location of the installation file.

Changing the File Name and Server Location

To change the name of the file to be installed or the location of the server on which the file resides, perform the following steps:

Step	Action
1	At the Boot config prompt, enter change boot-entry .
2	Press Return. The following message is displayed: Change which entry [1]?
3	Enter the index number associated with the entry you want to change. The default is 1. The index number is the first number, followed by a colon, in each item listed when you display information about the location of the installation file using the list boot-entries command. Refer to the Displaying File Names and Server Locations section for additional information.
4	Press Return. The following message is displayed: remote host [18.123.0.16]?
5	Enter the new IP address of the remote host on which the installation file resides. The previously configured address is the default.
6	Press Return. The following message is displayed: via gateway (0.0.0.0 if none) [0.0.0.0]?
7	Enter the IP address of the first hop router, if any. The default address is 0.0.0.0. The default should be used if there is no first hop router.
8	Press Return. The following message is displayed: timeout in seconds [10]?

Upgrading and Reinstalling Switch Software

Step	Action
9	Enter the desired TFTP timeout value. TFTP is the protocol the switch uses to download the installation file. If the TFTP request is lost, TFTP retransmits the request several times. The timeout value affects the amount of time that passes before the next retry. The default timeout value of 10 seconds is recommended. A value greater than 10 seconds should be used if the path over which the installation is to occur is typically slow.
10	Press Return. The following message is displayed: File name [user/lib/gw/gwimage.ldb]?
11	Enter the path and file name for the location on the remote server where the installation file is located.
12	Press Return. The specified values are set and the Boot config prompt is displayed.

Deleting a File Name and Server Location

To delete the name of the file to be installed and the location of the server on which the file resides, perform the following steps:

Step	Action
1	At the Boot config prompt, enter delete boot-entry .
2	Press Return. The following message is displayed: Delete which entry [1]?
3	Enter the index number associated with the entry you want to delete. The default is 1. The index number is the first number, followed by a colon, in each item listed when you display information about the location of the installation file using the list boot-entries command. Refer to the Displaying File Names and Server Locations section for additional information.
4	Press Return. The specified location is no longer configured for use during installations.
5	To verify that the entry is deleted, use the list boot-entries command. The index numbers of the remaining entries are renumbered to avoid leaving a gap in the list of entries. For example, if entry 2 is deleted, entry 3 becomes entry 2, and so on. Refer to the Displaying File Names and Server Locations section for additional information.

Upgrading and Reinstalling Switch Software

Canceling the Installation Procedure

You can cancel the installation procedure before writing the image to flash memory begins. You do so by pressing Ctrl/C. You cannot use Ctrl/C to cancel installation while the image is being written to flash.

Caution

Do not power cycle the module while the image is writing to flash. The module may become disabled.

Backing Up and Restoring the Switch

You can back up the switch's configuration settings. Switch configuration settings that are stored in NVRAM can be restored manually.

Automatic Image Recovery

The VNswitch 900 family of modules automatically tries to reload the switch's image if the currently installed image is corrupted. Although the current image is stored in NVRAM and does not need to be reloaded in the event of a power outage, the image can become corrupted due to an unusual event such as a power surge. The automatic download of a new image only occurs over the LDM port.

Location of Source Software

The source software used for automatic recovery must be stored on a network host or server. The location can be the same as that used for normal software upgrades or installations described in the [Upgrading and Reinstalling Switch Software](#) section. If the switch software image is corrupted, the switch must look to an external device to determine the location of the software it must reload. Such a device is referred to as a BootP server. In this discussion, the VNswitch module is referred to as a BootP client.

BootP Server and Clients

The BootP server contains a file that lists all the BootP clients for which the server is responsible, including the clients' IP addresses, and the locations and names of their boot files. This list of clients and boot file location information is maintained by the network administrator. If a switch's image is corrupted, the switch (BootP client) broadcasts a request to the BootP server. The request, in the form of a UDP packet, includes the client's MAC address. When the server receives the request, it looks up the client's address in its database of client information. If it locates the client's address in the database, the BootP server responds to the client, providing it with information about the location of the software it is to install. The switch then initiates a TFTP request for a download of the software from the boot server on which the software resides.

Note

The BootP server can use any BootP software available on a variety of operating systems.

Backing Up and Restoring the Switch

DIGITAL Online Services

VNswitch software is also available through the DIGITAL Network Products Home Page on the World Wide Web located at the following addresses:

North America	http://www.networks.digital.com
Europe	http://www.networks.europe.digital.com
Asia Pacific	http://www.networks.digital.com.au

Configuring Automatic Image Recovery

To configure your system for automatic recovery of the switch's executable software, you must configure the BootP server. You do not need to configure the BootP client (the VNswitch module). The image is automatically reloaded through the LDM port on the switch.

The network node you use as a BootP server can use any operating system that supports the BootP utility. You must properly configure the server to include a BootP database that contains a list of all the BootP clients (switches) for which the server is responsible, including the clients' IP addresses, and the locations and names of their boot files. Refer to the operating system documentation provided by the system vendor for configuration instructions.

Backing Up Configuration Settings

Default configuration settings and most of the settings you configure are stored in a configuration database in the VNswitch module's NVRAM. Although the configuration database does not need to be reloaded if there is a power outage, the database can become corrupted due to an unusual event such as a power surge. You can make a copy of the configuration database and back it up on a remote server or host. Should the switch's configuration database be corrupted, you can then copy the backup to the switch, minimizing the need to reconfigure custom settings you previously entered.

Use the Trivial File Transfer Protocol (TFTP) to back up the configuration database to a remote system, and to copy the backup to the switch if the original database on the module is corrupted. The VNswitch software includes a version of TFTP for this purpose.

Before You Begin

On some systems, you may need to create the file on the remote server or host to which the configuration database is backed up. This is the host file name you enter during the backup procedure specified in the [Backing Up the Configuration Database](#) section.

Backing Up and Restoring the Switch

You must also configure the switch's TCP/IP Host Services to include one IP address for the module, before attempting backup or restoral. Because the configured IP address is lost if NVRAM is corrupted, the address must be reconfigured before restoral. You can also reconfigure the switch's host name (optional). Refer to [Chapter 12](#) for information about configuring TCP/IP Host Services.

Backing Up the Configuration Database

To back up a switch's configuration database to a remote system, perform the following steps:

Step	Action
1	At the Boot config prompt, enter tftp put .
2	Press Return. The following message is displayed: local filename [CONFIG]? CONFIG is the default file name. It cannot be changed.
3	Press Return. The following message is displayed: remote host [0.0.0.0]?
4	Enter the IP address of the host to which you want to copy the configuration database. The default (0 . 0 . 0 . 0) is an invalid address, and is meant only to show the format of the address.
5	Press Return. A message similar to the following is displayed: host filename [0B070706.cfg]? The filename's prefix is derived from a portion of the module's IP address converted to ASCII.
6	Enter the path, followed by a unique file name, to identify the location on the remote system where you want to back up the configuration database. The default is 0B070706 .cfg. It is recommended that you give the file a name that is descriptive of the switch from which it originates. By doing so, you should more easily be able to distinguish between backup files derived from multiple modules. Example: /usr/local/tftp/switch11.cfg
7	Press Return. The following message is displayed when the transfer is completed: TFTP transfer complete, Status: OK Refer to the following table for a list of possible error messages and their meanings.

Backing Up and Restoring the Switch

The following error messages may be displayed if an error occurs during the transfer:

Error Message	Meaning
Unknown Error	Protocol failure
File Not Found	Specified host file does not exist
Access Violation	File protection error
Disk Full	File system full during write
Illegal Operation	Undefined TFTP operation requested
Unknown TID	Unexpected TFTP packet received
File Already Exists	File already exists
No Such User	TFTP not supported on host

Restoring a Configuration Database

The configuration database can be restored to a VNSwitch module of the same type (900EE, 900EX, and so on) from which it was backed up. The restoral process includes verification of a “magic number” that ensures the configuration database is being restored on a switch of the proper type. The message `Bad Magic Number` is displayed if you attempt to restore the database to a switch of the wrong type. The switch’s host name is also checked to make sure the configuration is restored on the correct switch. If the host name of the module to which the file is to be restored differs from that of the original switch, the message `Is this acceptable? (Yes or [No])` is displayed. No is the default.

Backing Up and Restoring the Switch

To restore a switch's configuration database by retrieving it from a remote system, perform the following steps:

Caution

Do not power off or reset the switch while restoring the configuration database. NVRAM may become corrupted, requiring you to reinstall the module's configuration database.

Step	Action
1	At the Boot config prompt, enter tftp get .
2	Press Return. The following message is displayed: Local filename [CONFIG]? config CONFIG is the default file name. It cannot be changed.
3	Press Return. The following message is displayed: Remote host [0.0.0.0]?
4	Enter the IP address of the host where the configuration database is backed up.
5	Press Return. The following message is displayed: Host filename []?
6	Enter the path and file name that identifies the location and configuration database you want to restore. Example: /usr/local/tftp/switch11.cfg
7	Press Return. The following message is displayed when the transfer is complete: TFTP transfer complete, status: OK An automatic RESTART will occur after writing CONFIG. Are you sure you want to write new CONFIG? (Yes or [No]): The following message is displayed if there is not enough memory to buffer the file on the switch: TFTP transfer complete, Status: Out of Memory

Backing Up and Restoring the Switch

Step	Action
8	Enter yes if you want to continue. Enter no if you want to abort the procedure.
9	If you entered yes , the following message is displayed: Updating CONFIG; Do Not Interrupt! If you enter no , the following message is displayed and the procedure is aborted: *** ERROR *** Write to Non-Volatile memory failed! CONFIG update aborted.

Note

You cannot transfer a configuration database using the remote system's version of TFTP-client.

Exiting and Canceling Backup or Restoration

The Boot config prompt is locked during backup or restoration of the configuration database. In order to:

- Exit the Boot config prompt and return to the Main prompt (Main>) while allowing the backup or restore to continue, press Ctrl/P.
- Cancel backup or restore, press Ctrl/C. The Boot config prompt is displayed.

Checking Available RAM

The VNswitch module's volatile memory is used to store information such as bridge tables and various buffers. A predefined amount of volatile memory is allocated from the switch's random access memory (RAM) at startup. If more volatile memory is required, additional RAM is variably allocated to volatile memory, as needed.

Refer to [Chapter 4](#) for more information about monitoring switch memory.

How to Check Available RAM

To check available RAM, perform the following steps:

Step	Action
1	At the Main prompt (Main>), enter memory .
2	Press Return. A message is displayed indicating the number of bytes in RAM that are busy, idle, and free.

Example

The following example shows how to check available RAM.

```
Main>memory
```

```
Number of bytes:  Busy = 28464,  Idle = 5632,  Free = 2926207
```

where,

Information Displayed	Description
Busy	Number of bytes in use (temporarily allocated) as volatile memory
Idle	Number of bytes previously allocated but freed and available for reuse
Free	Number of bytes never allocated from RAM

Note

The sum of Idle and Free memory equals the total available heap memory.

Capturing Restart or Crash Messages and Diagnostic Data

The VNswitch module logs messages and diagnostic data whenever the module is restarted or experiences a fatal error. You can display these messages to help troubleshoot problems and can download detailed diagnostic data for further problem analysis.

Displaying and Managing Restart or Crash Error Messages

The VNswitch module generates informational messages whenever it is restarted and a problem is encountered, and whenever a fatal error such as a bug halt occurs. The restart messages, generated by diagnostic routines during power-up, are recorded in a diagnostic log. Messages related to fatal errors (crashes) are recorded in a crash log. Both logs are stored in NVRAM and, therefore, the information they contain remains intact after a restart or power outage.

You display the messages in these logs to determine the results of diagnostics and to detect crashes that might occur when you are not present. (The occurrence of a crash may not be obvious because the switch typically restarts automatically.) However, a detailed interpretation of the messages is possible only with the assistance of DIGITAL Customer Services, who can assist troubleshooting problems you encounter. Each log retains up to 254 messages. If the maximum of 254 messages is reached, each new message overwrites (wraps) the oldest existing message. Refer to the [Determining When a Crash Occurs](#) section for detailed information about how to read the crash log to determine when a crash occurs.

You can also delete messages from the logs if you want to make detection of new messages easier. You can delete messages either one at a time or you can delete all messages from a log.

Displaying Diagnostic or Crash Log Messages

You can display diagnostic or crash log messages from either the Config prompt (`Config>`) or the Monitor prompt (`Monitor>`). To display diagnostic log or crash log messages, perform the following steps:

Step	Action
1	At either the Config prompt or the Monitor prompt, enter err-logs .
2	Press Return. The <code>Error-Log></code> prompt is displayed.
3	If you want to display the diagnostic log, enter list diagnostic-log . If you want to display the crash log, enter list crash-log .

Capturing Restart or Crash Messages and Diagnostic Data

Step	Action
4	Press Return. The list of messages contained in the specified log and the <code>ERROR-Log></code> prompt is displayed.

Determining When a Crash Occurs

You can find out when a crash occurs by determining when a message is logged in relation to other messages, and by inserting a date/time stamp or similar marker at specific points in the crash log.

Message Sequence

Two numbers are automatically assigned to each message, a crash log entry number and a historical message number. The crash log entry number is displayed on the line preceding a message. It is used to identify messages you want to delete. The historical message number is used to determine when a specific message is logged in relation to other messages. The historical message number is displayed on the same line as the message. Refer to [Figure 10-1](#) for an example.

The maximum number of crash log entries that can be recorded is 254. If this number is exceeded, the next message logged overwrites the oldest crash message starting at crash log entry number 1. However, when the historical message number reaches 254, it continues to increment by 1 to 255, 256, and so on as shown in [Figure 10-1](#), up to and including 9999. It is then reset to 1. Therefore, the crash message having the highest historical message number is the most recent message logged. Using [Figure 10-1](#) as an example, crash log entry number 5 with a historical entry number of 259 is a more recent entry than crash log entry number of 6 with a historical entry number of 6.

The crash log is cleared and the crash log entry number is reset to 1, if the total message record size is exceeded and a new message is logged. Because the length of each message is different, the speed with which the record becomes full varies, depending on the total number of characters of all messages logged.

If the time has been set using either the `Config>time set` or `Config>time host` commands, then a timestamp is saved with each crash message. If not, the uptime of the switch is saved.

Capturing Restart or Crash Messages and Diagnostic Data

Figure 10-1: Sample Crash Log

```
CRASH LOG ENTRY NUMBER 5 | Message 5 | Recorded (Uptime):  
0:01:48 (T2.0-022) Illegal Instruction at PC 0xA  
  
CRASH LOG ENTRY NUMBER 6 | Message 6 | Recorded: 1/5/  
1993 17:27:16 (T2.0-023) User induced bughlt via DVT
```

Crash Log Markers

You can insert a marker in the crash log so that you can identify when messages that follow a marker occur. The marker you enter is a free-form text entry. For example, if you enter the date and time, you will know that all messages that follow the marker occur after the specified date and time. If the marker you enter describes a particular system configuration change you make, you will know that all messages that follow the marker may be the result of the new configuration.

Markers appear as a message in the crash log, and are assigned a crash log entry number and historical message number. Refer to crash log entry number 8 in [Figure 10-1](#) for an example.

To enter a crash log marker, perform the following steps:

Step	Action
1	At the Config prompt, enter err-logs .
2	Press Return. The <code>Error-Log></code> prompt is displayed.
3	Enter writelog .
4	Press Return. The following message is displayed: Message [1-240 chars]?:
5	Enter the desired text.
6	Press Return. The text of the marker you enter is logged in the crash log and the <code>Error-Log></code> prompt is displayed.

Capturing Restart or Crash Messages and Diagnostic Data

Deleting All Diagnostic or Crash Log Messages

You can delete (clear) all messages from either the diagnostic or crash log, or from both logs. You can do so only from the Config prompt (`Config>`). The Monitor (`Monitor>`) prompt cannot be used to clear the logs. To delete all diagnostic log or crash log messages, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt, enter <code>err-logs</code> .
2	Press Return. The <code>Error-Log></code> prompt is displayed.
3	If you want to delete all messages from the diagnostic log, enter <code>clear diagnostic-log</code> . If you want to delete all messages from the crash log, enter <code>clear crash-log</code> . If you want to delete all messages from both logs, enter <code>clear all</code> .
4	Press Return. All messages in the specified log are deleted and the <code>Error-Log></code> prompt is displayed.

Deleting a Single Diagnostic or Crash Log Message

You can delete one message at a time from either the diagnostic or crash log. You can do so only from the Config prompt (`Config>`). The Monitor (`Monitor>`) prompt cannot be used to delete a message from the logs. To delete one message at a time from the diagnostic log or the crash log, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt enter <code>err-logs</code> .
2	Press Return. The <code>Error-Log></code> prompt is displayed.
3	If you want to delete a message from the diagnostic log, enter <code>delete diagnostic-log</code> . If you want to delete a message from the crash log, enter <code>delete crash-log</code> .

Capturing Restart or Crash Messages and Diagnostic Data

Step	Action
4	<p>Press Return.</p> <p>If you are deleting an entry from the diagnostic log, the following message is displayed:</p> <pre>diagnostic error-log entry number []?</pre> <p>If you are deleting an entry from the crash log, the following message is displayed:</p> <pre>Crash log entry number []?</pre>
5	<p>Enter the number of the message you want to delete. The crash log entry number is displayed on the line preceding a message. Refer to Figure 10-1 for examples of crash log entries and their associated numbers.</p> <p>Example: diagnostic error-log entry number []?235</p>
6	<p>Press Return. The error entry associated with the number you entered is deleted, the message Deleted Error Log Entry Number = 235 and the Error-Log> prompt are displayed.</p> <p>If the error entry you specified by number does not exist, the following message is displayed:</p> <pre>Error Log Entry Not Found</pre>

Downloading Diagnostic Data for Problem Analysis

You can capture a “snapshot” of a failed module’s status at the time of failure. This information can be used as a diagnostic tool by Digital Services representatives, should you require assistance.

Capture the diagnostic data by configuring your switch to automatically download (dump) the information to one or more host systems when the module is reset due to a hardware or software failure. The data is also downloaded when the Reset/Dump button located on the module’s front panel is pressed.

This section discusses the following topics:

- Configuring and enabling dump files
- Displaying dump status information
- Viewing dump file locations
- Modifying dump file locations

Capturing Restart or Crash Messages and Diagnostic Data

Configuring and Enabling Dump Files

You can configure up to eight remote locations to which diagnostic information is downloaded. It is recommended that you configure at least one location. Configuring multiple locations helps ensure that, in the event a download to one location is not successful, the download to another location is likely to succeed.

The VNswitch module dumps files over the Load/Dump/Management (LDM) port only. The LDM port is located on the front panel of the module.

To configure the location of the host or server that is to receive a dump file, perform the following tasks:

Task	Description
1	Configure an IP address for the LDM port and enable TCP/IP host services.
2	Create a dump file on the TFTP server.
3	Specify the name of the dump file, and the location of the server to which the file is to be downloaded.
4	Enable dumping on the switch.
5	Retain multiple dumps at a single location (enable unique-naming). This task is optional.
6	If routing is enabled, use the add address command to configure an IP address for the LDM port.
7	Perform a test dump.

Configuring an IP Address for the LDM Port

To be able to download information, the module must be configured with an in-band IP address and TCP/IP Host Services enabled.

If an IP address is not assigned to the LDM port, the LDM port defaults to the IP-HST address that has been configured for the module. You can specify a different IP address for the LDM port by using the **add address** command at the Boot config prompt.

Note

If routing is enabled, you must use the **add address** command at the Boot Config prompt to configure an IP address for the LDM port.

Capturing Restart or Crash Messages and Diagnostic Data

Refer to the [Enabling and Disabling Host Services](#) section for information on enabling TCP/IP Host Services.

Creating a Dump File on the TFTP Server

On most TFTP servers, you must create a file on the TFTP server with appropriate network access before a dump operation can be completed successfully.

Specifying the File Name and Server Location

To specify the name of the dump file and the location of the server to which the file is to be downloaded, perform the following steps:Enabling and Disabling Dumps

Step	Action
1	At the Boot config prompt, enter add dump-entry .
2	Press Return. The following message is displayed: remote host [0.0.0.0]?
3	Enter the IP address of the remote host to which the dump file is downloaded. The default address is 0.0.0.0.
4	Press Return. The following message is displayed: via gateway (0.0.0.0 if none) [0.0.0.0]?
5	Enter the IP address of the first hop router, if any. The default address is 0.0.0.0. The default should be used if there is no first hop router.
6	Press Return. The following message is displayed: timeout in seconds [10]?
7	Enter the desired TFTP timeout value. TFTP is the protocol the switch uses to download the dump file. If the TFTP request is lost, TFTP retransmits the request several times. The timeout value affects the amount of time that passes before the next retry. The default timeout value of 10 seconds is recommended. A value greater than 10 seconds should be used if the path over which the dump is to occur is typically slow.
8	Press Return. The following message is displayed: file name []?
9	Enter the file name to be assigned to the dump file. Example: file name []?/usr/tftp/switch11.dmp
10	Press Return. The specified values are set and the Boot config prompt is displayed.

Capturing Restart or Crash Messages and Diagnostic Data

Step	Action
11	Repeat steps 1 through 10 for each remote location you want to configure. You can configure a maximum of eight locations.

Enabling and Disabling Dumps

The switch can dump diagnostic data to the remote system location you specify only after you enable dumping. Disabled is the default setting. Dumping remains disabled until you enable it, and remains enabled until you disable it. (Refer to the [Displaying All Boot Config Settings](#) section for information about how to determine whether dumping is enabled on a switch.)

Enabling Dumping

To enable dumping of diagnostic data to the configured remote system location, perform the following steps:

Step	Action
1	At the Boot config prompt, enter enable dumping .
2	Press Return. Dumping is enabled and the Boot config prompt is displayed.

Disabling Dumping

To disable dumping of diagnostic data to the configured remote system location, perform the following steps:

Step	Action
1	At the Boot config prompt, enter disable dumping .
2	Press Return. Dumping is disabled and the Boot config prompt is displayed.

Retaining Multiple Dumps at a Single Location

You have the option of configuring the switch to enable multiple downloads to the same path on a server without overwriting earlier dumps. You do so by configuring the switch to assign a unique name to each dump file that is downloaded. If a unique name is not assigned to each file, then new dumps overwrite the previous file downloaded to the same location.

Capturing Restart or Crash Messages and Diagnostic Data

Note

Because most TFTP servers require that a dump file be created before the dump can occur, the **enable unique-naming** command may not be appropriate or useful.

If retaining multiple dumps (unique naming) is enabled, the switch appends a random suffix of one to five hexadecimal characters to the base file name you specify according to instructions in the [Specifying the File Name and Server Location](#) section. Disabled is the default setting. Retaining multiple dumps remains disabled until you enable it, and remains enabled until you again disable it. (Refer to the [Displaying All Boot Config Settings](#) section for information about how to determine whether unique naming is enabled on a switch.)

Enabling Multiple Dumps

To retain multiple dumps of diagnostic data to the same location on a server, perform the following steps:

Step	Action
1	At the Boot config prompt, enter enable unique-naming .
2	Press Return. Retaining multiple dump files is enabled and the Boot config prompt is displayed.

Disabling Multiple Dumps

To disable the retention of multiple dump files downloaded to the same location, perform the following steps:

Step	Action
1	At the Boot config prompt, enter disable unique-naming .
2	Press Return. Retaining multiple dump files downloaded to the same location is disabled and the Boot config prompt is displayed.

Performing a Test Dump

It is recommended that you perform a test dump of newly configured dump locations. Doing so helps verify that the diagnostic data is properly downloaded in the event of a hardware or software failure, or when the Reset/Dump button is pressed.

Capturing Restart or Crash Messages and Diagnostic Data

Note

For the **dump** command to work, dumping must first be enabled at the `boot config>` prompt.

To test the newly configured dump location, perform the following steps:

Step	Action
1	At the Main prompt (Main>), enter dump .
2	Press Return. The following message is displayed: Dumping will invoke a RESTART. Are you sure you want to dump memory? (Yes or [No])
3	If you want to continue the download and initiate a restart, enter Yes . If you do not want to continue the download, enter No .
4	Press Return. If you entered Yes , the contents of switch memory are downloaded to the remote host and file name you specified. The switch is restarted and the VNswitch module Installation Menu (Figure 2-1) is displayed. If you entered No , the Main prompt (Main>) is displayed.

Capturing Restart or Crash Messages and Diagnostic Data

Performing a Test Dump Using the Reset/Dump button

The Reset/Dump button is located on the module's front panel. When you press this button, the contents of the switch memory is downloaded to the host and filename you specified. The switch is restarted and the VNswitch module Installation Menu is displayed.

Note

If you press the Reset/Dump button without configuring and enabling dump file(s), the VNswitch module is restarted without downloading any diagnostic data.

Displaying Dump Status Information

You can display the status of the most recent dump attempted for a switch. The following information is displayed for each configured dump location:

Information Displayed	Description
Status	Possible outcomes include <code>Successful</code> , <code>Failed</code> , and <code>Not Attempted</code> .
Pathname	The path and file name on the remote server for which the download is configured.
IP Address	The IP address of the remote system to which the dump is downloaded is displayed. If a first hop router is used, the IP address of the first hop router is also listed.

Note

A maximum of eight remote dump locations can be configured for each switch.

To display the status and location of the most recent dump, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter <code>dump information</code> .

Capturing Restart or Crash Messages and Diagnostic Data

Step	Action
2	Press Return. The status and location information about the most recent dump is displayed.

Example

Monitor>dump information

```
1:Dump Failed to "/usr/router1.dmp" on 1.2.3.4 error = TFTP protocol error
2:Dump Not attempted to "/usr/tftp/router1.dmp" on 10.23.2.5
3:Dump Successful to "/usr/tftp/router1.dmp" on 1.2.3.4
```

Viewing Dump File Locations

You can view the locations to which diagnostic data (dump files) are downloaded. To view a list of the file names to be assigned to dump files, and the location of the server on which the files are to reside, perform the following steps:

Step	Action
1	At the Boot config prompt, enter list dump-entries .
2	Press Return. A reference number is displayed for each configured interface, followed by a colon. The reference number is then followed by the location to which each dump file is to be downloaded, including the path and name of the file on the remote system, the IP address of the remote host to which the download occurs, the IP address of the first hop router (if any), and the TFTP retransmission timer value.

Capturing Restart or Crash Messages and Diagnostic Data

Example

```
Boot config>list dump-entries
```

```
Dump to:
```

```
1: "/usr/local/sw1.dmp" on 1.2.3.4 via 0.0.0.0 for 10 secs
```

```
2: "/usr/tftp/sw1.dmp" on 13.12.2.3 via 1.1.2.7 for 10 secs
```

Modifying Dump File Locations

Modifying the remote locations to which diagnostic data is downloaded includes the ability to change and delete the currently configured location.

Changing the File Name and Server Location

To change the name assigned to a dump file or the location of the server to which the download is to occur, perform the following steps:

Step	Action
1	At the Boot config prompt, enter change dump-entry .
2	Press Return. The following message is displayed: Change which entry [1]?
3	Enter the reference number associated with the entry you want to change. The default is 1. The reference number is the first number, followed by a colon, in each item listed when you display information about dump file locations using the list dump-entries command. Refer to the Viewing Dump File Locations section for additional information.
4	Press Return. The following message is displayed: remote host [18.123.0.16]?
5	Enter the new IP address of the remote host to which the dump file is downloaded. The previously configured address is the default.
6	Press Return. The following message is displayed: via gateway (0.0.0.0 if none) [0.0.0.0]?
7	Enter the IP address of the first hop router, if any. The default address is 0.0.0.0. The default should be used if there is no first hop router.
8	Press Return. The following message is displayed: timeout in seconds [3]?

Capturing Restart or Crash Messages and Diagnostic Data

Step	Action
9	Enter the desired TFTP timeout value. TFTP is the protocol the switch uses to download the dump file. If the TFTP request is lost, TFTP retransmits the request several times. The timeout value affects the amount of time that passes before the next retry. The default timeout value of 10 seconds is recommended. A value greater than 10 seconds should be used if the path over which the dump is to occur is typically slow.
10	Press Return. The following message is displayed: File name [user/lib/gw/gwimage.ldb]?
11	Enter the path and file name for the location on the remote server to which the dump file is downloaded.
12	Press Return. The specified values are set and the Boot config prompt is displayed.

Deleting a File Name and Server Location

To delete the name to be assigned to the dump file and the location of the server to which the file is downloaded, perform the following steps:

Step	Action
1	At the Boot config prompt, enter delete dump-entry .
2	Press Return. The following message is displayed: Delete which entry [1]?
3	Enter the reference number associated with the entry you want to delete. The default is 1. The reference number is the first number, followed by a colon, in each item listed when you display information about the configured location for the dump file using the list dump-entries command. Refer to the Viewing Dump File Locations section for additional information.
4	Press Return. The specified location is no longer configured for use during downloads.
5	To verify that the entry is deleted, use the list dump-entries command. Refer to the Viewing Dump File Locations section for additional information.

Displaying All Boot Config Settings

You can view a report that shows all switch settings configured using the Boot config process, including the following information:

- IP address and subnet mask used during reload and dump operations
- Location and file name of boot files on remote systems
- Whether dumping is enabled
- Whether retaining multiple dumps at a single location (unique-naming) is enabled
- Locations and file names on remote systems to which diagnostic data is dumped

Refer to the Configuring Automatic Image Recovery, Capturing Diagnostic Data for Problem Analysis, and Retaining Multiple Dumps at a Single Location sections for detailed information about dumps, BootP software loads, and enabling multiple dumps, respectively.

To display all switch settings that are configured from the Boot config prompt, perform the following steps:

Step	Action
1	At the Boot config prompt, enter list all .
2	Press Return. All switch settings configured from the Boot config prompt are displayed.

Example

```
Boot config>list all
```

```
Interface addresses:
```

```
1: 192.9.1.1 on interface 0, mask 255.255.255.252
2: 192.9.223.39 on interface 5, mask 255.255.255.0
```

```
Boot files:
```

```
1: "/usr/bt/inst.ldb" on 192.9.1.2 through 0.0.0.0 for 3 secs
2: "/usr/bt/in.ldb" on 192.9.2.2 through 192.9.1.4 for 3 secs
```

```
Dumping disabled
```

```
Unique-naming disabled
```

```
Dump to:
```

```
1: "/usr/local/sw1.dmp" on 1.2.3.4 via 0.0.0.0 for 10 secs
2: "/usr/tftp/sw1.dmp" on 13.12.2.3 via 1.1.2.7 for 10 secs
```

Chapter 11

Event Logging and Reporting

Overview

Introduction

The Event Logging System (ELS) is a background process that records operational event messages for a DIGITAL VNswitch module. You can view this log of event messages from your console's CLI, or by making the events available to an SNMP-based agent such as the DIGITAL clearVISN MultiChassis Manager. You can also configure the ELS to record specific types of events and to eliminate others, depending on the level of operational detail you require. For example, you may want to view information that relates only to bridging, or information that relates only to communication between the switch and an IPX server on an Ethernet LAN.

Event messages are recorded by ELS and displayed on your console in abbreviated form. Refer to the *Event Logging System Messages Guide* for expanded descriptions of all event messages, as well as an explanation of the message, possible causes of the event (if applicable), and possible actions you can take to correct error conditions.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Event Messages and Related Concepts	11-2
Selecting Which Events Are Logged	11-8
Displaying the Event Log	11-41
Printing ELS Output	11-44

Event Messages and Related Concepts

You must understand how event messages are generated to be able to interpret the messages. Knowledge of the following concepts is also required if you want to narrow the scope of recorded messages, to focus on specific operations or problem areas.

Types of Events That Are Logged

The ELS records the following general types of events:

- System activity
- Status changes
- Service requests
- Data transmission and reception
- Data and internal errors

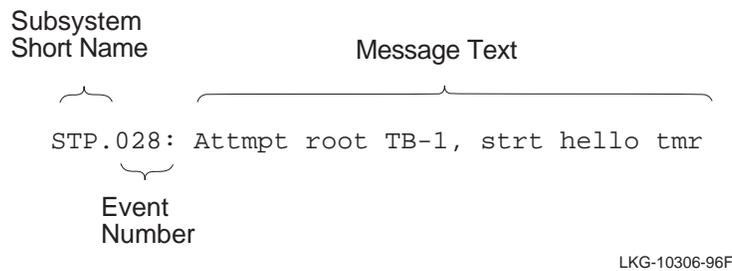
As events occur, ELS receives information from the switch that identifies the source and nature of the events. The information is incorporated into the resultant event message which ELS generates and records. You can use the information to monitor switch activity or to troubleshoot potential problems.

Elements of an Event Message

Event messages are composed of the following three elements:

- Subsystem
- Event number
- Message text

[Figure 11-1](#) is an example of a message generated by an event. It identifies the subsystem, event number, and message text components.

Figure 11-1: Sample Message Generated by an Event

Subsystem

The Event Logging System divides VNswitch module functionality into several operational subsystems. These include bridging, the Spanning Tree Protocol, and the Ethernet handler, for example. When an event occurs, a subsystem short name is added as a prefix to the event message. The short name identifies with which subsystem the event is related. In [Figure 11-1](#), for example, the event short name is STP, indicating the event is related to a Spanning Tree Protocol operation. When you display the log of event messages, the prefix should help you to more accurately monitor switch activity and isolate potential problems. Refer to [Table 11-1](#) for a list of subsystems and their associated short names.

You can configure the ELS to record only those events generated by one or more specific subsystems, or all subsystems. Recording events for one or a few specific subsystems can help you focus on events related to a particular operation or set of operations. Refer to the [Selecting Which Events Are Logged](#) section for information about recording events generated by one or more subsystems.

Note

The subsystems that are active on a given VNswitch module vary, depending on the specific hardware and software configured for the module.

Event Messages and Related Concepts

Table 11-1: Event Subsystems and Associated Short Names

Subsystem Description	Subsystem Short Name
Router base and network library	GW
Address Resolution Protocol	ARP
Internet Protocol	IP
Internet Control Message Protocol	ICMP
Transmission Control Protocol	TCP
User Datagram Protocol	UDP
BooTP relay agent	BTP
Trivial File Transfer Protocol	TFTP
Simple Network Management Protocol	SNMP
Source Routing Transparent Bridge	SRT
Spanning Tree Protocol	STP
Filter Library	FLT
IP Routing Information Protocol	RIP
Exterior Gateway Protocol	EGP
Open SPF-Based Routing Protocol	OSPF
OSPF Multicast extensions	MSPF
DECnet	DN
Xerox Networking Systems Protocol	XNS
Internetwork Packet Exchange Protocol	IPX
AppleTalk	APL
AppleTalk Phase 2	AP2
Apollo Domain Protocol	DDS
IP Protocol Net	IPPN
All subsystems	ALL

Event Messages and Related Concepts

Event Number

The Event Logging System automatically assigns a unique number to each event message generated by a subsystem. In [Figure 11-1](#), for example, the event number is 028. It is separated from the subsystem short name by a period. The short name and event number, together, identify an individual event. You can use the subsystem short name and event number as a parameter for specific ELS configuration and monitoring commands. Only the event indicated by the specified subsystem and event are affected by the ELS command.

Message Text

An abbreviated description of a specific event is provided in the text portion of an event message, as shown in [Figure 11-1](#). For example, the message text `Attmp root TB-1, strt hello tnr`, indicates that the instance of STP on VSD 1 is declaring itself as root, and has restarted the BPDU Hello Timer. The *Event Logging System Messages Guide* provides expanded descriptions of all abbreviated messages, as well as an explanation of the message, possible causes of the event (if applicable), and possible actions you can take to correct error conditions.

Some event messages include fields that display variable values, such as network and interface numbers, source addresses, and error codes. Refer to the *Event Logging System Messages Guide* for a detailed discussion about these variables.

Logging Levels and Event Types

The logging level is a further classification of messages according to the type of event that generated a message. For example, a particular type of event might typically be caused by an unusual internal error. In such a case, ELS associates the UI-ERROR (unusual internal error) logging level with the event. [Table 11-2](#) shows the full list of logging levels and the type of condition that generates events associated with the level.

Event Messages and Related Concepts

Table 11-2: Logging Levels and Associated Event Types

Logging Level	Event Type
UI-ERROR	Unusual internal errors
CI-ERROR	Common internal errors
UE-ERROR	Unusual external errors
CE-ERROR	Common external errors
ERROR	Includes all <i>error</i> levels above
U-INFO	Unusual informational comment
C-INFO	Common informational comment
INFO	Includes all <i>comment</i> levels above
STANDARD	Includes all <i>error</i> levels and all <i>informational comment</i> levels (recommended default)
P-TRACE	Per packet trace
U-TRACE	Unusual operation trace message
C-TRACE	Common operation trace message
TRACE	Includes all <i>trace</i> levels above
ALL	Includes <i>all</i> logging levels

The logging level can be used for informational purposes, and to provide you with the ability to further narrow the scope of recorded events. Assume, for example, the following event message is displayed on your console:

```
SNMP.005 no access; comm "community", hst source_address
```

If you are interested in knowing the type of event that generated the message, you would look up the event using the subsystem short name and event number (SNMP.005) in the *Event Logging System Messages Guide*. The guide lists the logging level as U-TRACE, indicating the message is the result of an “unusual operation packet trace.”

If you want to narrow the scope of events recorded by ELS to only those involving unusual operation packet traces, you would specify U-TRACE as the logging level when configuring ELS. Refer to the [Selecting Which Events Are Logged](#) section for information about how to do so.

Preconfigured Logging Criteria (Groups)

You can preconfigure customized lists of one or more event numbers that you want to record. The customized lists are referred to as groups. You can subsequently configure ELS to record all occurrences of events listed in a group by simply entering the name of the group. Configuring ELS using groups is most useful when you have combinations of events that you regularly need to record. Using groups helps eliminate the need to enter the desired event numbers individually.

Selecting Which Events Are Logged

You can configure ELS to log specific types of events and to eliminate other events, depending on the level of detail you require. For example, you may want to log and view events that relate only to STP. Or, you may want to log events generated by both the STP and Ethernet Handler subsystems, and only events associated with the P-TRACE logging level. When you configure ELS to log events associated with a particular subsystem, event number, logging level, or group, those settings remain in effect until you either clear all configuration settings, or clear selected settings.

Modes of Configuration

You can configure ELS in nonvolatile memory and in volatile memory. ELS configurations that you set in nonvolatile memory (NVRAM) require that you restart the module to take effect. These settings survive power outages and module restarts. This method of setting values is most useful when configuring selection criteria for events that you want to monitor on a regular basis. You configure ELS in NVRAM from the Config prompt.

ELS configurations you set in volatile memory do not require that you restart the module to take effect. These settings do not survive power outages and module restarts unless you save them to a reserved portion of NVRAM (Refer to the [Saving and Managing a Configuration in NVRAM](#) section later in this chapter). Setting ELS configurations in volatile memory is most useful when configuring selection criteria you want to remain in effect only temporarily, and that you are likely to change from moment to moment as you troubleshoot a particular problem. You configure ELS in volatile memory from the Monitor prompt.

Commands Used to Log Events

You can view the log of event messages from your console's CLI, or by sending the events to SNMP. The process of sending specific events over SNMP is referred to as trapping. If sent to SNMP, the events are viewed using an SNMP-based agent such as the DIGITAL clearVISN MultiChassis Manager. The tasks you perform to select which events are logged are the same whether you plan to view the messages via the CLI, or through an SNMP agent. However, the commands you use are different. [Table 11-3](#) lists the commands you use to record and clear events for viewing using the CLI, and the commands used to trap events and clear traps for SNMP.

Selecting Which Events Are Logged

Table 11-3: Commands Used to Log and Trap Event Messages

Command	Description
<u>display</u>	Specifies which events are logged by ELS so you can display them using the CLI.
<u>nodisplay</u>	Clears previously configured events so they are <i>not</i> logged by ELS for display using the CLI.
<u>trap</u>	Specifies which events are trapped and sent to SNMP.
<u>notrap</u>	Clears previously configured events so they are <i>not</i> trapped and sent to SNMP.

Refer to the [Displaying the Event Log](#) section for information about how to display the specified events through the CLI. For information about how to display events sent to SNMP, refer to the appropriate vendor documentation that supports the particular SNMP-based agent you are using.

Special Convention Used in This Section

The following sections present instructions for recording events you plan to view through the CLI and those instructions used to view events via SNMP. The steps you perform are the same for both viewing methods; however, the commands are different. A command used to trap events, or to clear a trap, are enclosed in parentheses and immediately follow the command used to record events to be viewed via the CLI.

Example

```
ELS config>display (trap) subsystem subsystem-shortname
```

You can use only one command (**display** or **trap**) at a time. Do not enter the parentheses when using the SNMP-related commands.

Selecting Which Events Are Logged

Configuring ELS in Nonvolatile Memory

This section describes how to configure ELS in nonvolatile memory (NVRAM). Configurations you set in nonvolatile memory require that you restart the module to take effect. These settings survive power outages and module restarts. This mode is most useful when configuring selection criteria for events you want to monitor on a regular basis.

The following tasks are presented in this section:

- Recording and clearing events by subsystem
- Recording and clearing events by subsystem and logging level
- Recording and clearing all occurrences of an event
- Recording and clearing events by group
- Clearing all previously configured events
- Setting the maximum number of traps per second
- Viewing current configuration settings

Recording and Clearing Events By Subsystem

You can configure ELS to record all events generated by one or more subsystems. For example, you might want to log all events related to the ARP subsystem. Once you configure ELS to record all events generated by a subsystem, that setting remains in effect until you clear it. Assume, for example, you first configure ELS to record events generated by the STP subsystem, and you later decide you want to record events from only the ARP subsystem. In this situation, you must both clear the STP subsystem setting *and* configure ELS to record ARP subsystem events. If you do *not* clear the STP subsystem, events from both the STP and ARP subsystems are recorded.

Recording Events By Subsystem

To record events by subsystem, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter display (trap) subsystem subsystem-shortname , where <i>subsystem-shortname</i> is the short name for the desired subsystem. Refer to Table 11-1 for a list of short names, or enter display (trap) subsystem ? at the ELS config> prompt to display a list of short names on your console.

Selecting Which Events Are Logged

Step	Action
4	<p>Press Return. The ELS is configured to record all messages generated by the specified subsystem using the Standard logging level, and the <code>ELS config></code> prompt is displayed.</p> <p>If ELS was previously configured to record events from other subsystems or using other logging levels, those settings also remain in effect until you clear them.</p>
5	<p>If you want to configure ELS to display events generated by yet another subsystem, repeat steps 3 and 4.</p>
6	<p>Restart the module if you want the new configuration settings to take effect.</p>

Example

```
ELS config>display subsystem arp
```

or

```
ELS config>trap subsystem arp
```

Clearing Events By Subsystem

To clear event recording by subsystem, perform the following steps:

Step	Action
1	<p>At the Config prompt (<code>Config></code>), enter els.</p>
2	<p>Press Return. The <code>ELS config></code> prompt is displayed.</p>
3	<p>Enter nodisplay (notrap) subsystem subsystem-shortname, where <i>subsystem-shortname</i> is the short name for the subsystem you want to clear.</p> <p>Refer to Table 11-1 for a list of short names, or enter nodisplay (notrap) subsystem ? at the <code>ELS config></code> prompt to display a list of short names on your console.</p>
4	<p>Press Return. The ELS is configured not to record messages generated by the specified subsystem, and the <code>ELS config></code> prompt is displayed.</p> <p>If ELS was previously configured to record events from other subsystems, those settings remain in effect until you clear them.</p>
5	<p>If you want to clear event recording for yet another subsystem, repeat steps 3 and 4.</p>

Selecting Which Events Are Logged

Step	Action
6	Restart the module if you want the new configuration settings to take effect.

Example

```
ELS config>nodisplay subsystem arp
```

or

```
ELS config>notrap subsystem arp
```

Recording and Clearing Events By Subsystem and Logging Level

You can configure ELS to record all events generated by a specific subsystem and associated with a particular logging level. For example, you might want to log all events related to STP that are caused by common external errors (the CE-ERROR logging level). Once you configure ELS to record all events generated by a subsystem and an associated logging level, those settings remain in effect until you clear them.

Assume, for example, you first configure ELS to record events generated by the STP subsystem and associated with the P-TRACE logging level, and you later decide you want to record STP events associated with only the U-TRACE logging level. In this situation, you must both clear the P-TRACE logging level from the STP subsystem *and* configure ELS to record STP events associated with the U-TRACE logging level. If you do *not* clear the P-TRACE logging level, STP events associated with both the P-TRACE and U-TRACE logging levels are recorded.

Recording Events By Subsystem and Logging Level

To record events by subsystem and logging level, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter display (trap) subsystem subsystem-shortname logging-level , where <i>subsystem-shortname</i> is the short name for the desired subsystem and <i>logging-level</i> is the desired logging level. Refer to Table 11-1 for a list of short names, or enter display (trap) subsystem ? at the ELS config> prompt to display a list of short names on your console. Refer to Table 11-2 for a list of logging levels, or enter display (trap) subsystem subsystem-shortname ? at the ELS config> prompt to display a list of logging levels.

Selecting Which Events Are Logged

Step	Action
4	<p>Press Return. The ELS is configured to record all messages generated by the specified subsystem and associated with the logging level, and the ELS <code>config></code> prompt is displayed.</p> <p>If ELS was previously configured to record events associated with other logging levels in the same or other subsystems, those settings also remain in effect until you clear them.</p>
5	<p>If you want to configure ELS to record all events associated with other logging levels in the same or other subsystems, repeat steps 3 and 4.</p>
6	<p>Restart the module if you want the new configuration settings to take effect.</p>

Example

```
ELS config>display subsystem stp ce-error
```

or

```
ELS config>trap subsystem stp ce-error
```

Clearing Events By Subsystem and Logging Level

To clear events by subsystem and logging level, perform the following steps:

Step	Action
1	<p>At the Config prompt (<code>Config></code>), enter els.</p>
2	<p>Press Return. The ELS <code>config></code> prompt is displayed.</p>
3	<p>Enter nodisplay (notrap) subsystem subsystem-shortname logging-level, where <i>subsystem-shortname</i> is the short name for the desired subsystem and <i>logging-level</i> is the logging level you want to clear.</p> <p>Refer to Table 11-1 for a list of short names, or enter nodisplay (notrap) subsystem ? at the ELS <code>config></code> prompt to display a list of short names on your console.</p> <p>Refer to Table 11-2 for a list of logging levels, or enter nodisplay (notrap) subsystem subsystem-shortname ? at the ELS <code>config></code> prompt to display a list of logging levels.</p>
4	<p>Press Return. The ELS is configured <i>not</i> to record messages generated by the specified subsystem and logging level, and the ELS <code>config></code> prompt is displayed.</p> <p>If ELS was previously configured to record events associated with other logging levels in the same or other subsystems, those settings remain in effect until you clear them.</p>

Selecting Which Events Are Logged

Step	Action
5	If you want to clear event recording associated with other logging levels in the same or other subsystems, repeat steps 3 and 4.
6	Restart the module if you want the new configuration settings to take effect.

Example

```
ELS config>nodisplay subsystem stp ce-error
```

or

```
ELS config>notrap subsystem stp ce-error
```

Recording and Clearing All Occurrences of an Event

You can configure ELS to record all occurrences of a particular event. For example, you might want to record each time STP restarts the BPDU Hello Timer. The event number for this event is STP.028.

Once you configure ELS to record all occurrences of an event, that setting remains in effect until you clear it. Assume, for example, you first configure ELS to record each time STP restarts the BPDU Hello Timer, and you later decide you want to record only when a BootP request is received on an interface (event number BTP.001). In this situation, you must both clear the STP.028 event number *and* configure ELS to record each time a BootP request is received on an interface. If you do *not* clear the STP.028 event number, all occurrences of both events (STP.028 and BTP.001) are recorded.

Recording All Occurrences of an Event

To record all occurrences of an event, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter display (trap) event subsystem.event# , where <i>subsystem.event#</i> is the event number for the event you want to record. Refer to the Viewing Current Configuration Settings section for information about how to display event numbers associated with one or all subsystems.

Selecting Which Events Are Logged

Step	Action
4	Press Return. The ELS is configured to record all occurrences of the specified event, and the <code>ELS config></code> prompt is displayed. If ELS was previously configured to record all occurrences of other individual events, those settings also remain in effect until you clear them.
5	If you want to configure ELS to record all occurrences of another individual event, repeat steps 3 and 4.
6	Restart the module if you want the new configuration settings to take effect.

Example

```
ELS config>display event stp.028
```

or

```
ELS config>trap event stp.028
```

Clearing All Occurrences of an Event

To clear all occurrences of an event, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter els .
2	Press Return. The <code>ELS config></code> prompt is displayed.
3	Enter nodisplay (notrap) event subsystem.event# , where <i>subsystem.event#</i> is the event number for the event you want to clear. Refer to the Viewing Current Configuration Settings section for information about how to display event numbers associated with one or all subsystems.
4	Press Return. The ELS is configured <i>not</i> to record all occurrences of the specified event, and the <code>ELS config></code> prompt is displayed. If ELS was previously configured to record all occurrences of other individual events, those settings remain in effect until you clear them.
5	If you want to clear recording of all occurrences of other individual events, repeat steps 3 and 4.
6	Restart the module if you want the new configuration settings to take effect.

Selecting Which Events Are Logged

Example

```
ELS config>nodisplay event stp.028
```

or

```
ELS config>notrap event stp.028
```

Recording and Clearing Events By Group

You can configure ELS to record all occurrences of each event that is included in a group. For example, you might want to regularly record each time STP restarts the BPDU Hello Timer (event number STP.028), and each time a BPDU is received (event number STP.001). A group is a customized list of one or more event numbers that you define as members of a group. Configuring ELS using groups is most useful when you have combinations of events that you need to record regularly and helps eliminate the need to enter the desired event numbers individually. The name you assign to a group when you create it should reflect the type of events the group contains so that you can more easily distinguish between groups.

Once you configure ELS to record all occurrences of events that are members of a group, that setting remains in effect until you clear it.

Creating a New Group and Adding Events to an Existing Group

You must create a group before you can configure ELS to record all occurrences of events in a group.

To create a group or add events to an existing group, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter add group-name subsystem.event-number , where group-name is the name of a new group you want to create or is the name of an existing group to which you want to add an event, and subsystem.event-number is the name of the event you want to add to the group. The group name must begin with an alphabetic character. All subsequent characters can be alphabetic or numeric. ELS is case sensitive with respect to group names. The name you assign to a group should reflect the type of events the group contains, so that you can more easily distinguish between groups. For example, a group that is used to list particular types of Spanning Tree Protocol events might be named stp2.
4	Press Return. The group is created or modified with the specified event entries, and the ELS config> prompt is displayed.

Selecting Which Events Are Logged

Step	Action
5	Repeat steps 3 and 4 for each event you want to add to a group.
6	Restart the module if you want the new configuration settings to take effect.

Example

```
ELS config>add stp2 stp.001
```

and

```
ELS config>add stp2 stp.028
```

Deleting an Event From a Group and Deleting an Entire Group

You can delete either a single event from an existing group or you can delete a group and all of its contents.

To delete an event from a group or to delete an entire group, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	<p>If you want to delete an event from a group, enter delete group-name subsystem.event-number, where <i>group-name</i> is the name of the group containing the event you want to delete, and <i>subsystem.event-number</i> is the name of the event you want to delete. ELS is case sensitive with respect to group names.</p> <p>If you want to delete an entire group, including its contents, enter delete group-name all, where <i>group-name</i> is the name of the group you want to delete.</p> <p>Refer to the Viewing Current Configuration Settings section for information about how to display a list of current groups and the events that compose each group.</p>
4	<p>Press Return. The event or group is deleted as specified.</p> <p>If you are deleting the last event in a group, a message is displayed to notify you of the fact.</p> <p>If you are not deleting the last event in a group, the ELS config> prompt is displayed.</p>

Selecting Which Events Are Logged

Step	Action
5	If you are deleting events from a group and you want to delete another event, repeat steps 3 and 4 for each event you want to delete from the group.
6	Restart the module if you want the deletion to take effect.

Example

```
ELS config>delete stp2 stp.001
```

and

```
ELS config>delete stp2 all
```

Recording All Occurrences of Events in a Group

To record all occurrences of events in a group, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter display (trap) group group-name , where <i>group-name</i> is the name of the group containing the events you want to record. Refer to the Viewing Current Configuration Settings section for information about how to display a list of current groups and the events that compose each group.
4	Press Return. The ELS is configured to record all occurrences of the events specified in the group, and the ELS config> prompt is displayed. If ELS was previously configured to record all occurrences of events in other groups, those settings also remain in effect until you clear them.
5	If you want to configure ELS to record all occurrences of events in another group, repeat steps 3 and 4.
6	Restart the module if you want the new configuration settings to take effect.

Selecting Which Events Are Logged

Example

```
ELS config>display group mygroup
```

or

```
ELS config>trap group mygroup
```

Clearing All Occurrences of Events in a Group

To clear all occurrences of events in a group, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter nodisplay (notrap) group group-name , where <i>group-name</i> is the name of the group for which you do <i>not</i> want to record events.
4	Press Return. ELS is configured <i>not</i> to record all occurrences of events in the specified group, and the ELS config> prompt is displayed. If ELS was previously configured to record all occurrences of events in other groups, those settings remain in effect until you clear them.
5	If you want to clear recording of events in another group, repeat steps 3 and 4.
6	Restart the module if you want the new configuration settings to take effect.

Example

```
ELS config>nodisplay group mygroup
```

or

```
ELS config>notrap group mygroup
```

Clearing All Previously Configured Events

You can clear ELS so that all events that it was previously configured to log are no longer recorded. This includes all events configured to be recorded by subsystem, by subsystem and logging level, by event number, and by group, and for both CLI and SNMP display. When you clear all previously configured events, no events are logged to ELS. Clearing ELS also resets the maximum number of traps per second to its default.

Selecting Which Events Are Logged

To clear ELS so that all events that it was previously configured to log are no longer recorded, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter clear .
4	Press Return. The following message is displayed: You are about to clear all ELS configuration information. Are you sure you want to do this (Yes or [No]):
5	If you want to clear all ELS configuration information, enter y . If you do not want to clear all ELS configuration information, enter n .
6	Press Return. The ELS config> prompt is displayed.
7	Restart the module if you want the new configuration setting to take effect. All events for which ELS was previously configured to log are no longer recorded.

Setting Maximum Number of Traps Per Second

You can configure ELS to limit the number of events that are trapped per second. This option is most useful when conditions result in such large numbers, or “bursts,” of events that you are overwhelmed by the data. Reducing the maximum number of traps per second effectively results in a sampling of the events.

The default value for the maximum number of traps per second is 0 (zero), meaning an unlimited number of traps per second is permitted. The maximum number of traps per second can be reset to its default setting by using the **clear** command. (Refer to the [Clearing All Previously Configured Events](#) section.)

To set the maximum number of traps per second, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter set pin .

Selecting Which Events Are Logged

Step	Action
4	Press Return. The following message is displayed: <code>events/second [0]?</code>
5	Enter the maximum number of events per second you want trapped. The range of acceptable values is 1 through 57600 traps per second. The default is 0 (zero), an unlimited number of events per second.
6	Press Return. The maximum number of events trapped is set, and the <code>ELS config></code> prompt is displayed.
7	Restart the module if you want the new configuration setting to take effect.

Viewing Current Configuration Settings

You can view several reports that detail the current ELS configuration settings in NVRAM. You may find these reports helpful if, for example, you want to check the settings before restarting the module and before displaying the event log. (Refer to the [Displaying the Event Log](#) section for information about how to display events.)

To view current configuration settings in NVRAM, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter <code>els</code> .
2	Press Return. The <code>ELS config></code> prompt is displayed.
3	Enter <code>list command-option</code> , where <i>command-option</i> is the command you must enter to display the desired information. Refer to Table 11-4 for a list of commands, and for a description of the type of information that is displayed when you enter the command.
4	Press Return. A report containing the desired information is displayed and the <code>ELS config></code> prompt is displayed.

Selecting Which Events Are Logged

Table 11-4: List Command Options and Descriptions

Command Option	Description
<u>groups</u>	Lists all group names and the events that compose each group.
<u>pin</u>	Displays the current maximum number of traps per second (pin).
<u>status</u>	Lists all event logging configurations by subsystem, subsystem and logging level, group, and event number. The information includes configurations for events that will be displayed through the CLI and those to be trapped for use by SNMP.
<u>subsystem</u>	Lists the short names for all possible subsystems, the number of different events that can be generated by the subsystem, and an expanded description of the subsystem short name.
<u>subsystem</u> <i>subsystem</i>	Lists all possible events that can be generated by a specified subsystem, where <i>subsystem</i> is the short name for the subsystem for which you want to list events. The information included for each event includes the event number, the logging level for the event, and a short description of the event.
<u>subsystems</u> <u>all</u>	Lists all possible events that can be generated by all subsystems. The information included for each event includes the event number, the logging level for the event, and a short description of the event.

Selecting Which Events Are Logged

Command Option	Description
all	<p>Lists the following configuration information for events to be displayed through the CLI, and those to be trapped for use by SNMP:</p> <ul style="list-style-type: none">• Short names for all possible subsystems, the number of different events that can be generated by each subsystem, and an expanded description of each subsystem short name• All event logging configurations by subsystem, subsystem and logging level, group, and event number• The (pin) value for the maximum number of traps per second

Configuring ELS in Volatile Memory

This section describes how to configure ELS in volatile RAM. Configurations you set in volatile memory *do not* require that you restart the module to take effect. These settings *do not* survive power outages and module restarts. Configuring ELS in volatile memory is most useful when configuring selection criteria you want to remain in effect only temporarily, and that you are likely to change from moment to moment as you troubleshoot a particular problem.

The following tasks are presented in this section:

- Recording and clearing events by subsystem
- Recording and clearing events by subsystem and logging level
- Recording and clearing all occurrences of an event
- Recording and clearing events by group
- Clearing historical entries from the log
- Setting the maximum number of traps per second
- Saving and managing a configuration in NVRAM
- Restoring default settings
- Viewing current configuration settings

Selecting Which Events Are Logged

Recording and Clearing Events By Subsystem

You can configure ELS to record all events generated by one or more subsystems. For example, you might want to log all events related to the ARP subsystem. Once you configure ELS to record all events generated by a subsystem, that setting remains in effect until you clear it. Assume, for example, you first configure ELS to record events generated by the STP subsystem, and you later decide you want to record events from only the ARP subsystem. In this situation, you must both clear the STP subsystem setting *and* configure ELS to record ARP subsystem events. If you do *not* clear the STP subsystem, events from both the STP and ARP subsystems are recorded.

Recording Events By Subsystem

To record events by subsystem, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter <code>els</code> .
2	Press Return. The <code>ELS></code> prompt is displayed.
3	Enter <code>display (trap) subsystem subsystem-shortname</code> , where <i>subsystem-shortname</i> is the short name for the desired subsystem. Refer to Table 11-1 for a list of short names, or enter <code>display (trap) subsystem ?</code> at the <code>ELS></code> prompt to display a list of short names on your console.
4	Press Return. The ELS is configured to record all messages generated by the specified subsystem, and the <code>ELS></code> prompt is displayed. If ELS was previously configured in volatile memory to record events from other subsystems, or if ELS was configured in NVRAM, those settings also remain in effect until you clear them.
5	If you want to configure ELS to display events generated by yet another subsystem, repeat steps 3 and 4.

Example

```
ELS>display subsystem arp
```

or

```
ELS>trap subsystem arp
```

Selecting Which Events Are Logged

Clearing Events By Subsystem

Clearing events in volatile memory causes ELS to stop recording any *new* events generated by the specified subsystem. Events already recorded (historical entries) are still displayed and may be visible at the upper portion of the log. Refer to the [Clearing Historical Entries From the Log](#) section for information about how to clear historical entries.

To clear event recording by subsystem, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter els .
2	Press Return. The <code>ELS></code> prompt is displayed.
3	Enter nodisplay (notrap) subsystem subsystem-shortname , where <i>subsystem-shortname</i> is the short name for the subsystem you want to clear. Refer to Table 11-1 for a list of short names, or enter nodisplay (notrap) subsystem ? at the <code>ELS></code> prompt to display a list of short names on your console.
4	Press Return. The ELS is configured not to record messages generated by the specified subsystem, and the <code>ELS></code> prompt is displayed. If ELS was previously configured in volatile memory to record events from other subsystems, or if other subsystems were configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear event recording for yet another subsystem, repeat steps 3 and 4.

Example

```
ELS>nodisplay subsystem arp
```

or

```
ELS>notrap subsystem arp
```

Recording and Clearing Events By Subsystem and Logging Level

You can configure ELS to record all events generated by a specific subsystem and associated with a particular logging level. For example, you might want to log all events related to STP that are caused by common external errors (the CE-ERROR logging level). Once you configure ELS to record all events generated by a subsystem and an associated logging level, those settings remain in effect until you clear them. Assume, for example, you first configure ELS to record events generated by the STP subsystem and associated with the P-TRACE logging level, and you later decide you

Selecting Which Events Are Logged

want to record STP events associated with only the U-TRACE logging level. In this situation, you must both clear the P-TRACE logging level from the STP subsystem *and* configure ELS to record STP events associated with the U-TRACE logging level. If you do *not* clear the P-TRACE logging level, STP events associated with both the P-TRACE and U-TRACE logging levels are recorded.

Recording Events By Subsystem and Logging Level

To record events by subsystem and logging level, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter display (trap) subsystem subsystem-shortname logging-level , where <i>subsystem-shortname</i> is the short name for the desired subsystem and <i>logging-level</i> is the desired logging level. Refer to Table 11-1 for a list of short names, or enter display (trap) subsystem ? at the ELS> prompt to display a list of short names on your console. Refer to Table 11-2 for a list of logging levels, or enter display (trap) subsystem subsystem-shortname ? at the ELS> prompt to display a list of logging levels.
4	Press Return. The ELS is configured to record all messages generated by the specified subsystem and associated with the specified logging level, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events associated with other logging levels in the same or other subsystems, or if other subsystems or logging levels were configured in NVRAM, those settings also remain in effect until you clear them.
5	If you want to configure ELS to record all events associated with other logging levels in the same or other subsystems, repeat steps 3 and 4.

Example

```
ELS>display subsystem stp ce-error
```

or

```
ELS>trap subsystem stp ce-error
```

Selecting Which Events Are Logged

Clearing Events By Subsystem and Logging Level

Clearing events in volatile memory causes ELS to stop recording any *new* events generated by the specified subsystem, and associated with the specified logging level. Events already recorded (historical entries) are still displayed and may be visible at the upper portion of the log. Refer to the [Clearing Historical Entries From the Log](#) section for information about how to clear historical entries.

To clear events by subsystem and logging level, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter nodisplay (notrap) subsystem subsystem-shortname logging-level , where <i>subsystem-shortname</i> is the short name for the desired subsystem and <i>logging-level</i> is the logging level you want to clear. Refer to Table 11-1 for a list of short names, or enter nodisplay (notrap) subsystem ? at the ELS> prompt to display a list of short names on your console. Refer to Table 11-2 for a list of logging levels, or enter nodisplay (notrap) subsystem subsystem-shortname ? at the ELS> prompt to display a list of logging levels.
4	Press Return. The ELS is configured <i>not</i> to record messages generated by the specified subsystem and logging level, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events associated with other logging levels in the same or other subsystems, or if other subsystems or logging levels were configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear event recording associated with other logging levels in the same or other subsystems, repeat steps 3 and 4.

Example

```
ELS>nodisplay subsystem stp ce-error
```

or

```
ELS>notrap subsystem stp ce-error
```

Selecting Which Events Are Logged

Recording and Clearing All Occurrences of an Event

You can configure ELS to record all occurrences of a particular event. For example, you might want to record each time STP restarts the BPDU Hello Timer. The event number for this event is STP.028.

Once you configure ELS to record all occurrences of an event, that setting remains in effect until you clear it. Assume, for example, you first configure ELS to record each time STP restarts the BPDU Hello Timer, and you later decide you want to record only when a BootP request is received on an interface (event number BTP.001). In this situation, you must both clear the STP.028 event number *and* configure ELS to record each time a BootP request is received on an interface. If you do *not* clear the STP.028 event number, all occurrences of both events (STP.028 and BTP.001) are recorded.

Recording All Occurrences of an Event

To record all occurrences of an event, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS config> prompt is displayed.
3	Enter display (trap) event subsystem.event# , where <i>subsystem.event#</i> is the event number for the event you want to record. Refer to the Viewing Current Configuration Settings and Statistics section for information about how to display event numbers associated with one or all subsystems.
4	Press Return. The ELS is configured to record all occurrences of the specified event, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events based on other criteria, or if other event selection criteria were configured in NVRAM, those settings also remain in effect until you clear them.
5	If you want to configure ELS to record all occurrences of another individual event, repeat steps 3 and 4.

Example

```
ELS>display event stp.028
```

or

```
ELS>trap event stp.028
```

Selecting Which Events Are Logged

Clearing All Occurrences of an Event

Clearing events in volatile memory causes ELS to stop recording any *new* events that match the specified event number. Events already recorded (historical entries) are still displayed and may be visible at the upper portion of the log. Refer to the [Clearing Historical Entries From the Log](#) section for information about how to clear historical entries.

To clear all occurrences of an event, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter nodisplay (notrap) event subsystem.event# , where <i>subsystem.event#</i> is the event number for the event you want to clear. Refer to the Viewing Current Configuration Settings and Statistics section for information about how to display event numbers associated with one or all subsystems.
4	Press Return. The ELS is configured <i>not</i> to record all occurrences of the specified event, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events based on other criteria, or if other event selection criteria were configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear recording of all occurrences of other individual events, repeat steps 3 and 4.

Example

```
ELS>nodisplay event stp.028
```

or

```
ELS>notrap event stp.028
```

Recording and Clearing Events By Group

You can configure ELS to record all occurrences of each event that is included in a group. For example, you might want to regularly record each time STP restarts the BPDU Hello Timer (event number STP.028), and each time a BPDU is received from a specified MAC address (event number STP.001). A group is a customized list of one or more event numbers that you define as members of a group. Configuring ELS using groups is most useful when you have combinations of events that you need to record regularly and helps eliminate the need to enter the desired event numbers individually.

Selecting Which Events Are Logged

The name you assign to a group when you create it should reflect the type of events the group contains so that you can more easily distinguish between groups. Refer to the [Configuring ELS in Nonvolatile Memory](#) section for information about how to create, modify, and delete groups.

Once you configure ELS to record all occurrences of events that are members of a group, that setting remains in effect until you clear it.

Recording All Occurrences of Events in a Group

To record all occurrences of events in a group, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS <code>config></code> prompt is displayed.
3	Enter display (trap) group <i>group-name</i> , where <i>group-name</i> is the name of the group containing the events you want to record.
4	Press Return. The ELS is configured to record all occurrences of the events specified in the group, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events based on other criteria, or if other event selection criteria were configured in NVRAM, those settings also remain in effect until you clear them.
5	If you want to configure ELS to record all occurrences of events in another group, repeat steps 3 and 4.

Example

```
ELS>display group mygroup
```

or

```
ELS>trap group mygroup
```

Clearing All Occurrences of Events in a Group

Clearing events in volatile memory causes ELS to stop recording any *new* events that match the specified group. Events already recorded (historical entries) are still displayed and may be visible at the upper portion of the log. Refer to the [Clearing Historical Entries From the Log](#) section for information about how to clear historical entries.

Selecting Which Events Are Logged

To clear all occurrences of events in a group, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter els .
2	Press Return. The <code>ELS></code> prompt is displayed.
3	Enter nodisplay (notrap) group <i>group-name</i> , where <i>group-name</i> is the name of the group for which you do <i>not</i> want to record events.
4	Press Return. ELS is configured <i>not</i> to record all occurrences of events in the specified group, and the <code>ELS></code> prompt is displayed. If ELS was previously configured in volatile memory to record events based on other criteria, or if other criteria are configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear recording of events in another group, repeat steps 3 and 4.

Example

```
ELS>nodisplay group mygroup
```

or

```
ELS>notrap group mygroup
```

Clearing Historical Entries From the Log

You can clear ELS so that all *historical* entries in the log, matching the specified criteria, are deleted. This includes all events configured to be recorded by subsystem, by subsystem and logging level, by event number, and by group, and for both CLI and SNMP display. When you clear historical entries, new events continue to be recorded to the log, even if they match the criteria used to clear entries.

Clearing Historical Entries By Subsystem

To clear ELS so that all historical entries are deleted from the log by subsystem, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter els .
2	Press Return. The <code>ELS></code> prompt is displayed.

Selecting Which Events Are Logged

Step	Action
3	Enter clear subsystem <i>subsystem-shortname</i> , where <i>subsystem-shortname</i> is the short name for the subsystem you want to clear. Refer to Table 11-1 for a list of short names, or enter list subsystem ? at the ELS> prompt to display a list of short names on your console.
4	Press Return. The ELS deletes all historical entries in the log that match the specified subsystem, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events from other subsystems, or if other subsystems were configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear historical entries for yet another subsystem, repeat steps 3 and 4.

Clearing Historical Entries By Subsystem and Logging Level

To clear historical entries by subsystem and logging level, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter clear subsystem <i>subsystem-shortname</i> logging-level , where <i>subsystem-shortname</i> is the short name for the desired subsystem and <i>logging-level</i> is the logging level you want to clear. Refer to Table 11-1 for a list of short names, or enter list subsystem ? at the ELS> prompt to display a list of short names on your console. Refer to Table 11-2 for a list of logging levels, or enter list subsystem <i>subsystem-shortname</i> ? at the ELS> prompt to display a list of events and associated logging levels.
4	Press Return. The ELS deletes historical entries from the log that match the specified subsystem and logging level, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events associated with other logging levels in the same or other subsystems, or if other subsystems or logging levels were configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear historical entries associated with other logging levels in the same or other subsystems, repeat steps 3 and 4.

Selecting Which Events Are Logged

Clearing All Historical Entries of a Specific Event

To clear all historical entries of a specific event, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter clear event subsystem.event# , where <i>subsystem.event#</i> is the event number for the event you want to clear. Refer to the Viewing Current Configuration Settings and Statistics section for information about how to display event numbers associated with one or all subsystems.
4	Press Return. The ELS deletes all historical entries of the specified event from the log, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events based on other criteria, or if other event selection criteria were configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear recording of all occurrences of other individual events, repeat steps 3 and 4.

Clearing Historical Entries By Group

To clear all historical entries associated with events in a group, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter clear group group-name , where <i>group-name</i> is the name of the group for which you do <i>not</i> want to record events.
4	Press Return. ELS deletes all historical entries associated with events in the specified group, and the ELS> prompt is displayed. If ELS was previously configured in volatile memory to record events based on other criteria, or if other criteria are configured in NVRAM, those settings remain in effect until you clear them.
5	If you want to clear recording of events in another group, repeat steps 3 and 4.

Selecting Which Events Are Logged

Setting Maximum Number of Traps Per Second

You can configure ELS to limit the number of events that are trapped per second. This option is most useful when conditions result in such large numbers, or “bursts,” of events that you are overwhelmed by the data. Reducing the maximum number of traps per second effectively results in a sampling of the events.

The default value for the maximum number of traps per second in volatile memory is equal to the maximum number of traps/second value set for NVRAM. (Refer to the [Configuring ELS in Nonvolatile Memory](#) section.)

To set the maximum number of traps per second, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter <code>els</code> .
2	Press Return. The <code>ELS></code> prompt is displayed.
3	Enter <code>set pin</code> .
4	Press Return. The following message is displayed: <code>events/second [0]?</code>
5	Enter the maximum number of events per second you want trapped. The range of acceptable values is 1 through 57600 traps per second. A value of 0 (zero), indicates an unlimited number of events per second.
6	Press Return. The maximum number of events trapped is set, and the <code>ELS></code> prompt is displayed.

Saving and Managing a Configuration in NVRAM

You can save an ELS volatile memory configuration in NVRAM. The configuration settings are saved in NVRAM, separate from the ELS NVRAM settings configured from the Config prompt. You may want to save volatile settings in NVRAM if, for example, you want to try a different ELS configuration, but want to retain the current settings for later use. You may also want to save volatile settings in NVRAM if you plan to continue to use the settings at a later time, and you want to ensure the settings are still available if a power outage or restart occurs.

You can retrieve and reload volatile ELS configuration settings you previously saved in NVRAM. Retrieving settings from NVRAM does not delete the settings from NVRAM. You must perform a separate operation to delete (remove) the saved NVRAM.

Selecting Which Events Are Logged

Saving Volatile Settings in NVRAM

To save a volatile ELS configuration in NVRAM, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter save .
4	Press Return. The volatile settings are stored in NVRAM, separate from ELS settings configured from the Config prompt. The ELS> prompt is displayed.

Retrieving Settings From NVRAM

To retrieve previously saved (volatile settings) from NVRAM, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter retrieve .
4	Press Return. The settings stored in NVRAM are retrieved and reloaded as an ELS volatile configuration, and the ELS> prompt is displayed. The NVRAM copy of the settings from which the reload occurred is retained. (Refer to the Deleting Settings Saved in NVRAM section for information about how to delete (remove) the NVRAM copy.)

Deleting Settings Saved in NVRAM

To delete (remove) previously saved (volatile) settings from NVRAM, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter remove .
4	Press Return. The settings stored in NVRAM are deleted and the ELS> prompt is displayed.

Selecting Which Events Are Logged

Restoring Default Settings

You can restore ELS volatile memory configuration defaults using the restore command. (Refer to [Appendix B](#) for a list of ELS volatile memory configuration default settings.) Restoring the default settings also stops ELS from recording events based on prior settings in volatile memory. However, ELS continues to record events that match selection criteria configured in NVRAM, if any.

To restore ELS volatile memory configuration defaults, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter restore .
4	Press Return. The default settings are restored and the ELS> prompt is displayed.

Viewing Current Configuration Settings and Statistics

Viewing Current Configuration Settings

You can view several reports that detail the current ELS configuration settings in volatile memory. Some reports include a count of specific events being recorded. You may find these reports helpful if, for example, you want to check the settings before displaying the event log. (Refer to the [Displaying the Event Log](#) section for information about how to display events.)

To view current configuration settings in volatile memory, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter list command-option , where <i>command-option</i> is the command you must enter to display the desired information. Refer to Table 11-5 for a list of commands, and for a description of the type of information that is displayed when you enter the command.
4	Press Return. A report containing the desired information is displayed and the ELS> prompt is displayed.

Selecting Which Events Are Logged

Table 11-5: List Command Options and Descriptions

Command Option	Description
<u>active</u> <i>subsystem-name</i>	<p>Lists the following information about a specific subsystem, where <i>subsystem-name</i> is the subsystem:</p> <ul style="list-style-type: none"> • Events configured for logging (active events). A <i>D</i> in the <i>Active</i> column of the report indicates the event is configured to be logged for display through the VNswitch module's CLI. A <i>T</i> in the <i>Active</i> column of the report indicates the event is configured to be trapped for remote (SNMP) management. • Number of occurrences of each event.
<u>event</u> <i>event-name</i>	<p>Displays the following information about the specified event, where <i>event-name</i> is the name of the event:</p> <ul style="list-style-type: none"> • Logging level. • Short form of the message. • Whether the event is configured for logging (active event). A <i>D</i> in the <i>Active</i> column of the report indicates the event is configured to be logged for display through the VNswitch module's CLI. A <i>T</i> in the <i>Active</i> column of the report indicates the event is configured to be trapped for remote (SNMP) management.
<u>groups</u>	<p>Lists all group names and the events that compose each group.</p>
<u>pin</u>	<p>Displays the current maximum number of traps per second (pin).</p>
<u>subsystems</u>	<p>Lists the short names for all possible subsystems, the number of different events that can be generated by the subsystem, and an expanded description of the subsystem short name.</p>

Selecting Which Events Are Logged

Command Option	Description
<u>subsystem</u> <i>subsystem</i>	Lists all possible events that can be generated by a specified subsystem, where <i>subsystem</i> is the short name for the subsystem for which you want to list events. The information included for each event includes the event number, the logging level for the event, and a short description of the event.
<u>subsystems</u> all	Lists all possible events that can be generated by all subsystems. The information included for each event includes the event number, the logging level for the event, and a short description of the event.
<u>all</u>	Lists the following configuration information: <ul style="list-style-type: none">• Short names for all possible subsystems, the number of different events that can be generated by each subsystem, and an expanded description of each subsystem short name.• All configured groups and the events specified for each group.• The (pin) value for the maximum number of traps per second.

Selecting Which Events Are Logged

Example

ELS>list active stp

Event	Active	Count
STP.003	D	0
STP.004	D	0
STP.005	D	0
STP.006	D	0
STP.007	D	0
STP.008	D	0
STP.009	D	0
STP.013	D	0
STP.014	D	0
STP.015	D	0
STP.016	D	0
STP.021	D	0
STP.022		4397
STP.024	D	0
STP.026	D	0
STP.028	D	4397
STP.029	D	0
STP.030	D	0
STP.032	D	0
STP.033	D	0
STP.034	D	0

Viewing Configuration Memory Statistics

You can view statistics about ELS configuration volatile memory categorized by subsystem. The report includes the following information for each subsystem:

- Subsystem short name
- Maximum number of events (vector) for the subsystem
- Current number of events that can be generated by the subsystem
- Number of bytes (string) used for all message storage by the subsystem
- Number of events that are active in the subsystem
- Dynamic memory in use by the subsystem

The report also includes totals for each of the above bulleted items, the maximum number of events (vector) of all subsystems, the maximum number of subsystems, and total heap memory size.

Selecting Which Events Are Logged

To view configuration memory statistics, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter els .
2	Press Return. The ELS> prompt is displayed.
3	Enter statistics .
4	Press Return. A report containing configuration memory is displayed, and the ELS> prompt is displayed.

Displaying the Event Log

The event log records all event messages until the buffer is full. Each new event then overwrites the oldest existing event in the log. The event log also displays new messages on your console as they occur. The categories of information and level of detail displayed varies, depending on how you configure ELS. (Refer to the [Selecting Which Events Are Logged](#) section for information about how to record and display selected information.)

You can display the log of recorded event messages using the CLI, the VNswitch web-based management application, or an SNMP-based agent such as the Digital clearVISN MultiChassis Manager. This section describes how to display the log using the CLI. Refer to the appropriate documentation supporting the particular SNMP agent you are using for information about displaying events.

Choosing the Method of Display

You can display the event log from the CLI either directly or indirectly. When displaying the event log directly, you view events only. You cannot perform any other task while viewing the log using this method.

Displaying the event log indirectly enables you to view events while entering commands from any CLI prompt (`Main>`, `Config>`, `Monitor>`, `Bridge>`, and so on). Although this method causes the command line to scroll off the screen as events are displayed, it enables you to more easily view the effect of configuration changes as you make them. This capability may also be of particular use if you need to view events that occur immediately after startup, but that might otherwise scroll off the screen if you first had to access the event log from the Main prompt (`Main>`), as described in the [Displaying and Exiting the Event Log](#) section. Displaying the event log indirectly applies only to a terminal or a printer attached to the setup port. It does not affect output accessed via remote devices.

You can set the method of display (direct or indirect) from either the `Config>` prompt or the `Monitor>` prompt. Setting the method of display from the `Config>` prompt takes effect immediately, without restarting the switch. Because it is stored in NVRAM, the setting survives resets and loss of power. Setting the method of display from the `Monitor>` prompt also takes effect immediately, without restarting the switch. However, because the setting is stored in volatile memory it does not survive resets and loss of power.

Displaying the Event Log

To choose the method of display, perform the following steps:

Step	Action
1	At either the Config> prompt or the Monitor prompt (Monitor>), enter set output method-of-display , where <i>method-of-display</i> is either the default or the console command option. The default command sets the method of display to direct (you can only view events). The console command sets the method of display to indirect (you can view events while entering commands). The default is the default command.
2	Press Return. The chosen method of display is set. If you set the method of display to default , refer to Displaying and Exiting the Event Log section for information about how to view the event log directly. If you set the method of display to console , the event log is automatically displayed at any CLI prompt you access.

Displaying and Exiting the Event Log

Only one user at a time can access the event log. If another user attempts to access the log while you are using it, the message `Current Process has been Redirected` is displayed and you are returned to the Main prompt (Main>). The user who accessed the log receives all redirected output from the log that did not yet display on your screen. The log is displayed at the other user's terminal until the user to whom the output is redirected cancels it.

Displaying the Event Log Directly

To display the event log directly (the method of display is set to **default**), perform the following steps:

Step	Action
1	At the Main prompt (Main>) enter Events .
2	Press Return. The event log is displayed. All existing event messages are displayed. New event messages are also added to the display as they occur.

Displaying the Event Log Indirectly

If you set the method of display to **console**, the event log is automatically displayed at any CLI prompt you access. If you log out after setting the method of display to **console** and then log in again, the stream of events that are output to the console are immediately displayed. To terminate the stream of output, at the Main prompt enter **halt 2** and press Return.

Exiting the Event Log

If you are viewing the event log directly, enter the intercept character to exit the event log. The default intercept character is Ctrl/P. Refer to [Chapter 2](#) for information about changing the intercept character.

If you are displaying the event log indirectly, you can cancel display of events by entering **set output default** command at either the Config or Monitor prompt. Refer to the [Choosing the Method of Display](#) section for information about how to do so.

Advanced Methods for Viewing Events

Refer to [Appendix A](#) for information about the following topics that may be useful when viewing the event log directly.

- Viewing output from multiple processes (Events and Config, for example) simultaneously (**divert** command)
- Canceling display of output to a console (**flush** command)
- Terminating process output (**halt** command)

Printing ELS Output

You can obtain a hard copy printout of event messages, including startup messages, for any module that is in a DEChub ONE docking station. You must configure the switch so that events are displayed directly, as described in the [Displaying the Event Log](#) section. Then, to print ELS messages, do the following:

Step	Action
1	From the Monitor> prompt, enter els .
2	From the ELS> prompt, enter set output console . ELS messages print immediately on the console printer.

Chapter 12

Configuring Remote Management

Overview

Introduction

You can configure, monitor, and manage DIGITAL VNswitch modules from a remote device using the CLI, a graphical interface, or both. To configure and manage the module remotely using the CLI, you must establish a remote console session. To configure and manage a module using a graphical interface, you must install a MIB-based management system such as the DIGITAL MultiChassis Manager (optional). MultiChassis Manager is a component of the DIGITAL clearVISN network management product.

You must configure certain network parameters and protocols before you can configure a module from a remote console session or from a MIB-based system. This chapter describes how to configure those network connections and protocols.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
Configuring the Module for Remote Console Sessions	12-2
Configuring the Module for MIB-Based Management Systems	12-6
Configuring TCP/IP Host Services	12-20
Displaying TCP/IP Host Services Settings	12-25
Monitoring and Managing TCP/IP Host Services	12-26
Connecting to Other Devices Using Telnet	12-33

Configuring the Module for Remote Console Sessions

This section describes how to configure the module to accept remote console sessions. Remote sessions are established by running Telnet on a remote workstation, PC, or terminal server, and connecting through any of the following interfaces:

- The OBM port on the hub
- The AUI port (DEChub ONE docking station only)
- Any VNswitch module network interface (Ethernet, Fast Ethernet, FDDI, ATM ELAN, ATM Bridge Tunnel, or the VNbus)

Refer to the VNswitch module installation and configuration documentation for information about cable installation.

You must perform the following tasks to configure the module for remote console sessions:

Task	Description
1	Configure TCP/IP Host Services or optionally, IP routing.
2	Configure the OBM port (optional).
3	Configure the AUI port (optional).

Configuring TCP/IP Host Services

You must configure TCP/IP Host Services to identify an IP address that can be used to Telnet to a VNswitch module and establish a remote session. To configure TCP/IP Host Services, refer to the [Configuring TCP/IP Host Services](#) section later in this chapter. If IP routing is configured, TCP/IP HST services is automatically disabled (not used).

Configuring and Viewing OBM Port Connections

If you plan to provide the ability to perform Out of Band Management, you must configure an OBM port. To do so, perform the following tasks:

Task	Description
1	Configure an OBM-IP address.
2	Configure OBM port speed.
3	Configure the Serial Line Interface Protocol (SLIP).

Configuring the Module for Remote Console Sessions

You can also view the current OBM port settings if you need to verify the address, port speed, and ready-to-send (RTS) enable/disable state (when applicable).

Configure the OBM-IP Address

You can set the OBM-IP address of the module from either the installation menu for the hub in which a module is installed, or by using the VNswitch module CLI.

Refer to the appropriate VNswitch module installation and configuration manual for information about how to set the OBM-IP address using the installation menu.

To set the IP address of the module using the CLI, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter set obm ip-address ip-address , where <i>ip-address</i> is the address of the module to which you want to provide access. Setting the address to 0.0.0.0 deletes the current OBM address.
2	Press Return. The OBM-IP address is set.
3	Restart the module if you want the new configuration settings to take effect.

Configuring the OBM Port Speed

The means by which you set the OBM port speed varies, depending on whether you are using a DEChub ONE docking station or a MultiSwitch 900.

MultiSwitch 900

You can set the OBM port speed only from the MultiSwitch 900 installation menu. Refer to the appropriate MultiSwitch 900 documentation for information about how to set the OBM port speed using the installation menu.

DEChub ONE Docking Station

You can set the OBM port speed on a DEChub ONE docking station from either the VNswitch module's installation menu, or by using the VNswitch module CLI.

Refer to the appropriate VNswitch module installation and configuration manual for information about how to set the OBM port speed using the installation menu.

Configuring the Module for Remote Console Sessions

To set the OBM port speed using the CLI, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter set obm speed , where <i>speed</i> is the port speed. The default port speed is 9600 baud.
2	Press Return. The OBM port speed is set.
3	Restart the module if you want the new configuration settings to take effect.

Configuring the OBM RTS State

The means by which you set the OBM port Ready-To-Send (RTS) state varies, depending on whether you are using a DEChub ONE docking station or a MultiSwitch 900.

MultiSwitch 900

You can set the OBM port RTS state (Enabled or Disabled) on a MultiSwitch 900 only from the MultiSwitch 900 installation menu. Refer to the appropriate MultiSwitch 900 documentation for information about how to set the OBM port speed using the installation menu.

DEChub ONE Docking Station

You can set the OBM port RTS state (Enabled or Disabled) on a DEChub ONE docking station from either the VNswitch module's installation menu, or by using the VNswitch module CLI.

Refer to the appropriate VNswitch module installation and configuration manual for information about how to set the OBM port RTS state using the installation menu.

To set the OBM port RTS state using the CLI, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter set obm rts , where <i>rts</i> is either enabled or disabled . The default RTS state is disabled.
2	Press Return. The OBM port RTS state is set.
3	Restart the module if you want the new configuration settings to take effect.

Configuring the Module for Remote Console Sessions

Configuring SLIP

You must configure your Telnet client device to use the Serial Line Interface Protocol (SLIP). No SLIP configuration settings are required on the VNswitch module. Refer to the appropriate terminal server documentation for information about how to do so.

Viewing OBM Port Connections

You can view the current OBM port settings if you need to verify the address, port speed, and ready-to-send (RTS) enable/disable state (when applicable).

To view the current settings, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt, enter <code>list obm ip-address</code> .
2	Press Return. The OBM address, port speed, and RTS state is displayed, and the <code>Config></code> prompt is displayed.

Example

```
Config>list obm
```

```
OBM Address: 0.0.0.0  
OBM Port Speed: 9600  
OBM Port RTS disabled
```

Redirecting Ethernet Interface 11

If you plan to perform in-band management through the AUI port on a DEChub ONE docking station, you must connect the module's front panel Ethernet interface 11 to the docking station's Attachment Unit Interface (AUI) port. If TCP/IP Host Services is configured, you can then log in to the module through the AUI port. Refer to [Chapter 5](#) for information about how to redirect Ethernet interface 11 to the docking station's AUI port.

Configuring the Module for MIB-Based Management Systems

This section describes how to configure the module for remote management using MIB-based management systems.

You must perform the following tasks to configure the module for MIB-based management:

Task	Description
1	Configure TCP/IP Host Services.
2	Configure SNMP.

Configuring TCP/IP Host Services

You must configure TCP/IP Host Services to identify the IP address to which SNMP protocol messages are sent. To configure TCP/IP Host Services, refer to the [Configuring TCP/IP Host Services](#) section later in this chapter.

Configuring SNMP

This section describes how to configure and manage the Simple Network Management Protocol (SNMP) to support MIB-based management systems. SNMP is a communications protocol used to configure devices on the network and collect management information from them. The information collected includes such information as traffic overloads, data throughput, and errors. You can then view the collected information using any MIB-based management system such as the DIGITAL clearVISN MultiChassis Manager product.

This section assumes you are familiar with SNMP concepts and MIB-based management systems.

Adding and Deleting a Community

SNMP includes a default community named “public.” The public community is assigned write-read-trap access by default. You should change the access assigned to the public community to read-trap and add at least one other community with write-read-trap access. [Setting Community Access](#) section later in this chapter for information about how to set access.

Configuring the Module for MIB-Based Management Systems

To add or delete a community to the list of SNMP communities, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter snmp .
2	Press Return. The <code>SNMP Config></code> prompt is displayed.
3	If you want to add a new community, enter add community name , where <i>name</i> is the name you want to assign to the new community. If you want to delete an existing community, enter delete community name , where <i>name</i> is the name of the community you want to delete.
4	Press Return. The following message is displayed: Community name []?
5	Enter a name for the community. The name can be composed of up to 32 alphanumeric characters and cannot include spaces, tabs, or Escape key sequences.
6	Press Return. The community is added or deleted as specified and the <code>SNMP Config></code> prompt is displayed.

Adding and Deleting an Address for a Community

You can specify one or more addresses for each community. However, you can add addresses only one at a time. If you do not specify an address for a community, requests are accepted from any host. The address is also used to specify the hosts that receive the traps. If no address is specified, no traps are generated.

To add an address for a community, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter snmp .
2	Press Return. The <code>SNMP Config></code> prompt is displayed:
3	If you are adding an address, enter add address . If you are deleting an address, enter delete address .
4	Press Return. The following message is displayed: Community name []?

Configuring the Module for MIB-Based Management Systems

Step	Action
5	Enter the name of the community for which you want to assign or delete an address. The name can be composed of up to 32 alphanumeric characters and cannot include spaces, tabs, or Escape key sequences.
6	Press Return. The following message is displayed: IP Address [0.0.0.0]?
7	Enter the IP address you want to add to or delete from the community.
8	Press Return. If you are adding an address, the following message is displayed: IP Mask [255.255.255.255]? If you are deleting an address, go to step 10.
9	Enter the IP mask for the address.
10	Press Return. The address is added or deleted as specified, and the SNMP Config> prompt is displayed.

Creating and Managing MIB Views

You can add a MIB view or delete an existing view.

Adding Views

You can add a new view or add a portion (subtree) of the MIB to an existing view. More than one subtree can be added to an existing view.

To create a new view or add a subtree to an existing view, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter snmp .
2	Press Return. The SNMP Config> prompt is displayed.
3	Enter add subtree .
4	Press Return. The following message is displayed: View name []?

Configuring the Module for MIB-Based Management Systems

Step	Action
5	<p>If you want to create a new view, enter a name that is not already being used.</p> <p>If you want to add a subtree to an existing view, enter the name of an existing view.</p> <p>The name can be composed of up to 32 alphanumeric characters and cannot include spaces, tabs, or Escape key sequences.</p>
6	<p>Press Return. The following message is displayed:</p> <pre>MIB OID name []?</pre>
7	<p>Enter the numeric value of the MIB Object ID (OID) for which you want to create a new view or add to an existing view.</p> <p>Example: MIB OID name []?1.3.6.1.4.1.2</p>
8	<p>Press Return. The view is created or added to an existing view, and the <code>SNMP Config></code> prompt is displayed.</p>

Deleting Subtrees and Views

To delete a subtree and view, perform the following steps:

Step	Action
1	<p>At the Config prompt (<code>Config></code>), enter <code>snmp</code>.</p>
2	<p>Press Return. The <code>SNMP Config></code> prompt is displayed.</p>
3	<p>Enter <code>delete subtree</code>.</p>
4	<p>Press Return. The following message is displayed:</p> <pre>View name []?</pre>
5	<p>Enter the name of the view you want to delete or the name of the view containing the subtree you want to delete.</p>
6	<p>Press Return. The following message is displayed:</p> <pre>MIB OID name []?</pre>
7	<p>Enter the numeric value of the MIB Object ID (OID) associated with the view you want to delete.</p> <p>Example: MIB OID name []?1.3.6.1.4.1.2</p>

Configuring the Module for MIB-Based Management Systems

Step	Action
8	Press Return. If you are deleting the last subtree in the view, the subtree and its associated view are deleted, and the <code>SNMP Config></code> prompt is displayed. If subtrees other than the one you are deleting are associated with the specified view, only the specified subtree is deleted and the <code>SNMP Config></code> prompt is displayed. The specified view and all remaining subtrees associated with the view remain intact.

Setting and Managing Community Views and Access

You can assign one or more views and one of three access types to a community.

Setting Community Views

To assign a view to a community, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter snmp .
2	Press Return. The <code>SNMP Config></code> prompt is displayed.
3	Enter set community view .
4	Press Return. The following message is displayed: <code>Community name []?</code>
5	Enter the name of the community to which you want to assign a view.
6	Press Return. The following message is displayed: <code>View name []?</code>
7	If you want to assign all MIB views (subtrees) to the community, enter the name of the view you want to assign to the community followed by a space and all . If you want to assign a selected MIB view (subtree) to the community, enter the name of the view you want to assign to the community followed by a space and the subtree number.
8	Press Return. The specified views are assigned to the community and the <code>SNMP Config></code> prompt is displayed.

Configuring the Module for MIB-Based Management Systems

Setting Community Access

The default community named “public” is assigned write-read-trap access by default. You should change the access assigned to the public community to read-trap, and add at least one other community with write-read-trap access.

To set read, write, or trap access to a community, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter snmp .
2	Press Return. The <code>SNMP Config></code> prompt is displayed.
3	Enter set community access <i>command-option</i> , where <i>command-option</i> is one of the community access options listed in Table 12-1 .
4	Press Return. The specified access option is assigned to the community and the <code>SNMP Config></code> prompt is displayed.

Table 12-1: Community Access Options

Command Option	Description
read-trap	Sets read access and trap generation for the specified community.
write-read-trap	Sets write and read access and trap generation for the specified community.
trap-only	Sets the trap port to the specified community.

Setting the Trap Port

To specify the User Datagram Protocol (UDP) trap port to which traps are sent, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter snmp .
2	Press Return. The <code>SNMP Config></code> prompt is displayed.
3	Enter set trap_port .
4	Press Return. The following message is displayed: UDP port [162]?

Configuring the Module for MIB-Based Management Systems

Step	Action
5	Enter the UDP port number for the trap port that is to receive traps. The default is the standard trap port (162).
6	Press Return. The specified UDP port number is assigned and the SNMP Config> prompt is displayed.

Enabling and Disabling SNMP

To enable or disable SNMP, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter snmp .
2	Press Return. The SNMP Config> prompt is displayed.
3	If you want to enable SNMP, enter enable snmp . If you want to disable SNMP, enter disable snmp . The factory default is enabled.
4	Press Return. SNMP is enabled or disabled as specified, and the SNMP Config> prompt is displayed.
5	Restart the module for the new configuration settings to take effect.

Enabling and Disabling Traps

You can enable or disable one trap at a time or all traps simultaneously. To enable or disable a trap, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter snmp .
2	Press Return. The SNMP Config> prompt is displayed.
3	If you want to enable a trap, enter enable trap trap-type . If you want to disable a trap, enter disable trap trap-type . The <i>trap-type</i> must be one of the command options specified in Table 12-2 .
4	Press Return. The following message is displayed: Community name []?

Configuring the Module for MIB-Based Management Systems

Step	Action
5	Enter the name of the community to which you want to assign the trap type.
6	Press Return. The specified trap type is enabled or disabled for the specified community, and the <code>SNMP Config></code> prompt is displayed.

Table 12-2: Trap Type Options

Command Option	Description
<u>cold_start</u>	A cold start trap (0) occurs when the transmitting device is reinitializing and the agent's configuration or the protocol entity implementation may be altered.
<u>warm_start</u>	A warm start trap (1) occurs when the transmitting device is reinitializing but the agent's configuration or the protocol entity implementation does not change.
<u>link_down</u>	A link down trap (2) occurs when there is a failure in one of the communication links represented in the agent's configuration.
<u>link_up</u>	A link up trap occurs when a previously inactive network link becomes active.
<u>auth_fail</u>	An authentication failure trap occurs when the sending entity is the addressee of a protocol message that is not properly authenticated.
<u>egp</u>	EGP neighbor loss traps occur when an EGP neighbor and peer are marked down and no longer a peer. The <code>egpNeighborLoss</code> trap-PDU contains the name and value of the <code>egpNeighAddr</code> instance for the affected neighbor as the first element of its variable bindings.
<u>enterprise</u>	Enterprise-specific traps occurrences vary, depending on the event. The <code>specific-trap</code> field identifies the particular trap that occurs.
<u>all</u>	Enables or disables all traps specified in this table.

Configuring the Module for MIB-Based Management Systems

Displaying and Monitoring SNMP

This section describes how to display current SNMP configuration parameters and how to monitor statistics about the number of defined variables and the size of the MIB.

Monitoring Defined Variables and MIB Size

To display information about the number of defined variables and the size of the MIB, perform the following steps:

Step	Action
1	Access the Monitor prompt (Monitor>).
2	Enter <code>statistics</code> .
3	Press Return. The desired information and the Monitor prompt (Monitor>) are displayed.

Example

```
Monitor>statistics
```

```
Number of defined variable handlers = 776  
Size of MIB = 35496 bytes
```

Displaying Current Configuration Parameters

You can display information about the following parameters assigned to communities:

- Access
- IP address and mask
- Trap types
- Views and subtrees associated with a views
- All the above information in one report

Configuring the Module for MIB-Based Management Systems

Displaying Selected Parameters

To display selected parameters assigned to communities, perform the following steps:

Note

You can display parameters by community from both the `SNMP Config>` prompt, as described in this section, and from the Monitor prompt (`Monitor>`). To display the same information from the Monitor process, access the Monitor prompt and enter `list community option` as described in the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter <code>snmp</code> .
2	Press Return. The <code>SNMP Config></code> prompt is displayed.
3	Enter <code>list community option</code> , where <i>option</i> is one of the command options specified in Table 12-3 .
4	Press Return. The information associated with the specified option is displayed and the <code>SNMP Config></code> prompt is displayed.

Table 12-3: Community Command Options

Command Option	Description
<code>access</code>	Displays the access assigned to each community.
<code>address</code>	Displays the IP address and IP mask assigned to each community.
<code>traps</code>	Displays the trap types enabled for each community.
<code>view</code>	Displays the views assigned to each community.

Configuring the Module for MIB-Based Management Systems

Displaying All Parameters in a Report

To display all parameters assigned to communities, perform the following steps:

Note

You can display all parameters from both the `SNMP Config>` prompt, as described in this section, and from the `Monitor` prompt (`Monitor>`). To display the same information from the `Monitor` process, access the `Monitor` prompt and enter **`list all`**.

Step	Action
1	At the <code>Config</code> prompt (<code>Config></code>), enter <code>snmp</code> .
2	Press Return. The <code>SNMP Config></code> prompt is displayed.
3	Enter <code>list all</code> .
4	Press Return. The report and the <code>SNMP Config></code> prompt is displayed.

Configuring the Module for MIB-Based Management Systems

Example

SNMP Config>**list all**

SNMP is disabled

Trap UDP port: 162

Community Name	Access
public	Read, Trap
salem	Trap Only

Community Name	IP Address	IP Mask
public	All	N/A
salem	All	N/A

Community Name	Enabled Traps
public	None
salem	Cold Restart

Community Name	View
public	All
salem	lancaster

View Name	Sub-Tree
lancaster	1.3.6.1.4.1.2
	1.3.6.1.4.1
	1.3.6.1.4.1.3

Identifying the Location of the Module

You can specify the physical location of the VNswitch module if it is serving as an SNMP node. For example, you may want to indicate that a module is located in Building 2, third floor, room 2. The information you specify is available to SNMP-based management agents as a MIB object. The information you can enter to identify the location is limited to 80 characters.

To identify the location of the module for MIB-based management systems, perform the following steps:

Step	Action
1	At the Config> prompt, enter set location .
2	Press Return. The following message is displayed: Location of this node []?

Configuring the Module for MIB-Based Management Systems

Step	Action
3	Enter the location of the module. (Building 2 , for example.) The location description you enter can be up to 80 characters long.
4	Press Return. The information you enter is set.

Example

```
Config>set location Building 2, 3rd floor
```

Identifying a Contact Person Responsible for the Switch

You can specify the name of a contact person who is responsible for the VNswitch module, if it is serving as an SNMP node. For example, you may want to indicate that Jane Doe is the switch administrator. You may also decide to provide a phone number to contact the administrator. The contact information you specify is available to SNMP-based management agents as a MIB object. The information you can enter to identify the contact person is limited to 80 characters.

To identify the contact person for MIB-based management systems, perform the following steps:

Step	Action
1	At the Config> prompt, enter set contact-person .
2	Press Return. The following message is displayed: Contact person for this node []?
3	Enter the information needed to identify the person responsible for the module (Jane Doe, ext. 5555 , for example). The contact description you enter can be up to 80 characters long.
4	Press Return. The information you enter is set.

Example

```
Config>set contact-person Jane Doe, ext. 5555
```

Configuring the Module for MIB-Based Management Systems

Displaying Contact Person and Location Information

You can display the name of a contact person who is responsible for the VNswitch module and the location of the module for a module serving as an SNMP node. You may want to do so, for example, to verify that the information is correct.

To display contact person and location information for MIB-based management systems, perform the following steps:

Step	Action
1	At the <code>Config></code> prompt, enter list configuration .
2	Press Return. A report is displayed, including information about the contact person for, and location of, the module.

Configuring TCP/IP Host Services

VNswitch modules can serve as IP end nodes in a network. As an IP end node, a module supports several IP protocols including, for example, Telnet, SNMP, TFTP, and Ping. These protocol functions are available through TCP/IP Host Services. Telnet and SNMP provide two vehicles for remotely managing modules. TFTP is used to back up and restore switch configurations. Ping is used to diagnose network problems. Refer to [Chapter 10](#) for information about backing up and restoring switch configurations.

When the VNswitch module is to serve as an IP end node, you must perform the following tasks before any of these IP protocols (Telnet, SNMP, and so on) can be used:

Task	Description
1	Configure the module's in-band IP address and subnet mask using either TCP/IP HST services or IP routing. The following procedure explains the steps involved when you use TCPI/IP HST services. To configure IP routing, see the <i>DIGITAL VNswitch 900 Series Router Management</i> manual.
2	Select the VSD on which you want a VNswitch IP end node to reside.
3	Configure a default gateway.
4	Enable or disable router discovery and RIP listening.
5	Enable or disable TCP/IP Host Services.
6	Restart the module.

Configuring the Module's In-Band IP Address and Subnet Mask

You can set, change, or delete the in-band IP address of the module from either the DEChub ONE docking station or MultiSwitch 900 installation menu, or by using the VNswitch module CLI.

Refer to the appropriate module installation and configuration manual for information about how to set the in-band address using the installation menu.

When changing the IP address, the VSD associated with the IP-HST address might need to be changed. The VNswitch CLI allows you to set both of these values in a single operation. However, the DEChub 900 ONE or DEChub 900 installation menus do not allow the configuration of the VSD associated with the IP-HST address. Changing the IP-HST address takes effect immediately.

Assigning a VNswitch IP End Node to a VSD

Assume, for example, each of three VSDs represents a different IP subnet. You must select the VSD on which the IP end node is to reside. If you select VSD 1, then you must assign an IP address and subnet mask to the IP end node that is valid for VSD 1's subnet. IP nodes directly attached to VSD 1's subnet can initiate an ARP request for, and connect to, the module directly. IP nodes connected to other subnets (VSD 2, for example) must connect to the module through a router.

You do not need to assign an IP address to the VNswitch module or assign the address to a VSD if you do not want to use any of the IP end node protocols' features (remote management, network backup and restoration, and so on).

Restrictions

Assigning an IP host to a VSD affects only in-band management (IP packets received over Ethernet, Fast Ethernet, FDDI, or ATM network interfaces). IP packets received through the OBM port of a MultiSwitch 900 or a DEChub ONE docking station are not affected. Packets received through an OBM port are not subject to VSD or bridge filtering.

To set the IP address and subnet mask of the module using the CLI, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter hst .
2	Press Return. The TCP/IP-Host config> prompt is displayed.
3	Enter set ip-host and press Return. The following message is displayed: IP-Host address [16.20.48.48]?

Configuring TCP/IP Host Services

Step	Action
4	Enter the IP address of the module and press Return. The following message is displayed: Address mask [255.0.0.0]?
5	Enter the address subnet mask and press Return. The default subnet mask is 255.0.0.0 if the address is a Class A IP address, 255.255.00 if the address is a Class B address, and 255.255.255.0 if the address is a Class C address. The IP address and address mask are set, and the TCP/IP-Host config> prompt is displayed.

Configuring a Default Gateway

A default network gateway may be required if routers are included in your network. The default network gateway is used when trying to send packets to IP destinations that are not on the same subnet as the VNswitch module. All packets on this router whose destination address is not found in the routing table, are routed to this destination.

The default gateway setting is not required if Router Discovery is enabled. It is also not typically needed if RIP listening is enabled. For example, if your system receives RIP messages containing a default route, then you do not need to configure a default gateway.

To set or change default gateways for the module using the CLI, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter hst .
2	Press Return. The TCP/IP-Host config> prompt is displayed.
3	Enter add default-gateway and press Return. The following message is displayed: Default-Gateway address [0.0.0.0]?
4	Enter the IP address of the router to be used as the default gateway, and press Return. The address of the default gateway is added and the TCP/IP-Host config> prompt is displayed.

Enabling and Disabling Router Discovery and RIP Listening

You may want to enable Router Discovery or RIP Listening if routers are included in your network. Router Discovery is the ability of the VNswitch module to learn of default routers by receiving ICMP Router Discovery messages. RIP Listening is the

Configuring TCP/IP Host Services

process of building routing table entries listening to the RIP protocol. Although you can enable both Router Discovery and RIP Listening at the same time, typically only one or the other is needed.

To enable or disable Router Discovery or RIP Listening using the CLI, perform the following steps:

Step	Action
1	At the Config prompt (<code>Config></code>), enter hst .
2	Press Return. The <code>TCP/IP-Host config></code> prompt is displayed.
3	If you want to enable or disable Router Discovery, enter either enable or disable , followed by router-discovery . The default is enabled. If you want to enable or disable RIP Listening, enter either enable or disable , followed by rip-listening . The default is disabled.
4	Press Return. The process is enabled or disabled as specified, and the <code>TCP/IP-Host config></code> prompt is displayed.

Enabling and Disabling Host Services

TCP/IP Host Services must be enabled to establish in-band remote console sessions. You should disable host services only if you want to limit access to the module via a local session through the setup port.

Configuring TCP/IP Host Services

To enable or disable TCP/IP Host Services using the CLI, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter hst .
2	Press Return. The TCP/IP-Host config> prompt is displayed.
3	If you want to enable Host Services, enter enable services . If you want to disable Host Services, enter disable services . The default is enabled.
4	Press Return. TCP/IP Host Services is enabled, and the TCP/IP-Host config> prompt is displayed Note: Configuration of IP-HST services takes effect immediately.

Note

If IP routing is enabled, IP-HST services are disabled automatically. Once routing is enabled, connectivity to the VNSwitch must be made through one (or more) of the configured routing interfaces. (See the *DIGITAL VNswitch 900 Series Router Management* manual.) When routing is enabled for the first time, any configuration information for IP-HST services can be automatically copied to the default routing interface.

Displaying TCP/IP Host Services Settings

You can display all TCP/IP Host Services configuration settings. To display the information, perform the following steps:

Step	Action
1	At the Config prompt (Config>), enter hst .
2	Press Return. The TCP/IP-Host config> prompt is displayed.
3	Enter list all , and press Return. The settings are displayed and the TCP/IP-Host config> prompt is displayed.

Example

```
TCP/IP-Host config>list all

IP-Host IP address : 16.20.48.48
Address Mask      : 255.0.0.0
Default Gateway IP-address(es)
1.2.3.4
TCP/IP-Host Services Enabled.
TCP/IP-Host Services offered on VSD 1, DEFAULT.
RIP-LISTENING Disabled.
Router Discovery Enabled
```

Monitoring and Managing TCP/IP Host Services

This section describes how to use TCP/IP Host Services to monitor and manage network connections. You can perform the following tasks using TCP/IP Host Services:

- Display routing tables.
- Display the bridge's interface addresses.
- Test a network connection using the **ping** command.
- Display a list of routers.
- Display a path to a destination device.
- Display the VSD on which TCP/IP host services is available.

Displaying the Routing Table

You can display the routing table generated through Router Discovery and RIP Listening. Refer to the [Enabling and Disabling Router Discovery and RIP Listening](#) section for information about configuring these functions.

To display the routing table, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter hst .
2	Press Return. The TCP/IP Host> prompt is displayed.
3	Enter dump .
4	Press Return.
5	The routing table and the TCP/IP Host> prompt are displayed.

Example

```
TCP/IP Host>dump
```

```
Type  Dest net      Mask      Cost    Age  Next hop(s)
Dir*  11.7.0.0      FFFF0000  1       0    VLAN/n
Routing table size: 768 nets (61440 bytes), 1 nets known
```

Displaying the Bridge's Interface Addresses

You can display the transparent bridge's interface number and IP address mask. To display the bridge's interface addresses, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter hst .
2	Press Return. The TCP/IP Host> prompt is displayed.
3	Enter interface addresses .
4	Press Return.
5	The bridge's interface addresses and the TCP/IP Host> prompt are displayed.

Example

TCP/IP Host>**interface addresses**

```
interface addresses
Interface IP Address(es) Mask(s)
VLAN/n 11.7.7.10 255.255.0.0
```

Testing a Network Connection Using the Ping Command

You can use the **ping** command to help isolate problems in an internetwork environment. When initiated, **ping** causes the bridge to send ICMP Echo Requests once per second to a specified destination, and to detect a response.

This process is done continuously, incrementing the ICMP sequence number with each additional packet. Matching received ICMP Echo responses are reported with their sequence number and the round trip time. The granularity (time resolution) of the round trip time calculation is (depending on the platform) usually about 20 milliseconds.

Note

The number of data bytes in the ICMP message, excluding the ICMP header, is 56 bytes, and the TTL used is 60 hops.

Monitoring and Managing TCP/IP Host Services

To **ping** a device, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter hst .
2	Press Return. The TCP/IP Host> prompt is displayed.
3	Enter ping ip-address , where <i>ip-address</i> is the address of the device for which you want to perform a test.
4	Press Return. The test is initiated.
5	Press any key. The process is completed. The results of the response and the TCP/IP Host> prompt are displayed.

Example

```
TCP/IP Host>ping 16.20.48.79
```

```
PING 16.20.48.79: 56 data bytes
----16.20.48.79 PING Statistics----
3 packets transmitted, 0 packets received, 100% packet
loss
```

Note

If IP routing is enabled, you can enter the ping command from the IP Protocol Console (IP>), accessed from the Monitor> prompt.

Displaying a List of Routers

You can display a list of currently known routers that were learned through one of the following operations:

- Static configuration (Refer to the [Configuring a Default Gateway](#) section.)
- Received ICMP redirects
- ICMP Router Discovery messages, if configured (Refer to the [Enabling and Disabling Router Discovery and RIP Listening](#) section.)
- RIP updates, if configured (Refer to the [Enabling and Disabling Router Discovery and RIP Listening](#) section.)

Monitoring and Managing TCP/IP Host Services

To display a list of currently known routers, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter hst .
2	Press Return. The TCP/IP Host> prompt is displayed.
3	Enter routers .
4	Press Return.
5	The list of routers and the TCP/IP Host> prompt are displayed.

Example

```
TCP/IP Host>routers
```

```
routers
```

```
Type Router          Priority    Lifetime Accessed
```

Displaying the Path to a Destination Device

You can display the complete path (hop by hop) from the VNswitch module to a particular destination device. This is an implementation similar to the UNIX traceroute tool. For each successive hop, the process sends out three probes and prints the IP address of the responder, together with the round trip time associated with the response. If a particular probe does not receive a response, an asterisk is printed. Each line in the display relates to this set of three probes, with the left-most number indicating the distance from the router executing the command (in router hops).

To display the path to a destination, perform the following steps:

Step	Action
1	At the Monitor prompt (Monitor>), enter hst .
2	Press Return. The TCP/IP Host> prompt is displayed.
3	Enter trace address , where address is the IP address of the device to which you want to display the path.
4	Press Return.

Monitoring and Managing TCP/IP Host Services

Step	Action
5	The trace is initiated. The trace is completed when the destination is reached, an ICMP Destination Unreachable is received, or the path length reaches 32 router hops. The <code>TCP/IP Host></code> prompt is redisplayed when the trace is completed.

Note

If routing is enabled, you perform this function using the **traceroute** command of the IP Protocol Console (IP>), available from the `Monitor>` prompt.

Interpreting Unexpected Results

When the probe receives an unexpected result, the following indications can be printed:

- **!N** — Indicates an ICMP Destination Unreachable (network unreachable) was received.
- **!H** — Indicates an ICMP Destination Unreachable (host unreachable) was received.
- **!P** — Indicates an ICMP Destination Unreachable (protocol unreachable) was received.

Since the probe is a UDP packet sent to a strange port, a port unreachable is expected. The exclamation point (!) indicates that the destination was reached, but the reply sent by the destination was received with a TTL of 1. This result typically indicates an error in the destination, prevalent in some versions of UNIX. The destination is inserting the probe's TTL in its replies. This results in a number of lines consisting of a number of asterisks before the destination is finally reached.

Displaying the VSD on Which TCP/IP Host Services Is Available

You can display the number and name of the VSD on which Host Services is available. This is the VSD to which you assigned the VNswitch module's IP Host address. The default is the default VSD number 1. (Refer to the [Configuring a Default Gateway](#) section for information about assigning the module's IP Host address to a VSD.)

To display the number and name of the VSD on which Host Services is available, perform the following steps:

Step	Action
1	At the Monitor prompt (<code>Monitor></code>), enter hst .
2	Press Return. The <code>TCP/IP Host></code> prompt is displayed.
3	Enter vsd .
4	Press Return.
5	The number and name of the VSD and the <code>TCP/IP Host></code> prompt are displayed.

Monitoring and Managing TCP/IP Host Services

Example

```
TCP/IP Host>vsd
```

```
TCP/IP-Host Services offered on VSD 1, DEFAULT.  
VNSwitch not in Hub--VSDs are effectively local!
```

Connecting to Other Devices Using Telnet

You can connect to any remote system from a VNswitch module's local or remote session by using Telnet. To do so, perform the following steps:

Step	Action
1	Access the Main prompt (Main>).
2	Enter telnet ip-address terminal-type , where <i>ip-address</i> is the IP address of the device you want to access, and <i>terminal-type</i> is an optional parameter you can enter to identify the type of terminal you are using. The <i>terminal-type</i> can be up to 80 characters long, and helps define the characteristics of the Telnet session for the remote device.
3	Press Return. The VNswitch module attempts to establish a session with the remote device.

Chapter 13

Monitoring Network Activity

Overview

Introduction

The DIGITAL VNswitch provides two methods for monitoring network activity:

- **Mirror interface** — This feature allows you to diagnose network activity by sampling incoming and outgoing packets from a specified VNswitch interface.
- **RMON agent** — This feature allows you to configure the VNswitch so that it independently monitors its own MIB variables and network traffic.

This chapter explains how to configure and monitor the mirror interface and the RMON agent.

In This Chapter

The following topics are covered in this chapter:

Topic	Page
About the VNswitch Mirror Feature	13-2
Mirror Restrictions and Mirroring Effect on VNswitch Performance	13-3
Configuring the Mirror Interface	13-4
Mirroring Multiple VNswitch Modules	13-8
Mirror Commands	13-10
About the VNswitch RMON Agent	13-11
RMON Command Line Interface	13-14
RMON Example	13-18

About the VNswitch Mirror Feature

The VNswitch 900 mirror feature lets you diagnose network activity by sampling incoming and outgoing packets from a specified VNswitch interface. Packets received or transmitted on specified interfaces can be copied to the mirror destination.

The mirror destination interface is a dedicated interface that accepts only mirrored packets. These mirrored packets are derived from specific VNswitch Ethernet-like interfaces and can be forwarded to any destination interface on the VNswitch with an Ethernet interface. By attaching an external network analyzer to the destination interface, you can evaluate packets and troubleshoot the network.

Mirror Restrictions and Mirroring Effect on VNswitch Performance

Mirror Restrictions

The following restrictions apply to the mirror feature:

- Routing must be disabled on VLANs containing interfaces that are mirrored.
- All mirror interfaces must be Ethernet-like interfaces (LAN Bridge Tunnels, LANEs, 10 Mb/s Ethernet, 100 Mb/s Ethernet).
- Destination interfaces must be dedicated to mirroring.
- Ethernet mirror interfaces cannot be configured simultaneously as a receive/transmit interface and a destination interface.
- Packets filtered by hardware, a runt packet for example, are not copied to the destination interface.
- The mirroring interfaces must reside in a VLAN containing only Ethernet-like ports, including: 10 Mb/s Ethernet, 100 Mb/s Fast Ethernet, Ethernet Bridge Tunnels, LAN Emulation (LANE).

For example, a VNswitch 900EF installed in a hub with the VNbus contains 12 Ethernet ports and one FDDI port. When powered up, the VNswitch 900EF defaults all ports including the FDDI port to one VLAN. For the mirror to operate in the default VLAN, the FDDI port must be removed and assigned to another VLAN.

Mirroring Effect on VNswitch Performance

Mirroring does not adversely affect the performance of the VNswitch. If more packets are sent to the destination than can be handled, the packets are dropped. Although the packets will not be available for viewing, operation of the VNswitch is unaffected.

Configuring the Mirror Interface

Accessing the Mirror Prompt

You can access the mirror feature either at the Configuration prompt (`Config>`), or the Monitor prompt (`Monitor>`). To access the mirror feature from the `Config>` prompt, perform the following steps:

Step	Action
1	To access the mirror prompt, at the <code>Config></code> prompt, enter <code>mirror</code> .
2	Press Return. The Mirror prompt (<code>Mirror Config></code>) is displayed.

Disabling Bridging on the Destination Interface

Before the destination interface is activated, bridging on the specified destination interface must be disabled. If you try to enable a mirror destination while bridging for the destination interface is active, you are prompted with the following message:

```
The interface that you are setting as the mirror destination is currently
functioning as an active bridge port. It is a requirement of the mirror
functionality that bridging is disabled. Would you like to proceed? If
you proceed, bridging will be disabled on this port and it will be set
as the mirror destination.
```

Once bridging is disabled on the destination interface and mirroring is enabled and you attempt to enable bridging on the destination interface, you are prompted with the following message:

```
The interface that you are attempting to enable is currently set up as
the destination mirror, to proceed, this interface must be disabled as
the destination mirror. Would you like to proceed? If you answer yes,
bridging will be enabled and this interface will no longer be the mirror
destination.
```

Interaction with SNMP

You cannot activate a mirror destination port and have the bridge port enabled simultaneously. Therefore, if a port is configured as the mirror destination and someone attempts to use SNMP to enable the bridge port (for example, by setting the `dot1dStpPortEnable.X` bit) the SNMP set will fail. The only way to reconfigure the mirror port destination is through the local CLI of the VNswitch.

Enabling Bridging After Using the Mirror Interface

Disabling the mirror destination does not automatically re-enabling bridging. To re-enable bridging on the destination interface, use the `BRIDGE config> set port` command. However, when the mirror destination is disabled, the you are prompted to determine if bridging should be enabled. If you answer YES, bridging is re-enabled.

Configuring the Mirror Interface

When you disable the destination port, you are prompted with the following message:
Bridging is currently disabled on interface n. Do you want to enable bridging?

Setting the Mirror Interface

You can configure the mirror interface to mirror the following types of packets:

- Receive packets
- Transmit packets
- Both receive and transmit packets
- None

In addition to setting a mirror interface, you must also set a mirror destination interface using the **enable** command. Only one destination interface per VNswitch module is allowed.

The steps in the following example set a mirror interface to mirror the packets received and transmitted on Ethernet interface 2, and sets the destination interface to 5:

Step	Action
1	At the <code>Mirror Config></code> prompt, enter set both 2 .
2	Press Return. A mirror interface on interface 2 is set to transmit and receive.
3	To set a mirror destination, at the <code>Mirror Config></code> prompt, enter enable 5
4	Press Return. A mirror destination is set on interface 5.

Setting Multiple Mirror Interfaces

With the mirror feature, you can simultaneously set multiple interfaces with the **set** command. For example, you can set all 12 Ethernet ports on the VNswitch 900EX to a mirror interface with one command.

The steps in the following example set mirror interfaces to transmit on Ethernet interfaces 1 through 5, 7 through 12, and enables the destination interface 6:

Step	Action
1	At the <code>Mirror Config></code> prompt, enter set transmit 1-5, 7-12 . Press Return.

Configuring the Mirror Interface

Step	Action
2	At the <code>Mirror Config></code> prompt, enter enable 6 . Press Return. A mirror destination is set on interface 6.

Changing the Mirror Destination Port

In the example above, Ethernet interfaces 1 through 5 and 7 through 12 are set to mirror transmit packets, and the destination is enabled for Ethernet interface 6. The steps in the following example change the destination from interface 6 to interface 3, and mirrors traffic transmitted on interface 6:

Step	Action
1	At the <code>Mirror Config></code> prompt, enter set neither 3 . Press Return.
2	At the <code>Mirror Config></code> prompt, enter set transmit 6 . Press Return.
3	At the <code>Mirror Config></code> prompt, enter enable 3 . Press Return. The mirror destination interface is now changed to interface 3.

Listing Mirror and Destination Interfaces

You can display all the receive and transmit mirror interfaces, and the destination interface using the **list** command. To list all the mirror interfaces, perform the following steps:

Step	Action
1	At the <code>Mirror Config></code> prompt, enter list .
2	Press Return. All mirror interfaces are displayed.

Clearing Mirror and Destination Interfaces

You can clear all the mirror interfaces and the destination interface with the **clear** command. To clear all the mirror interfaces, perform the following steps:

Step	Action
1	At the <code>Mirror Config></code> prompt, enter clear .
2	Press Return. All mirror interfaces and the destination port are cleared.

Configuring the Mirror Interface

You can also clear specific mirror interfaces using the **set** command. To clear specific mirror interfaces, perform the following steps:

Step	Action
1	At the <code>Mirror Config></code> prompt, enter set neither 3,4
2	Press Return. Packets on mirror interfaces 3 and 4 are not copied to the destination interface.

Mirroring Multiple VNswitch Modules

You can configure your mirror destination interface for each VNswitch to connect to a repeater. This section gives an example of how to mirror multiple VNswitch modules with one network analyzer on a DIGITAL DECreeper. The configuration criteria for this example are:

- The VNswitch module must be configured to conform to the requirements in the [Mirror Restrictions](#) section.
- The VNswitch modules must reside in the same MultiSwitch 900.
- A DIGITAL DECreeper must be installed in the MultiSwitch 900.
- DIGITAL clearVISN software must be available.

To configure multiple VNswitch modules using one network analyzer ([Figure 13-1](#)), perform the following steps:

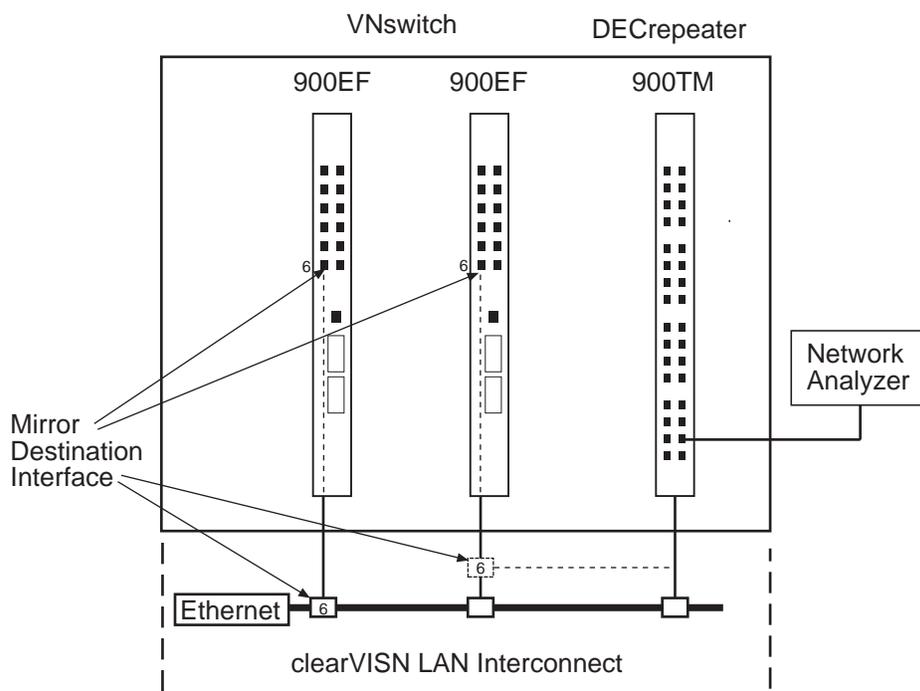
Step	Action
1	Configure your first VNswitch to the mirror interfaces and destination interface you want to use.
2	Configure each additional VNswitch to the mirror interfaces and destination interface you want to use.
3	Ensure that your DIGITAL DECreeper has enough available ports to handle each destination interface and a port to connect a network analyzer.
4	Using DIGITAL clearVISN software, connect the destination interface from the first VNswitch to an available repeater port.
5	Repeat step 4 for each VNswitch you want to configure.
6	Connect a network analyzer to an available repeater port.

Your network analyzer is now capable of receiving packets from each VNswitch connected to the repeater.

To avoid losing packets to the network analyzer, ensure that the combined traffic from all destination interfaces connected by clearVISN to the DECreeper does not exceed 10 Mb/s.

Mirroring Multiple VNswitch Modules

Figure 13-1: Using clearVISN to Connect a Mirror Destination Interface



LKG-10396-97MF

Mirror Commands

The following table provides a list of the mirror commands, with a description and example of each:

Command	Description
<u>C</u>lear	Clears the destination interface, all the transmit interface, and all the receive interfaces. Example: Mirror Config> <u>c</u>lear
<u>E</u>nable/ <u>D</u>isable	Enables/disables the mirror destination interface to which mirror packets will be sent. Only one destination interface per VNswitch are allowed. Example 1: Mirror Config> <u>e</u>nable 5 Enables interface 5 as the destination interface. Example 2: Mirror Config> <u>d</u>isable Disables the destination interface.
<u>L</u>ist	Displays the destination interface, all the transmit interfaces and all the receive interfaces. Example: Mirror Config> <u>l</u>ist
<u>S</u>et	Sets a mirror interface on an Ethernet port. The interface can be set to mirror receive packets, transmit packets, both or none. When an interface is set to none, the mirror interface is disabled. Example 1: Mirror Config> <u>s</u>et <u>t</u>ransmit 1-5 Copies transmit packets from interfaces 1 through 5 to the destination interface. Example 2: Mirror Config> <u>s</u>et <u>r</u>ecieve 2,3,4 Copies receive packets from interfaces 2, 3, and 4 to the destination interface. Example 3: Mirror Config> <u>s</u>et <u>b</u>oth 1-12 Copies both transmit and receive packets from interfaces 1 through 12 to the destination interface. Example 4: Mirror Config> <u>s</u>et <u>n</u>either 6 Disables mirror interface 6.

About the VNswitch RMON Agent

The VNswitch Remote Network Monitoring (RMON) agent allows you to configure the VNswitch so that it independently monitors its own MIB variables and network traffic. The VNswitch RMON agent supports the Alarm and Event MIB groups and adheres to RMON MIB RFC 1757 for Ethernet objects.

RMON Alarm and Event Groups

The alarm group allows you to configure the VNswitch so that it monitors its own MIB variables. If the value of a monitored variable crosses its configured thresholds, the RMON agent generates an event. The event group associates an event with a set of actions. Two actions are defined: generate an SNMP trap message and add an entry to the event group log table.

You can configure alarms and events from a network management application, such as a MIB browser, that uses SNMP. You can also use the VNswitch CLI to read and write MIB variables in the alarm and event groups.

You can separately configure the Event Logging System (ELS) to generate an ELS event whenever an alarm generates an RMON event. Otherwise, the RMON event group and ELS are independent of each other.

Alarms and events are stored in NVRAM (nonvolatile RAM) memory and are preserved if you power cycle the VNswitch. You can delete individual table entries by using the RMON CLI **delete** command or you can use SNMP to set the table entry status to invalid. You can also use the **clear** RMON command from the Config prompt to delete all RMON table entries.

If you use SNMP to create, delete, or modify alarm and event table rows, you must follow the conventions for EntryStatus as specified in the RMON MIB (RFC 1757). You are not required to follow the EntryStatus conventions when you configure alarm and event table rows from the CLI. The CLI correctly transitions row status.

The number of alarm and event table entries is limited to 256.

You cannot write more than one alarm table or event table row at a time in a single SNMP set PDU (Protocol Data Unit). If you do not specify all the values for a row in a set PDU, the default values specified in [Table 13-1](#) and [Table 13-2](#) are used. The CLI uses these default values only the first time you enter an alarm or event from the CLI. Then the CLI uses the values you last entered as a default.

About the VNswitch RMON Agent

[Table 13-1](#) shows the variables and the default values for the **set alarm** command.

Table 13-1: RMON Set Alarm Command Parameters

Alarm Variable	Default Value
Alarm Index	1
Alarm Status	valid
Alarm Interval	1
Alarm Variable	[0.0]
Alarm Sample Type	deltaValue
Alarm Value	0
Startup Alarm	risingOrFallingAlarm
Rising Alarm Threshold	0
Falling Alarm Threshold	0
Rising Event Index	0
Falling Event Index	0
Alarm Owner	Null string
Alarm Description	Null string

About the VNswitch RMON Agent

[Table 13-2](#) shows the variables and the default values for the **set event** command.

Table 13-2: RMON Set Event Command Parameters

Event Variable	Default Value
Event Index	1
Event Status	valid
Event Description	Null string
Event Type	log-and-trap
Event Community	Null string
Event Owner	Null string

RMON Command Line Interface

You can configure RMON alarms and events from a network management application, such as a MIB browser, that uses SNMP. You can also use the VNSwitch CLI to read and write MIB variables in the alarm and event groups.

Accessing the RMON Configuration Process

To access the RMON configuration process using the CLI, perform the following steps:

Step	Action
1	From the Main prompt (Main>), enter the following command: Main> config The Config> prompt is displayed.
2	At the Config> prompt, enter the following command: Config> rmon
3	Press Return. The following prompt is displayed: RMON Config>

RMON Command Line Interface

Once you have entered the RMON configuration process, you can execute the commands in [Table 13-3](#).

Table 13-3: RMON Configuration Commands

Command	Command Parameters
a dd	alarm event
s et	alarm event log-table-max
d el e te	alarm event log
l ist	alarm all event log

Accessing the RMON Monitor Process

To access the RMON monitor process through SNMP, perform the following steps:

Step	Action
1	From the Main prompt (Main>), enter the following command: Main> monitor The Monitor> prompt is displayed.
2	At the Monitor> prompt, enter the following command: Monitor> rmon
3	Press Return. The following prompt is displayed: RMON user console RMON>

RMON Command Line Interface

Once you have entered the RMON monitor process, you can execute the commands in [Table 13-4](#).

Table 13-4: RMON Monitor Commands

Command Type	Command Parameters
<u>d</u>delete	log table
<u>l</u>ist	alarm all event log
<u>e</u>xit	

Clearing RMON Configuration Information

To clear all configuration information for RMON, perform the following steps:

Step	Action
1	From the Main prompt (Main>), enter the following command: Main> config The Config> prompt is displayed.
2	At the Config> prompt, enter the following command: Monitor> clear rmon
3	Press Return. The following message is displayed: You are about to clear all RMON configuration information *** WARNING *** This will invoke an automatic RESTART Are you sure you want to do this (Yes or [No]):
4	If you are certain you want to clear the information, enter Yes .

Displaying RMON Statistics

To display RMON information, perform the following steps:

Step	Action
1	From the Main prompt (Main>), enter the following command: Main> monitor The Monitor> prompt is displayed.
2	At the Monitor> prompt, enter the following command: Monitor> interface statistics n where <i>n</i> is the interface number for which you want statistics.
3	Press Return. RMON statistics information is displayed in the format shown in the example, along with other interface information.

Example

```
Monitor>interface statistics 1
RMON statistics Interface: 1
Octets          256  Drop Events          0
Pkts             4  CRC Align Errs      0
Broadcast Pkts  0    Undersize Pkts      0
Multicast Pkts  0    Oversize Pkts       0
Pkts 64 Octets  4    Fragments           0
Pkts 65-127 Octets 0  Jabbers             0
Pkts 128-255 Octets 0  Collisions          0
Pkts 256-511 Octets 0
Pkts 512-1023 Octets 0  Interval 30 secs Utilization 152 per cent
Pkts 512-1023 Octets 0  Interval 1800 secs Utilization 10 per cent
Monitor>
```

RMON Example

The following sections provide a configuration example. In this example, whenever the number of received SNMP packets increases, the VNswitch generates a *risingAlarm* trap to the community *rmon-trap* that has an IP address of *16.20.48.46*.

Configuring an SNMP Community

The following procedure describes how to configure SNMP with a new community named *rmon-trap* with an IP address of *16.20.48.46*:

Step	Action
1	At the <code>Config></code> prompt, enter snmp . SNMP Config> add comm rmon-trap SNMP Config> add addr rmon-trap IP address [0.0.0.0]? 16.20.48.46 IP Mask [255.255.255.255]?
2	Press Return. The SNMP community <i>rmon-trap</i> is now configured.

Configuring an RMON Trap

The following sections describe how to configure RMON so that it generates a rising alarm trap if the MIB variable *snmpInPkts.0* (oid 1.3.6.1.2.1.11.1.0) increases by more than zero over a 1-second interval. This is accomplished by creating an event entry and an alarm entry.

Creating an Event Entry

To create an event entry with an event type *snmp-trap* and an event community *rmon-trap*, perform the following steps:

Step	Action
1	At the Config> prompt, enter rmon .
2	At the RMON Config> prompt, enter add event .
3	Enter the event description, type, community, and owner. Event Description []? Enter the ascii string for event description and press Return. Event Type [log-and-trap]? snmp-trap . Enter the trap type and press Return. Event Community []? rmon-trap Event Owner []? Enter the ascii string for the event owner and press Return.
4	Press Return. The following is displayed: Creating Event with index <i>n</i> The RMON event entry is complete.

RMON Example

Creating an Alarm Entry

To create an alarm entry for the above event entry, perform the following steps:

Step	Action
1	At the RMON Config> prompt, enter add alarm . Enter the alarm values and descriptions as requested by the prompts. Alarm Interval [1]? Alarm Variable []? 1.3.6.1.2.1.11.1.0 Alarm Sample Type [deltaValue]? Startup Alarm [risingOrFallingAlarm]? risingalarm Rising Alarm Threshold [0]? 1 Falling Alarm Threshold [0]? Rising Event Index [0]? n (Use the value from the add event entry.) Falling Event Index [0]? Alarm Owner []? Alarm Description []?
2	Press Return. The following is displayed: Creating Alarm with index n The RMON alarm entry is complete.

Summary

In the RMON example, whenever the VNswitch receives an SNMP request (snmpInPkts.0 increases by 1 or more), it sends an RMON risingAlarm trap.

Appendix A

Advanced Console Management

Overview

Introduction

This appendix discusses local and remote console session management tasks that may be of interest to advanced DIGITAL VNswitch module users.

In This Appendix

The following topics are covered in this appendix:

Topic	Page
Process Aliases and IDs	A-2
Viewing Process Status and PIDs	A-3
Viewing Output from Multiple Processes	A-5
Canceling Display of Output to a Console Session	A-6
Terminating Process Output	A-7

Process Aliases and IDs

The VNswitch prompts are each associated with an operational process, each of which is assigned a process ID (PID). The Config and Monitor prompts and Events are each associated with a single PID. The VNswitch Main prompt (Main>) is associated with three processes, each of which is associated with a PID. The processes associated with VNswitch prompts, and their PIDs, are shown in the following table:

VNswitch Prompts	Process Names	PIDs
Main	LOpCon	1
	MOpCon and ROpCon	9 and 10
Config	Config	6
Monitor	Monitor	5
Events	Events	2

LOpCon, MOpCon, and ROpCon are the main operator control processes that manage communication between consoles and lower level processes. LOpCon is the process that supports a (single) locally connected console session through the setup port on a DEChub ONE docking station. MOpCon is the process that supports a (single) locally connected console session through the setup port on a MultiSwitch 900. ROpCon is the process that supports up to two remote console sessions. The functions supported by LOpCon, MOpCon, and ROpCon are the same.

Viewing Process Status and PIDs

You can view information about any process including the name of the process, the process ID (PID), the status of the process, the console to which output is currently routed, and the IP address of any user logged in remotely via ROpCon.

To view process status information and PIDs, perform the following steps:

Step	Action
1	At the Main prompt (Main>), enter status .
2	Press Return. A list of processes is displayed, as shown in the following example.

Example

Main>**status**

Pid	Name	Status	TTY	Comments
1	LOpCon	RDY	TTY0	
2	Events	DET	--	Event Logger
3	Tasker	RDY	--	Tasker
4	Debugger	IOW	--	
5	Monitor	DET	--	To Monitor Brouter
6	Config	DET	--	To Configure Brouter
7	SNMP	IDL	--	
8	DEChub	IDL	--	
9	ROpCon	IDL	TTY1	
10	ROpCon	IDL	TTY2	

Refer to the [Process List Output](#) section for a description of each item in the process list.

Process List Output

The process list includes the following information.

Output	Description
Pid	Process ID (PID).
Process	Name of the process. Tasker, MOSDDT, SNMP, and DEChub are processes reserved for use by DIGITAL Services.

Viewing Process Status and PIDs

Output	Description
Status	Status of the process. Refer to the status descriptions in the following table for a list of process status conditions and descriptions.
TTY	Output terminal, if any, to which the process is currently connected. TTY0 indicates a local console. TTY1 or TTY2 indicate Telnet consoles. "Sink" indicates that a process was flushed. Two dashes (- -) indicate that a process was halted.
Comments	The user's login IP address provided when a user is logged in using Telnet (ROpCon).

The following table lists the process status conditions that can be displayed in the Status column, and a description of each.

Status	Description
IDL	The process is idle and waiting for completion of an external event, such as asynchronous I/O.
RDY	The process is ready to run and is waiting to use the CPU.
IOW	The process is waiting for synchronous I/O to complete.
DET	The process has output ready to be displayed and is either waiting to be attached to a display console, or waiting to have its output diverted to a specified console.
FZN	The process is "frozen" due to an error. This usually means the process is trying to use a device that is faulty or incorrectly configured.

Viewing Output from Multiple Processes

You can display the output from several processes simultaneously, at any terminal, by diverting output to the terminal. This capability is most commonly used to redirect event log output messages to a specific terminal while running another process at the same time.

To divert (display) process output to a particular console, perform the following steps:

Step	Action
1	Determine the PID of the process you want to divert, and the number of the console (tty#) to which you want to divert the output. (Refer to the Viewing Process Status and PIDs section.)
2	At the Main prompt (Main>), enter divert , followed by the process ID and the console number. For example, entering divert 2 1 redirects Event Logging System messages generated by the Events process (PID 2) to a remote console (tty1). (Refer to Chapter 11 for information about the Event Logging System.)
3	Press Return. The output from the specified process is diverted to the console indicated.

Canceling Display of Output to a Console Session

You can cancel the display of output from a specific process to your console session. Doing so cancels only the display of output to your session. Other users accessing the same process can still view output from the process. You must use the **config**, **monitor**, **events**, or **divert** command to reaccess the process and redisplay output. (Refer to the [Viewing Output from Multiple Processes](#) section for information about using the **divert** command.)

To cancel (flush) process output to a particular console, perform the following steps:

Step	Action
1	Determine the PID of the process you want to stop displaying. (Refer to the Viewing Process Status and PIDs section.)
2	At the Main prompt (Main>), enter flush , followed by the process ID. For example, entering flush 2 cancels the display of Event Logging System messages generated by the Events process (PID 2).
3	Press Return. The output from the specified process is no longer displayed on your console.

Terminating Process Output

You can terminate all output from a specific process to all console sessions. Doing so cancels not only the display of output to your console session, but also to console sessions of other users. You must use the **config**, **monitor**, **events**, **divert**, or **flush** commands to reaccess the process and redisplay output. (Refer to the [Viewing Output from Multiple Processes](#) section for information about using the **divert** command. Refer to the [Canceling Display of Output to a Console Session](#) section for information about using the **flush** command.)

To terminate process output to all console sessions, perform the following steps:

Step	Action
1	Determine the PID of the process you want to stop displaying. (Refer to the Viewing Process Status and PIDs section.)
2	At the Main prompt (<code>Main></code>), enter halt , followed by the process ID. For example, entering halt 2 terminates the display of Event Logging System messages generated by the Events process (PID 2).
3	Press Return. The output from the specified process is terminated.

Appendix B

Plug and Play Default Settings

Overview

Introduction

The DIGITAL VNswitch family of modules features plug and play operation. Once installed, the module begins handling network traffic running transparent bridging on each port, by default. This appendix lists the default settings that are used when the switch is first installed. You may need to alter the default settings to maximize network performance, to enable special functions such as Virtual Local Area Networks (VLANs), and to set IP addresses for remote management.

In This Appendix

The following topics are covered in this appendix:

Topic	Page
Switch-wide Default Settings	B-3
Nonspecific Interface Default Settings	B-4
Ethernet Interface Default Settings	B-5
Fast Ethernet Interface Default Settings	B-6
FDDI Interface Default Settings	B-7
ATM Interface Default Settings	B-8
VNbus Default Settings	B-11
Bridge Default Settings	B-12
Maintenance Default Settings	B-14

Topic	Page
ELS Configuration Default Settings	B-15
Remote Management Default Settings	B-17

Switch-wide Default Settings

Table B-1: lists the default settings for parameters that affect switch-wide operations.

Table B-1: Switch-wide Defaults

Setting	Default
Enabling and disabling ID and password prompts	Disabled
Inactivity timer	0 (zero) — Turns off the inactivity timer, so that an established session remains active no matter how long a console is inactive. The inactivity timer affects only those consoles linked to modules on which ID and password prompting is enabled.
Hostname	No host name
Time zone offset	0 (zero) — Indicates no offset from GMT.
Time host synchronization frequency	0 (zero) — Indicates the switch is not to poll the time host to synchronize clocks.
Routing	Disabled

Nonspecific Interface Default Settings

Table B-1: lists the default settings common to all interface types (Ethernet, Fast Ethernet, FDDI, and ATM).

Table B-1: Nonspecific Interface Defaults

Setting	Default
Enabling and disabling an interface	Enabled
ARP Refresh Timer	20 minutes
ARP Auto-Refresh	Disabled

Ethernet Interface Default Settings

Table B-1: lists the default settings for all Ethernet interface configurations.

Table B-1: Ethernet Defaults

Setting	Default
Full-duplex or half-duplex mode	Half (half-duplex mode)
Redirecting interface 11 to an AUI port	Front

Fast Ethernet Interface Default Settings

Table B-1: lists the default settings for all Fast Ethernet interface configurations.

Table B-1: Fast Ethernet Defaults

Setting	Default
Auto-negotiation (enabled or disabled)	Enabled (supported only on 100BaseTX MMI PHY cards)
Setting advertised connection capabilities	100BaseTX/full-duplex mode enabled 100BaseTX/half-duplex mode enabled 10BaseT/full-duplex mode enabled 10BaseT/half-duplex mode enabled
Full-duplex or half-duplex	Half (half-duplex mode)

FDDI Interface Default Settings

Table B-1: lists the default settings for all FDDI interface configurations.

Table B-1: FDDI Defaults

Setting	Default
Station type (DAS or SAC)	DAS
Link Error Rate Cutoff	8 (0.00000001 bits per second)
Link Error Rate Alarm	8 (0.00000001 bits per second)
Status Report Policy (SMT notification)	On
Maximum Token Rotation Time (t-max-lower-bound)	167 milliseconds Note: The software rounds the specified value to the nearest multiple of 5.24288 milliseconds. For example, the default value of 167 milliseconds is interpreted by the software as 167.77216 milliseconds.
Requested Token Rotation Time (t_req)	8 milliseconds Note: The software rounds the specified value to the nearest multiple of 20.48 microseconds. For example, the default value of 8 milliseconds is interpreted by the software as 7.9872 milliseconds.
Valid Transmission Time (tvx-timer)	2.5 milliseconds Note: The software rounds the specified value to the nearest multiple of 20.48 microseconds. For example, the default value of 2.5 milliseconds is interpreted by the software as 2.51904 milliseconds.
Ring purger	Off
Full-duplex (or half-duplex mode)	On (Full-duplex mode)

ATM Interface Default Settings

The following tables list the default settings for all ATM interface configurations. Table B-1: shows the type of ATM logical interfaces (ELAN or Bridge Tunnel) that are automatically established after installation, and the conditions under which they are active. Table B-1: lists the ATM physical interface default settings. Table B-1: lists the default parameter settings for each type of default connection shown in Table B-1:.

ATM Interface Default Settings

Table B-1: Interface Defaults at Startup

Type of Connection Established/Logical Interface	VNswitch Module	Conditions Under Which Connection Is Established
ELAN/13 ELAN/2 ELAN/3	VNswitch 900EA VNswitch 900FA VNswitch 900XA	The VNswitch module ATM interface is connected to a DIGITAL GIGAswitch/ATM.
FDDI ATM Bridge Tunnel/27 FDDI ATM Bridge Tunnel/16 FDDI ATM Bridge Tunnel/17	VNswitch 900EA VNswitch 900FA VNswitch 900XA	The VNswitch module ATM interface is connected to a DIGITAL GIGAswitch/FDDI AGL2.
Ethernet ATM Bridge Tunnel/28 Ethernet ATM Bridge Tunnel/17 Ethernet ATM Bridge Tunnel/18	VNswitch 900EA VNswitch 900FA VNswitch 900XA	The VNswitch module ATM interface is connected to the ATM interface of another VNswitch 900EA module.

Table B-1: ATM Physical Interface Default Parameters

Setting	Default
Transmission type	SONET
Transmission timing	local (front panel interface) loop (interface to the DEChub ONE backplane)
Payload scrambling	On (enable)
Testing the ModPHY card (loopback)	Off

ATM Interface Default Settings

Table B-1: ATM Logical Interface Default Parameters

Setting	Default
ELAN Name	No name
Identifying the LES address	Automatic (mode)
Expected ARP Response Time	1 second
LE_ARP Request Retries	1
Aging Time (for LE_ARP cache entries)	300 seconds
Forward Delay Time (for LE_ARP cache entries)	15 seconds
Maximum Frame Size	1516 octets
Maximum Unknown Frame Count	1 frame
Maximum Unknown Frame Time	1 second
Control Timeout	120 seconds
VCC Timeout	20 minutes
Flush Timeout	4 seconds
Path Switching Delay	6 seconds

VNbus Default Settings

Table B-1: lists the default settings for the VNbus interface.

Table B-1: VNbus Defaults

Setting	Default
Interface status (enabled or disabled)	Enabled

Bridge Default Settings

Table B-1: lists the default settings for all bridge configurations.

Table B-1: Bridge Defaults

Setting	Default
Rate Limiting (per filter basis)	No (disabled)
Switch-wide Rate Limiting	Disabled
Switch-wide Rate Limiting frames/second	400
Bridge Priority	32768
Port Path Cost	0 (Causes STP to calculate the Path Cost)
Port Priority	128
Bridge-Max-Age	20
Bridge-Hello-Time	2
Bridge-Forward-Delay	15
Bridge port STP state	Enabled
Manual Mode	No (disabled)
Bridge port state	Enabled
IP fragmentation	Enabled (not configurable)
IPX translation	Disabled
No Frame Interval	300 seconds
Aging Time	300 seconds

Bridge Default Settings

Setting	Default
VSDs	<p>All bridge ports on a VNswitch module are members of a default VSD. The default VSD is numbered VSD 1 and is assigned the name "Default." The VNbus tag number for the default VSD is 65. The number, name, and VNbus tag cannot be changed.</p> <p>When more than one VNswitch module is resident in a MultiSwitch 900, all ports on all VNswitch modules are, by default, members of the same (default) VSD. Ports that are members of the default VSD operate as a traditional bridge without VLANs, except that the bridge spans the VNbus.</p>

Maintenance Default Settings

Maintenance Default Settings

Table B-1: lists the default settings for maintenance procedures.

Table B-1: Maintenance Defaults

Setting	Default
Diagnostics downloads (dumps)	Disabled

ELS Configuration Default Settings

Table B-1: lists Event Logging System default configurations for nonvolatile memory (NVRAM). Table B-1: lists Event Logging System default configurations for volatile memory.

Table B-1: ELS NVRAM Configuration Defaults

Setting	Default
Display by subsystem	No subsystems
Trap by subsystem	No subsystems
Display by subsystem and logging level	No subsystems or logging levels
Trap by subsystem and logging level	No subsystems or logging levels
Display by event number	No events
Trap by event number	No events
Display by group	No groups
Trap by group	No groups
Maximum traps per second	0 (Indicates an unlimited number of traps per second is permitted.)
Set output (displaying events directly from the log, not from a CLI prompt)	Default (Events can be viewed by entering Events at the Main prompt.)

Table B-1: ELS Volatile Memory Configuration Defaults

Setting	Default
Display by subsystem	No subsystems
Trap by subsystem	No subsystems
Display by subsystem and logging level	No subsystems or logging levels
Trap by subsystem and logging level	No subsystems or logging levels

ELS Configuration Default Settings

Setting	Default
Display by event number	No events
Trap by event number	No events
Display by group	No groups
Trap by group	No groups
Maximum traps per second	Equal to the current setting for maximum traps per second applied to NVRAM configuration

Remote Management Default Settings

Table B-1: lists the default settings for remote management.

Table B-1: Remote Management Configuration Defaults

Setting	Default
Enabling and disabling access via modem	Disabled
OBM port speed	9600 baud
OBM RTS state (Enabled or Disabled)	Disabled
Router Discovery	Enabled
RIP Listening	Disabled
(Host) Services	Enabled
Address subnet mask (for Host Services)	255 . 0 . 0 . 0, if the address is a Class A IP address 255 . 255 . 00, if the address is a Class B IP address 255 . 255 . 255 . 0, if the address is a Class C IP address
SNMP	Enabled
One SNMP community	Public
Access to SNMP public community	Write-read-trap

Appendix C

VNswitch Packet Counters

Overview

Introduction

This appendix provides an overview of the VNswitch counters and the effect of packets on counters as packets flow through the VNswitch.

In This Appendix

The following topics are covered in this chapter:

Topic	Page
Packet Counter Overview	C-2
Interface Counters	C-6
Bridge Port Counters	C-9
Management Interfaces	C-14

Packet Counter Overview

The VNswitch contains packet counters that allow you to observe the amount and types of traffic being processed. The counters keep track of sent and received traffic, in categories that indicate how many packets have reached various outcomes (terminated, dropped, bridged, routed, flooded, fragmented, and so on).

Packet counters exist at four internal layers to help you trace packets as they flow within the VNswitch:

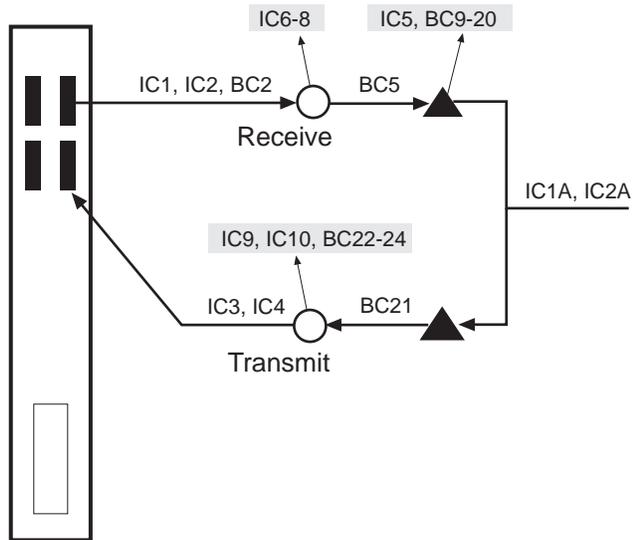
- Logical interfaces
- Bridge ports
- VLAN interfaces
- IP router

When a packet is received by an entity within a layer, the packet is either dropped, processed, or passed on to one or more entities within the next layer. Dropped packets cause an error counter within that layer to increment. In some situations, such as bridge ports and interfaces, entities are tightly coupled and dropped packets can increment error counters in two layers. The effects of the packets are not seen in counters at any layers it does not reach. Packets that are successfully processed at a layer increment non-error counters within that layer. Packets sent to the next layer increment non-error counters in that subsequent layer as well.

[Figure C-1](#) illustrates the logical interface and bridge port layers and shows the relationships between logical interfaces and bridge ports. The figure shows which counters are incremented for the various paths a packet can take within a VNswitch.

Packet Counter Overview

Figure C-1: Packet Flow



- ▲ = Bridge Port
- = Logical Interface
- = Dropped packets

IC n = Interface Counters 1-10

BC n = Bridge Port Counters 1-24

LKG-10694-97WF

Packet Counter Overview

Packets arriving at the VNswitch enter the physical interface and can then travel through each of the four layers. Physical interfaces are the connection jacks for cables, and have a one-to-one or one-to-many (in the case of an ATM physical interface) relationship with interfaces. Logical interfaces, (shown as circles in [Figure C-1](#)), are the lowest layer where counters are used. All packets received and sent are counted by logical interface counters. Error counters for interfaces catch some basic types of errors appropriate to the level of decoding the packet has undergone at this point (for example, a bad FCS) or other errors that are not necessarily associated with a specific higher-level protocol (for example, buffer overflow). If such an error is detected on an interface, the packet being sent or received is discarded, and the appropriate interface error counter is incremented. Otherwise, it is passed to a bridge port (where bridging runs on all interfaces).

Packets arriving at a bridge port (dark triangle in [Figure C-1](#)) are first subject to the effects of bridging. They may be dropped for numerous reasons (destination address filtering, STP port state, and so on), each of which causes a single bridge error or increment a dropped packet counter for that port. If a packet is not dropped, its destination address determines whether it is unicast to another port, flooded out all ports, terminated, and/or delivered to routing. If the packet is bridged out other ports, the bridge attempts to translate and enqueue the packet for sending, if necessary. A failure in this process causes a dropped packet and the error counter increments for the received port. A success means that the packet is sent out other ports and counted by them as well. If a received packet is not dropped or sent out by bridging, it is terminated (such as an STP BPDU) and/or submitted to routing.

VLAN interfaces (VIs) receive all packets destined to routing. VIs, which are groups of bridge ports, are paired one-to-one with VLANs. VIs submit packets for routing on behalf of any ports within their VLAN. VI counters keep track of the total number of packets submitted to routing from their VLAN. Outbound packets sent by routing also go through a VI for transmission on a VLAN. VI transmit counters increment once for each packet sent by routing, although multiple packets may be sent on one or more ports (whose counters are incremented as well). Packets sent or received on VIs cannot be dropped by the VI. All errors, overflows, and so on, at the router layer are detected and counted in other layers. Packets reaching the router may be terminated and are counted by routing. The VNswitch IP counters count transmitted, received, and error packets across all VIs and do not display this information on a per-VI basis.

Counter Relationships

Some simple relationships exist between interface counters, bridge port counters, and VI counters. For a given packet in [Figure C-1](#), the following relationships exist. Refer to the [Interface Counters](#) and [Bridge Port Counters](#) sections for a complete list of the counters.

Receive Relationships

Example 1

$$IC1 + IC2 = BC2$$

(Unicast packets received + Multicast packets received = Total frames received by the interface)

Example 2

$$BC2 \geq BC5$$

(Total frames received by interface \geq Frames submitted to bridging.)

Some packets received on an interface may be dropped before being submitted to bridging.

Example 3

$$BC5 \geq BC6$$

(Frames submitted to bridging \geq Frames submitted to routing)

Only some packets submitted to bridging will be submitted to routing. The rest are either dropped or bridged.

Example 4

$$BC6 \leq IC1A + IC2A$$

(Frames submitted to routing \leq Unicast packets received + Multicast packets received)

The packets submitted to routing by a port's VI may represent only a portion received by that VI since the VI's VLAN may contain other ports.

Transmit Relationships

Example 5

$$IC3 + IC4 \leq BC20$$

Unicast packets transmitted + Multicast packets transmitted \leq Translation failures

Some packets sent by bridging may be dropped at the interface layer.

Interface Counters

Table C-1 describes each interface counter (IC) associated with the VNswitch. Interface counters IC1A, IC2A, IC3A, and IC4A are identical to IC1, IC2, IC3, and IC4, except the A counters represent VI counters.

Table C-1: Interface Counters and Descriptions

Counter Number	Name and Definition
IC1,IC1A*	Unicast packets received Unicast packets received on an interface.
IC2,IC2A*	Multicast packets received Multicast packets received on an interface.
IC3,IC3A*	Unicast packets transmitted Unicast packets transmitted on an interface.
IC4,IC4A*	Multicast packets transmitted Multicast packets transmitted on an interface.
IC5	Input overflow drops Input queue overflows on an interface (received packet is dropped). <u>Note:</u> Some of these packets may be counted in bridge port counters BC14 and BC15.
IC6	Input error drops Invalid packets/errored packets received on an interface (packet is dropped). Examples of this are packets which are of an invalid length (are shorter than the length given within the packet or are too short to contain the required packet fields).
IC7	Input unknown protocol drops Packets received on an interface which are destined to an unknown protocol (packet is dropped).
IC8	Input congestion control drops Packets received on an interface which were flow controlled (dropped) due to congestion on the packet's receive and transmit interfaces. <u>Note:</u> Some of these packets may be counted in bridge port counter BC22 as well.
IC9	Output overflow drops Outbound packets dropped on an interface due to an output queue overflow. <u>Note:</u> Some of these packets may be counted in bridge port counter BC22 as well.

Interface Counters

Counter Number	Name and Definition
IC10	<p>Output error drops</p> <p>Outbound packets dropped on an interface because of one of the following:</p> <ul style="list-style-type: none">• The interface is down or going down while the packet is transmitted• The data link hardware chip aborts sending the packet• The maximum bridge latency of the packet for the associated bridge port has expired.• The packet was too large to be transmitted on the interface. <p><u>Note:</u> Some of these packets may be counted in bridge port counters BC23 and BC24 as well.</p>

*. The *a* suffix on a counter number refers to VI counters.

Bridge Port Counters

Table C-2 lists and describes the bridge port counters.

For most of the dropped packet counters, note that the counter is incremented only if the packet dropped completely—meaning that it is not sent out *any* ports.

For example, assume the case that a packet arrives with an unknown DA and requires flooding and translation to other ports. If the packet can be successfully sent out at least one of these ports, counters BC9, BC10, BC11, BC12, BC13, BC14, BC15, and BC16 will not be incremented. Counter BC20, however, will be incremented if the packet *fails* to make it out at least one port.

Bridge Port Counters

Table C-2: Bridge Port Counters and Descriptions.

Counter Name	Name and Description
BC1	Port restarts Number of times this bridge port's associated interface has passed a self-test, indicating that the port has restarted.
BC2	Total frames received by interface Number of packets received on this bridge port's associated interface (including those eventually discarded due to errors). This is the same as the sum of unicast and multicast packets received on the ports associated interface.
BC3	IP frames fragmented Number of IP packets received on this port's associated interface that were fragmented due to a smaller frame size on the transmit interface's datalink.
BC4	IP frames not fragmented Number of IP packets received on this port's associated interface that bridging attempted to fragment but could not because the DONT_FRAGMENT bit was set in the IP header. These packets were discarded.
BC5	Frames submitted to bridging Number of packets received on this bridge port. This includes packets that may eventually be bridged, routed, filtered, or dropped.
BC6	Frames submitted to routing Number of IP packets received on this port's associated interface that were delivered to routing.
BC7	Frames with unknown destination address The number of packets received on this bridge port for which the bridge's forwarding database had no destination address.
BC8	Frames causing learning transactions Number of packets received on this bridge port whose source address (SA) is not learned or was known on another port and thus caused a learning transaction.

Bridge Port Counters

Counter Name	Name and Description
BC9	<p>Source address filter drops</p> <p>Number of packets completely dropped (not sent out any port) because of their source address. A packet's source address (SA) can be filtered for two reasons:</p> <ul style="list-style-type: none">• A user-configured source address (SA) filter was defined.• The source address (SA) is a multicast addresses (these are always filtered).
BC10	<p>Destination address filter drops</p> <p>Number of packets completely dropped (not sent out any port) because of their destination address. A packet can be filtered for one of two reasons:</p> <ul style="list-style-type: none">• A user-configured destination address (DA) filter was defined.• The destination address (DA) was received on the source address (SA)'s port (natural bridge filtering).
BC11	<p>Protocol filter drops</p> <p>Number of packets completely dropped (not sent out any port) because of a user-configured protocol filter.</p>
BC12	<p>Address rate limiting drops</p> <p>Number of packets completely dropped (not sent out any port) because of a user-configured limit on the number of packets per second that may be sent to the destination address (DA).</p>
BC13	<p>Protocol rate limiting drops</p> <p>Number of packets completely dropped (not sent out any port) because of a user-configured limit on the number of packets per second that may be sent with this protocol type.</p>
BC14	<p>Input buffer overflows</p> <p>Number of packets completely dropped (not sent out any port) due to a lack of input buffers.</p> <p><u>Note:</u> Some of these packets may be counted in interface counter IC5 as well.</p>
BC15	<p>Input queue overflows</p> <p>Number of packets completely dropped (not sent out any port) due to congestive input loss.</p> <p><u>Note:</u> Some of these packets may be counted in interface counter IC5 as well.</p>

Bridge Port Counters

Counter Name	Name and Description
BC16	Source or destination port blocked drops Number of packets completely dropped (not sent out any port) due to either the source or destination port not being in the forwarding state.
BC17	Terminating queue overflows Number of terminating packets dropped because of an overflow in the termination queue.
BC18	Fragmentation queue overflows Number of packets received with a known unicast destination address (DA) which were dropped because they required fragmentation before being sent out the transmit port, and the fragmentation queue overflowed.
BC19	Translate flood queue overflows Number of packets received that needed to be flooded to multiple ports, some of which required translation, but were not successfully sent out all ports due to an overflow in the translate flood queue.
BC20	Translation failures Number of packets completely dropped (not sent out any port) because the packet could not be translated to a different data link. For example, the length of the packet was less than the minimum required for the received data link.
BC21	Frames sent by bridging Number of packets transmitted on this bridge port. This includes packets that were bridged from other ports as well as locally originated (including routed) packets. It includes packets that may eventually be dropped at the bridge port's associated interface, but not packets that were dropped by the bridge port itself.
BC22	Transmit queue overflows Number of packets that were dropped and not sent on this bridge port due to an overflow in the transmit queue.
BC23	Transmit errors Number of packets that were dropped and not sent on this bridge port due to an error in transmission. <u>Note:</u> Some of these packets may be counted in interface counter IC10.
BC24	Too big to send on port drops

Counter Name	Name and Description
	<p>Number of packets that were dropped and not sent on this bridge port because they were too large. The packet was not fragmented due to one of the following:</p> <ul style="list-style-type: none"> • Fragmentation is disabled. • The packet contains an unfragmentable protocol. • There were errors within the packet that prevented it from being fragmented. • The “do not fragment” bit was set within the packet. <p><u>Note:</u> Some of these packets may be counted in interface counter IC10 as well.</p>

Counter Relationships

Some simple relationships exist between interface counters, bridge port counters, and VI counters. For a given packet in [Figure C-1](#), the following relationships exist.

Receive Relationships

$$IC1 + IC2 = BC2$$

Sum of the unicast and multicast packets received on an interface total all of them.

$$BC2 \geq BC5$$

Some packets received on an interface may be dropped before being submitted to bridging.

$$BC5 \geq BC6$$

Only some packets submitted to bridging will be submitted to routing. The rest are either dropped or bridged.

$$BC6 \leq IC1a + IC2a$$

The packets submitted to routing by a port’s VI may represent only a portion received by that VI since the VI’s VLAN may contain other ports.

Transmit Relationships

$$IC3 + IC4 \leq BC20$$

Some packets sent by bridging may be dropped at the interface layer.

Management Interfaces

The counters supported on the VNswitch can be viewed through various screens on the CLI. These CLI displays are shown below with the counters (IC* and BC*) labeled.

CLI

Example 1

```

Monitor>statistics
Ifc Name                Unicast          Multicast         Packets
                        Pkts Rcv        Pkts Rcv         Trans
0   VNbus/0              2                0                 362
1   Eth/1                2                0                 2
2   Eth/2                2                0                 2
3   Eth/3                2                0                 2
4   Eth/4                4                12                6
5   Eth/5                2                0                 2
6   Eth/6                3                0                 3
7   Eth/7                (IC1) 3         (IC2) 0         (IC3)+(IC4) 3
8   Eth/8                3                0                 3
9   Eth/9                3                0                 3
10  Eth/10               3                0                 3
11  Eth/11               3                0                 3
12  Eth/12               3                0                 3
13  E100/13              3                0                 3
14  E100/14              3                0                 3
15  VLAN/15              0                0                 0
16  VLAN/16              0                0                 0
17  VLAN/17              (IC1A) 0        (IC2A) 0        (IC3A) + (IC4A) 0
...
46  VLAN/46              0                0                 0
Monitor>

```

IC* = Interface Counters

Management Interfaces

Example 2

Monitor>error

Nt	Interface	Input Discards	Input Errors	Input Unk Proto	Input Flow Drop	Output Discards	Output Errors
0	VNbus/0	0	0	0	0	0	0
1	Eth/1	1	0	0	0	0	0
2	Eth/2	1	0	0	0	0	0
3	Eth/3	1	0	0	0	0	0
4	Eth/4	8	0	0	0	0	0
5	Eth/5	1	0	0	0	0	0
6	Eth/6	2	0	0	0	0	0
7	Eth/7	2	0	0	0	0	0
8	Eth/8	(IC5)2	(IC6)0	(IC7)0	(IC8)0	(IC9)0	(IC10)0
9	Eth/9	2	0	0	0	0	0
10	Eth/10	2	0	0	0	0	0
11	Eth/11	2	0	0	0	0	0
12	Eth/12	2	0	0	0	0	0
13	E100/13	0	0	0	0	0	0
14	E100/14	0	0	0	0	0	0
15	VLAN/15						

...VLAN's post no errors (instead refer to the underlying interfaces).

46 VLAN/46

Monitor>

IC* = Interface Counters

Management Interfaces

Example 3

```
Monitor>interface statistics 0 (VNBus)

Monitor>int stat 0

```

Ifc	Ifc' Name	Self-Test Passed	Self-Test Failed	Maintenance Failed
0	0 VNbus/0	2	1	0

```

VNbus-Encapsulation MAC/data-link on VNbus interface

VNbus not connected to hub backplane

-----
Total Receive Frames:                2
Terminated Rcv Mcast Frames:         0 (IC2)
Terminated Rcv Ucast Frames:         2 (IC1)
-----
Total Transmit Frames:               1737
Terminated Xmt Mcast Frames:         1735 (IC4)
Terminated Xmt Ucast Frames:         2 (IC3)
-----
Clock Failure Errors:                 0
Receive Start Cell Errors:            0
Receive Cell Parity Errors:           0
-----
Transmit Interface Errors:            0
Transmit Cell Parity Errors:          0
DME Abort Cell Errors:                0
-----
Monitor>

```

IC* = Interface Counters

Management Interfaces

Example 4

Monitor>interface statistics 1 (ETH)

Ifc	Ifc' Name	Self-Test Passed	Self-Test Failed	Maintenance Failed
1	1 Eth/1	0	7234	0

Ethernet/IEEE 802.3 MAC/data-link on Ethernet/802.3 interface

Physical address	0000F860DE02	Network Speed	10 Mb/s
PROM address	0000F860DE02	Duplex mode	HALF
Shared hardware errors	0		

Input statistics:

failed, frame too long	0	failed, FCS error	0
failed, alignment error	0	failed, FIFO verrun	0
runts	0	packets missed	0
internal hardware errors	0		

Output statistics:

deferred transmission	0	single collisions	0
multiple collisions	0	total collisions	0
excess collisions	0	failed, FIFO underrun	0
failed, carrier sense err	0	SQE test err	0
late collision	0	internal hardware errors	0

MAU Status:

connector	RJ45 s
type	10BASE-T HALF DUPLEX
state	OPERATIONAL
media available state	NOT AVAILABLE
media avail state exits	0
jabber state	NO JABBER
jabber state enters	0
false carriers	0

Monitor>

Example 5

```
BRIDGE>list count port 2
Counters for port 2, interface Eth/1:
Port restarts: 0 (BC1)
Total frames received by interface: 2 (BC2)
IP frames fragmented: 0 (BC3)
IP frames not fragmented: 0 (BC4)
Frames submitted to bridging: 2 (BC5)
Frames submitted to routing: 0 (BC6)
Frames with unknown dest address: 0 (BC7)
Frames causing learning transactions: 0 (BC8)
Dropped, source address filtering: 0 (BC9)
Dropped, dest address filtering: 0 (BC10)
Dropped, protocol filtering: 0 (BC11)
Dropped, address rate limiting: 0 (BC12)
Dropped, protocol rate limiting: 0 (BC13)
Dropped, input buffer overflow: 0 (BC14)
Dropped, input queue overflow: 0 (BC15)
Dropped, source or dest port blocked: 1 (BC16)
Dropped, terminating queue overflow: 0 (BC17)
Dropped, fragmentation queue overflow: 0 (BC18)
Dropped, translate flood queue overflow: 0 (BC19)
Dropped, translation failure: 0 (BC20)
Frames sent by bridging: 2 (BC21)
Dropped, transmit queue overflow: 0 (BC22)
Dropped, transmit error: 0 (BC23)
Dropped, too big to send on port: 0 (BC24)

BRIDGE>

IC* = Interface Counters
BC* = Bridge Port Counters
```

Index

A

A/M PHY ports 5-19

Accessing

- Boot config prompt 10-2
- bridge configuration prompt 7-2
- bridge monitor prompt 8-2
- Config prompt 2-8
- Main prompt 2-7
- Monitor prompt 2-8
- network interface prompts 5-7
- VSD configuration prompt 9-5

Adding

- a community 12-6 to 12-7
- a community address 12-7 to 12-8
- MIB views 12-8 to 12-9
- users 3-2 to 3-3

Address entries

- permanent 7-5
- static 8-18
- See also* Address types, defined

Address filtering

- by destination address 7-5, 8-18
- by source address 7-5, 8-18

Address resolution parameters, setting 5-61 to 5-62

Address Resolution Protocol

- adding ARP cache entries manually 5-73 to 5-74
- changing ARP cache entries manually 5-73 to 5-74
- configuring 5-73 to 5-78
- defined 5-73
- deleting learned ARP cache entries 5-74
- deleting manually entered ARP cache entries 5-74 to 5-75
- displaying current configuration settings 5-77 to 5-78
- displaying operational statistics 6-43
- listing protocols with registered addresses 6-43
- listing registered interfaces 6-43
- managing how long learned ARP cache entries are retained 5-75 to 5-77
- monitoring 6-43 to 6-45

See also ARP cache entries

Address types, defined 1-14

Address, determining

- entry type 8-8
- if a multicast address 8-8
- last port on which seen 8-8
- port on which learned 8-8
- ports on which allowed 8-8
- rate limiting state 8-8

Administrative users, defined 1-15

Advertised connections, setting for Fast Ethernet interface 5-13 to 5-14

Aging parameters, monitoring 8-3

Alarm value, setting for FDDI logical interface 5-22

ARP cache entries

- adding and changing 5-73 to 5-74
- deleting 5-74 to 5-75, 6-43
- displaying 6-43
- enabling and disabling auto-refresh 5-76 to 5-77
- managing how long learned entries are retained 5-75 to 5-77
- setting the refresh timer 5-75 to 5-76

ARP. *See* Address Resolution Protocol

Assigning

- hostname 4-6
- VNswitch IP end node to a VSD 9-20 to 9-21

ATM bridge tunnel

- configuring 5-51 to 5-54
- defined 5-51
- types of 5-52

ATM logical interfaces

- configuring 5-47 to 5-71
- configuring a LAN emulation client 5-56 to 5-71
- configuring an ATM bridge tunnel 5-51 to 5-54
- connecting to front panel or backplane interface 5-32
- displaying type and state 5-50
- enabling and disabling 5-48 to 5-50
- factory defaults B-10
- types established at startup 5-47 to 5-48

ATM ModPHY card

- testing 5-35
- when not to initiate loopback test 5-35
- ATM Physical Configuration prompt 5-31
- ATM physical interface 6-22
 - configuration commands for DS1 interfaces 5-40
 - configuration commands for DS3 interfaces 5-45
 - configuration commands for OC3 interfaces 5-37
 - configuring DS1 interfaces for T1/E1 links 5-38
 - configuring DS3 interfaces for T3/E3 links 5-41
 - configuring OC3 interfaces 5-32
 - displaying configuration parameters and counters 5-36 to 6-24
 - enabling and disabling payload scrambling 5-34
 - interval defined 6-24
 - monitoring 6-20 to 6-24
 - setting transmission timing 5-33
 - setting transmission type 5-33
 - testing ModPHY card 5-35
- Attachment Unit Interface port
 - configuring for remote management 12-5
 - redirecting Ethernet interface to 5-9
- AUI port. *See* Attachment Unit Interface port
- Automatic image recovery, configuring 10-13 to 10-14
- Auto-negotiation
 - Ethernet interfaces 5-8
 - Fast Ethernet interfaces, enabling and disabling on 5-11 to 5-12
 - restarting 5-13
 - when initiated 5-13
- Auto-prompts
 - defined 1-12
 - example 2-16
 - how to use 2-16 to 2-17

B

- B/S PHY ports 5-19
- Backing up

- configuration database 10-13 to 10-18
- executable software 10-13 to 10-14
- Backplane, redirecting Ethernet interfaces to 5-9
- Backup
 - exiting and canceling 10-18
- Boot config prompt
 - accessing 10-2
 - exiting 10-2
- BootP client software 1-17, 2-5
- BootP server and client 10-13 to 10-14
- BPDUs. *See* Hello messages
- Bridge configuration parameters
 - displaying 7-41 to 7-44
 - factory defaults B-12 to B-13
 - setting 7-3 to 7-38
- Bridge configuration prompt
 - accessing 7-2
 - exiting 7-2
- Bridge failure, detecting 7-24 to 7-25
- Bridge monitor prompt
 - accessing 8-2
 - exiting 8-2
- Bridge operational states, monitoring 8-3
- Bridge port
 - activity counters 8-4 to 8-5
 - defined 1-3 to 1-5
 - displaying the name of the interface associated with a 8-6
 - displaying the number of the interface associated with a 8-6
 - displaying the state of a 8-6
 - enabling and disabling 7-32
 - enabling fast start on 7-27
 - profile, monitoring 8-3
- Bridge port versus interface numbering 1-7
- Bridge priority, setting 7-22
- Bridging Ethernet and FDDI networks 7-33 to 7-38

C

- Cable length
 - configuring for DS1 interfaces 5-38
- Canceling
 - backup 10-18

- output to a console session A-6
- restoration 10-18
- software installation procedure 10-12

Changing IDs and passwords 3-5 to 3-7

Clearing interface counters 6-41

CLI

- command shortcuts 2-14
- defined 1-11
- entering commands 2-14

Clock time, setting 4-11

Command line interface. *See* CLI

Command shortcuts 2-14

Commands, entering 2-14

Community

- adding and deleting 12-6 to 12-7
- adding and deleting an address for 12-7 to 12-8
- assigning views 12-10
- default 12-6
- setting access 12-11

Concepts and terminology 1-2

Configuration Bridge Protocol Data Units. *See* Hello messages

Configuring

- Address Resolution Protocol (ARP) 5-73 to 5-78
- automatic image recovery 10-13 to 10-14
- dump files 10-25 to 10-33
- Event Logging System 11-8 to 11-40
- installation file locations 10-8 to 10-9
- spanning tree protocol 7-21 to 7-27
- transparent bridge 7-1 to 7-44
- VLANs 9-1 to 9-21
- VNbus to support VLANs 9-18 to 9-19

Configuring a LAN emulation client

- aging the LE_ARP cache 5-62 to 5-63
- assigning an ELAN name 5-57
- displaying current ATM configuration parameters 5-68 to 5-71
- identifying the LES address 5-58 to 5-61
- setting address resolution parameters 5-61 to 5-62
- setting control frame timeout 5-65
- setting data VCC timeout 5-65 to 5-66
- setting flush timeout 5-67
- setting LE_ARP request retries 5-62
- setting maximum frame size 5-66 to 5-67
- setting maximum unknown frame count 5-64
- setting path switching delay 5-68

Configuring filters

- MAC address as a static entry 8-18 to 8-21
- permanent address 7-5 to 7-9
- protocol 7-10 to 7-20

Configuring interfaces

- ATM bridge tunnel 5-51 to 5-54
- ATM interfaces 5-31 to 5-71
- ATM logical interfaces 5-47 to 5-71
- Ethernet logical interfaces 5-8 to 5-10
- Fast Ethernet logical interfaces 5-11 to 5-17
- FDDI logical interfaces 5-18 to 5-30

Configuring remote management 12-1 to 12-33

- in-band IP address and subnet mask 12-21 to 12-22
- MIB-based 12-6 to 12-19
- Out of Band Management IP address 12-3
- Out of Band Management port 12-2 to 12-5
- Out of Band Management port speed 12-3 to 12-4
- Out of Band Management RTS state 12-4
- remote console sessions 12-2 to 12-5
- Serial Line Interface Protocol 12-5
- Simple Network Management Protocol 12-6 to 12-19
- TCP/IP Host Services 12-20

Connecting the switch console 1-16

Connecting to other devices using Telnet 12-33

Console sessions

- advanced management procedures A-1 to A-7
- canceling display of output to A-6
- setting inactivity timer 4-5
- starting and terminating 2-2 to 2-5
- terminating process output A-7
- See also* Local session *and* Remote session

Counters

- packet counter overview C-2

Counters, clearing interface 6-41

Crash counts, monitoring 4-18

Crash log 10-20

Creating

- filters, default protocol 7-18 to 7-19
- filters, permanent address 7-6 to 7-8
- filters, protocol 7-11 to 7-14
- filters, static MAC address 8-19 to 8-21
- LAN segment on VNbus 9-18
- VSDs 9-6 to 9-11

Cutoff value for FDDI logical interface 5-23

D

- DAS 5-18
- DEChub ONE docking station
 - redirecting Ethernet interface to AUI port 5-9
- Default gateway, configuring using TCP/IP Host Services 12-22
- Default protocol filters
 - creating and modifying 7-18 to 7-19
 - defined 7-18
 - deleting 7-20
- Default settings B-1 to B-17
- Deleting
 - a community 12-6 to 12-7
 - a community address 12-7 to 12-8
 - a single permanent address filter 7-8 to 7-9
 - a single protocol filter by protocol type 7-15
 - all permanent address filters 7-9
 - all protocol filters for all frame types 7-17 to 7-18
 - all protocol filters of a particular frame type 7-16 to 7-17
 - ARP cache entries 6-43
 - crash log messages 10-23 to 10-24
 - default protocol filters 7-20
 - installation file locations 10-10 to 10-11
 - MIB views 12-8 to 12-10
 - permanent address filters 7-8 to 7-9
 - protocol filters 7-14 to 7-18
 - restart diagnostic messages 10-23 to 10-24
 - static MAC address filters 8-21
 - users 3-9 to 3-10
 - VSDs 9-15
- Destination address filtering 7-5, 8-18
- Destination device, displaying path to 12-29 to 12-31
- Digital

- clearVISN MultiChassis Manager,
 - establishing LAN segment on VNbus using 9-18
 - online services 10-14
- Disabled bridge ports and STP packets 7-32
- Disabling
 - a bridge port 7-32
 - ATM physical interface payload scrambling 5-34
 - dumps 10-27
 - interfaces 5-4
 - IPX translation 7-38
 - manual mode 7-31
 - multicast rate limiting 7-7 to 7-8, 7-14, 8-21
 - multiple dumps to the same location 10-28
 - RIP listening 12-22 to 12-23
 - router discovery 12-22 to 12-23
 - Simple Network Management Protocol 12-12
 - spanning tree protocol 7-26 to 7-27
 - switchwide rate limiting 7-4
 - TCP/IP Host Services 12-23
 - traps 12-12 to 12-13
 - VNbus port 9-19
- Displaying
 - ARP cache entries 6-43
 - ARP configuration settings 5-77 to 5-78
 - ARP operational statistics 6-43
 - Boot config settings 10-34
 - bridge configuration parameters 7-41 to 7-44
 - bridge interface addresses 12-27
 - crash log messages 10-20 to 10-22
 - dump file locations 10-31 to 10-32
 - dump status 10-30 to 10-31
 - event messages (the event log) 11-41 to 11-43
 - general switch information 4-7 to 4-10
 - installation file locations 10-9 to 10-10
 - output from multiple processes A-5
 - path to destination device 12-29 to 12-31
 - process ID A-3 to A-4
 - process status A-3 to A-4
 - restart diagnostic messages 10-20 to 10-22
 - users 3-4
 - VSD configuration parameters 9-16 to 9-17

- VSD on which TCP/IP Host Services is available 12-31 to 12-32
- Displaying default protocol filters
 - type applied to a bridge port 8-10 to 8-14
- Displaying interface
 - configuration parameters, ATM logical interface 5-68 to 5-71
 - configuration parameters, ATM physical interface 5-36
 - configuration parameters, Ethernet 5-10
 - configuration parameters, Fast Ethernet 5-16 to 5-17
 - configuration parameters, FDDI 5-29 to 5-30
 - counters, ATM physical interface 6-20 to 6-24
 - input and output queues 6-36 to 6-37
 - number and type 5-3, 6-2 to 6-4
 - packet buffer data 6-31 to 6-33
 - packet error statistics 6-34 to 6-35
 - packet statistics 6-38 to 6-39
 - test results and MAC type 6-40
 - type and state, ATM logical interface 5-50
- Displaying protocol filters
 - types applied to a bridge port 8-10 to 8-14
- Displaying remote management
 - list of known routers 12-28 to 12-29
 - OBM port settings 12-5
 - routing table 12-26
 - SNMP configuration parameters 12-14 to 12-17
 - TCP/IP Host Services settings 12-25
- Displaying VSD parameters, warning messages associated with 9-17
- Documentation
 - list of associated documents xxv
- Dropped buffer 6-37
- DS1 interfaces
 - configuration commands 5-40
 - configuring cable length 5-38
 - configuring line type 5-39
 - displaying current settings for 5-39
 - enabling and disabling payload scrambling 5-34
 - testing the ModPHY card 5-35
- DS3 interfaces
 - configuration commands 5-45
 - displaying current settings for 5-43
 - enabling and disabling payload scrambling 5-34
 - enabling and disabling PLCP 5-42
 - setting line attenuation 5-42
 - setting line type 5-42
 - setting transmission power 5-43
 - testing the ModPHY card 5-35
- Dual Attachment Station. *See* DAS
- Dump
 - displaying IP address and subnet mask used during 10-34
 - displaying location and name of boot file 10-34
 - displaying whether enabled or disabled 10-34
 - displaying whether multiple dumps at a single location is enabled 10-34
 - testing a 10-28 to 10-29
- Dump files
 - configuring and managing 10-24 to 10-33
 - creating 10-26
 - displaying location of 10-31 to 10-32
 - modifying location of 10-32 to 10-33
 - specifying file name and server location 10-26 to 10-27
- Dump status, displaying 10-30 to 10-31
- Duplicate MAC addresses 7-45
 - configuring 7-46
 - configuring permanent 7-46
 - configuring static 7-46
 - deleting permanent 7-48
 - deleting static 7-48
 - displaying 7-47
- during 2-5
- Dynamic address
 - defined 1-14
 - length of time retained 7-40

E

- ELAN name, assigning 5-57
- ELS. *See* Event Logging System
- Enabling
 - ATM physical interface payload scrambling

- 5-34
- bridge port 7-32
- dumps 10-27
- ID and password prompting 3-8
- interfaces 5-4 to 5-6
- IPX translation 7-38
- manual mode 7-30
- multicast rate limiting 7-7 to 7-8, 7-14, 8-21
- multiple dumps to the same location 10-28
- RIP listening 12-22 to 12-23
- router discovery 12-22 to 12-23
- Simple Network Management Protocol 12-12
- SMT notification 5-24
- spanning tree protocol 7-27
- switchwide rate limiting 7-4
- TCP/IP Host Services 12-23
- traps 12-12 to 12-13
- VNbus port 9-19
- Error counts, monitoring interface 6-31 to 6-39
- Error logs
 - displaying 10-20
- Establishing LAN segment on VNbus
 - using clearVISN MultiChassis Manager 9-18
 - using DEChub 900 Auto-LAN feature 9-18
- Ethernet logical interfaces
 - and auto-negotiation 5-8
 - configuring 5-8 to 5-10
 - displaying current configuration parameters 5-10
 - factory defaults B-5
 - monitoring 6-5 to 6-9
 - redirecting to backplane 5-9
 - redirecting to DEChub ONE AUI port 5-9
 - setting full-duplex and half-duplex mode 5-8 to 5-9
- Event Logging System
 - defined 11-1
 - displaying event messages (the event log) 11-41 to 11-43
 - logging levels and event types 11-5 to 11-6
 - preconfigured logging criteria (groups) 11-7
 - printing event messages 11-44
 - selecting which events are logged 11-8 to 11-40

- types of events logged 11-2
- Event message
 - elements of 11-2 to 11-5
 - event numbers 11-5
 - message text 11-5
 - printing 11-44
 - subsystems 11-3 to 11-4
 - types 11-2
- Event number, defined 11-5
- Event types, defined 11-5 to 11-6
- Exiting
 - backup or restoration 10-18
 - prompts 2-10

F

- Factory defaults
 - ATM interfaces at startup B-8 to B-9
 - ATM logical interfaces B-10
 - ATM physical interfaces B-9
 - bridge B-12 to B-13
 - common to all interface types B-4
 - Ethernet logical interfaces B-5
 - Event Logging System B-16
 - FDDI logical interfaces B-7
 - maintenance procedures B-14
 - remote management B-17
 - resetting FDDI 5-28
 - resetting to 4-2 to 4-4
 - switchwide B-3
 - VNbus B-11
- Fast Ethernet logical interfaces
 - configuring 5-11 to 5-17
 - displaying configuration parameters 5-16 to 5-17
 - monitoring 6-10 to 6-14
 - restarting auto-negotiation 5-13
 - setting advertised connections 5-13 to 5-14
 - setting full-duplex and half-duplex mode manually 5-14 to 5-15
 - setting line speed 5-15 to 5-16
- Fast start 7-27
 - disabling 7-28
 - displaying the status of 7-28
 - enabling 7-28

FDDI link error rate, format for specifying 5-22 to 5-23

FDDI logical interfaces

- configuring 5-18 to 5-30

- disconnecting nodes causing excessive link errors 5-23 to 5-24

- displaying configuration parameters 5-29 to 5-30

- enabling and disabling SMT notification 5-24

- factory defaults B-7

- monitoring 6-15 to 6-19

- purging of bad frames 5-26 to 5-27

- setting full-duplex and half-duplex mode 5-27

- setting link error rate alarm 5-22 to 5-23

- setting maximum token rotation time 5-25

- setting requested token rotation time 5-25

- setting the station type 5-22

- setting valid transmission timer 5-26

FDDI station types supported 5-18

Flow control 6-37

Flush timeout, setting 5-67

Forward delay, setting 7-25 to 7-26

Full-duplex mode

- setting Ethernet interface to 5-8 to 5-9

- setting Fast Ethernet interface to 5-14 to 5-15

- setting FDDI interface to 5-27

G

Groups, defined 11-7

H

Half-duplex mode

- setting Ethernet interface to 5-8 to 5-9

- setting Fast Ethernet interface to 5-14 to 5-15

- setting FDDI interface to 5-27

Hello messages

- adjusting frequency of 7-24

- defined 7-23

Hello time, setting 7-25 to 7-26

Help 2-21

Hexadecimal values for common protocols 7-13 to 7-14

Hostname, assigning 4-6

I

ICMP Counters 6-46

ICMP Echo Requests 12-27

Image

- configuring automatic recovery 10-14

- location of source files 10-13

Inactive sessions, setting maximum time for 4-5

In-band IP address, configuring 12-21 to 12-22

Input and output queues, displaying interface 6-36 to 6-37

Installation file locations

- configuring 10-8 to 10-9

- deleting 10-10 to 10-11

- displaying 10-10

- modifying 10-10 to 10-11

Installation of VNswitch software, canceling 10-12

Installing VNswitch modules 1-16

Integrity of SONET or SDH payload envelope 5-34

Intercept character 2-10

Interface versus bridge port numbering 1-7

Interfaces

- availability at installation 5-1

- bridge addresses, displaying 12-27

- counters, clearing 6-41

- defined 1-3

- name associated with a bridge port 8-6

- nonspecific factory defaults B-4

- number associated with a bridge port 8-6

- prompts, accessing 5-7

- prompts, exiting 5-7

- registered with Address Resolution Protocol 6-43

- testing 6-42

Interval, defined for ATM interface activity counters 6-24

IP address 2-5

- configuring OBM 12-3

- configuring the module's in-band 12-21 to

- 12-22
- IP end node, VNswitch module as an 12-20
- IP protocols, supported 12-20
- IPX translation
 - determining whether to enable or disable 7-34 to 7-36
 - disabling 7-38
 - enabling 7-38
 - rules 7-34

L

- LAN emulation clients, configuring 5-56 to 5-71
- LDM port. *See* Load/Dump/Management port
- LE_ARP cache
 - aging the 5-62 to 5-63
 - setting aging time 5-63
 - setting forward delay time 5-63
- LE_ARP request retries, setting the number of 5-62
- LES address, identifying 5-58 to 5-61
- Line attenuation
 - setting for DS3 interfaces 5-42
- Line speed, setting Fast Ethernet interface 5-15 to 5-16
- Line type
 - configuring for DS1 interfaces 5-39
 - setting for DS3 interfaces 5-42
- Link errors
 - disconnecting nodes causing excessive 5-23 to 5-24
 - rate alarm, setting 5-22 to 5-23
 - rate cutoff 5-23 to 5-24
- list command options for OC3 interfaces 6-22
- Load/Dump/Management port 10-25
- Local session
 - advanced management procedures A-1 to A-7
 - canceling output to a A-6
 - defined 1-2, 2-2
 - starting and terminating 2-3 to 2-4
 - terminating process output A-7
- Log files
 - crash 10-20
 - diagnostic 10-20

- displaying 10-20
- error 10-20
- Logging levels, defined 11-5 to 11-6
- Logical interface, defined 1-5
- Login IDs at installation 3-1
- Loop avoidance 7-25
- Loopback test 5-35
 - when not to initiate 5-35
- LOpCon A-2

M

- MAC address
 - configuring duplicate 7-46
 - duplicate 7-45
- MAC address static entry, configuring 8-18 to 8-21
- Main prompt, returning to 2-10
- Maintenance 10-1 to 10-34
 - factory defaults B-14
- Managing TCP/IP Host Services 12-26 to 12-32
- Manual mode
 - displaying state 8-6
 - enabling and disabling 7-30 to 7-31
- MAU
 - installing in DEChub ONE 5-9
 - supported media types 5-9
- Max age, setting 7-25 to 7-26
- Maximum
 - frame size, setting 5-66 to 5-67
 - token rotation time, setting 5-25
 - unknown frame count, setting 5-64
- Medium Attachment Unit. *See* MAU
- Memory
 - checking available RAM 10-19
 - monitoring 4-16 to 4-17
 - types of 1-13
- Message text for ELS, defined 11-5
- Messages
 - crash log 10-20 to 10-24
 - from restart diagnostics 10-20 to 10-24
- MIB
 - configuring for management using 12-6 to 12-19
 - statistics, monitoring ATM physical interface

- 6-24
- views, adding and deleting 12-8 to 12-10
- Mirror 13-2
 - accessing the prompt 13-4
 - changing the destination port 13-6
 - commands 13-10
 - disabling bridging on the destination interface 13-4
 - enabling bridging after mirror use 13-4
 - interaction with SNMP 13-4
 - restrictions 13-3
- Modifying
 - default protocol filters 7-18 to 7-19
 - dump file locations 10-32 to 10-33
 - installation file locations 10-10 to 10-11
 - permanent address filters 7-6 to 7-8
 - protocol filters 7-11 to 7-14
 - static MAC address filter 8-19 to 8-21
 - VSDs 9-12 to 9-14
- Monitor users, defined 1-15
- Monitoring
 - Address Resolution Protocol 6-43 to 6-45
 - ATM interface 6-20 to 6-28
 - ATM physical interface 6-20 to 6-24
 - crash counts 4-18
 - Ethernet interface 6-5 to 6-9
 - Fast Ethernet interface 6-10 to 6-14
 - FDDI interface 6-15 to 6-19
 - general bridging operation 8-3 to 8-4
 - ICMP Counters 6-46
 - interface packet statistics and error counts 6-31 to 6-39
 - network interfaces 6-1 to 6-45
 - restart/reload data 4-18
 - Simple Network Management Protocol defined variables and MIB size 12-14
 - spanning tree protocol 8-15 to 8-17
 - switch memory 4-16 to 4-17
 - TCP/IP Host Services 12-26 to 12-32
 - transparent bridge 8-1 to 8-21
- MOPCon A-2
- Multicast rate limiting 7-7 to 7-8
 - determining whether enabled for a MAC address 8-8

- enabling and disabling 7-7 to 7-8, 7-14, 8-20 to 8-21
- MultiChassis Manager
 - creating ATM link over the backplane using 5-32
 - creating Ethernet LANs over the backplane using 5-9
- Multiple processes, displaying output from A-5

N

- NetWare operating system, setting Ethernet frame type 7-37
- Network
 - frame count, and Ring Purger 5-26
 - interfaces, defined 1-3
 - interfaces, monitoring 6-1 to 6-45
 - manager. *See* VNswitch network manager
 - topology, detecting changes in 7-23 to 7-26
- No Frame Interval
 - defined 7-39
 - setting 7-39
- No Owner Frames. *See* NOFs
- NOFs 5-26
- Nonvolatile RAM (NVRAM), defined 1-13

O

- OBM port. *See* Out of Band Management port
- OC3 interfaces
 - configuration commands 5-37
 - configuring 5-32
 - displaying current settings for 5-36
 - enabling and disabling payload scrambling 5-34
 - list command options 6-22
 - setting the transmission type 5-33
 - testing the ModPHY card 5-35
- Online services 10-14
- Operations users, defined 1-15
- Out of Band Management port
 - configuring 12-2 to 12-5
 - configuring OBM port speed 12-3 to 12-4
 - configuring OBM RTS state 12-4
 - configuring OBM-IP address 12-3
 - displaying port settings 12-5

P

Packet

- buffer data, displaying 6-31 to 6-33
- error statistics, displaying 6-34 to 6-35
- statistics, displaying 6-38 to 6-39
- statistics, monitoring 6-31 to 6-39

Packet counters

- bridge port counters C-9
- interface counters C-6
- receive relationships C-5
- transmit relationships C-5

Passwords

- at installation 3-1
- changing 3-5 to 3-7
- enabling and disabling prompting for 3-8

Path Cost 7-22

Path switching delay, setting 5-68

Payload scrambling, enabling and disabling on ATM interface 5-34

Permanent address entries. *See* Address entries

Permanent address filters

- and destination address filtering 7-5
- and source address filtering 7-5
- configuring 7-5 to 7-9
- creating and modifying 7-6 to 7-8
- deleting 7-8 to 7-9
- forwarding using only manually created filters 7-30 to 7-31

Permanent address, defined 1-14

PHY ports

- identifying 5-19 to 5-21
- relation to station type 5-19 to 5-20

Physical interface

- defined 1-3

PID A-2

PLCP

- enabling and disabling 5-42

Plug and play

- defined 1-2
- factory defaults B-1 to B-17

Plug-and-play 1-17

Port ID, displaying 8-6

Port priority 7-22

Ports

bridge ports defined 1-3 to 1-5

UDP trap port 12-11 to 12-12

See also Out of Band Management port

Preinstallation planning 1-16

Process

- defined A-2
- displaying status of a A-3 to A-4

Process aliases

- defined A-2

Process ID (PID)

- defined A-2
- displaying A-3 to A-4

Prompts

- accessing ATM frontpanel and backplane interface prompts 5-31
- accessing Boot config prompt 10-2
- accessing bridge configuration prompt 7-2
- accessing bridge monitor prompt 8-2
- accessing interface prompts 5-7
- accessing the Config prompt 2-8
- accessing the Main prompt 2-7
- accessing the Monitor prompt 2-8
- accessing VSD configuration prompt 9-5
- exiting 2-10
- exiting interface prompts 5-7

Protocol filtering, supported frame types 7-10

Protocol filters

- configuring 7-10 to 7-20
- creating and modifying 7-11 to 7-14
- deleting 7-14 to 7-18
- deleting a single filter by protocol type 7-15
- deleting all of a particular frame type 7-16 to 7-17
- deleting all of all frame types 7-17 to 7-18
- hexadecimal values for common protocols 7-13 to 7-14

Purging bad frames from FDDI ring 5-26 to 5-27

R

RAM, checking available 10-19

Rate limiting. *See* Switchwide rate limiting, Multicast rate limiting

Redirecting Ethernet interface 11 12-5

Redirecting, Ethernet interfaces to backplane 5-9

- Registered address, defined 1-14
- Reinstalling VNswitch software 10-4 to 10-12
 - configuring installation file locations 10-8 to 10-11
 - location of installation files 10-4
 - using a preconfigured network file location 10-4 to 10-5
 - using an unconfigured network file location 10-5 to 10-7
- Remote management factory defaults B-17
- Remote session
 - advanced management procedures A-1 to A-7
 - canceling output to a A-6
 - configuring for 12-2 to 12-5
 - defined 1-2
 - starting and terminating 2-5
 - terminating process output A-7
- Requested token rotation time 5-24 to 5-25
 - setting FDDI 5-25
- Reserved address, defined 1-14
- Reserving VSDs 9-11
- Resetting
 - factory defaults 4-2 to 4-4
 - FDDI interface defaults 5-28
- Responsibilities of VNswitch network manager 1-16
- Restart/Reload data, monitoring 4-18
- Restarting auto-negotiation 5-13
- Restarting the switch 10-3
 - and clearing of bridge tables 10-3
 - and dropped packets 10-3
- Restoring
 - configuration database 10-16 to 10-18
 - executable software (image) 10-13 to 10-14
 - exiting and canceling 10-18
- Returning to Main prompt 2-10
- Ring purger 5-26 to 5-27
 - and unexpectedly high network frame count 5-26
- RIP listening, enabling and disabling 12-22 to 12-23
- RMON
 - alarm and event groups 13-11
 - clearing configuration information 13-16

- configuring an RMON trap 13-18
 - configuring an SNMP community 13-18
 - displaying statistics 13-17
- Root bridge, influencing selection of 7-22
- Root port, influencing selection of 7-22 to 7-23
- ROpCon A-2
- Router discovery, enabling and disabling 12-22 to 12-23
- Routers, displaying a list of 12-28 to 12-29
- Routing table, displaying 12-26
- RTS state, configuring OBM 12-4

S

- SAC 5-18
- SDH 5-32
- Security
 - enabling prompting for ID and password 3-8
 - types of 1-15
- Self-test 6-42
- Serial Line Interface Protocol, configuring 12-5
- Setting
 - aging time 7-40
 - ATM physical interface transmission timing 5-33
 - ATM physical interface transmission type 5-33
 - bridge forward delay 7-25 to 7-26
 - bridge hello time 7-25 to 7-26
 - bridge max age 7-25 to 7-26
 - bridge priority 7-22
 - clock time 4-11
 - Ethernet full-duplex and half-duplex modes 5-8 to 5-9
 - Fast Ethernet advertised connection 5-13 to 5-14
 - Fast Ethernet full-duplex and half-duplex modes 5-14 to 5-15
 - Fast Ethernet line speed 5-15 to 5-16
 - FDDI full-duplex and half-duplex modes 5-27
 - FDDI link error rate alarm 5-22 to 5-23
 - FDDI requested token rotation time 5-25
 - FDDI station type 5-22
 - FDDI valid transmission timer 5-26

- maximum token rotation time 5-25
- no frame interval 7-39
- path cost 7-23
- port priority 7-23
- session inactivity timer 4-5
- switchwide rate limiting 7-3 to 7-4
- the time unused addresses are retained 7-40

Simple Network Management Protocol

- adding and deleting a community 12-6 to 12-7
- adding and deleting a community address 12-7 to 12-8
- adding and deleting community views 12-8 to 12-10
- configuring 12-6 to 12-19
- displaying configuration settings 12-14 to 12-17
- displaying contact person and module location 12-19
- enabling and disabling 12-12
- enabling and disabling traps 12-12 to 12-13
- identifying a contact person 12-18
- identifying physical location of the module 12-17 to 12-18
- monitoring defined variables and MIB size 12-14
- setting community access 12-11
- setting community views 12-10
- specifying the UDP trap port 12-11 to 12-12

Single Attachment Concentrator. *See* SAC

SLIP. *See* Serial Line Interface Protocol

SMT notification, enabling and disabling 5-24

SNMP. *See* Simple Network Management Protocol

Software installation, canceling 10-12

SONET 5-32

Source address filtering 7-5, 8-18

Spanning tree protocol

- adjusting frequency of hello messages 7-24 and VSDs 7-21
- configuring 7-21 to 7-27
- defined 7-21
- detecting a failed bridge or link 7-24 to 7-25
- detecting changes in network topology 7-23 to 7-26
- disabled bridge ports 7-32
- enabling and disabling 7-26 to 7-27
- influencing selection of root bridge 7-22
- influencing selection of root port 7-22 to 7-23
- loop avoidance 7-25
- monitoring 8-15 to 8-17
- state, displaying 8-6
- undetectable network failure 7-39

SRFs 5-24

Starting and terminating console sessions 2-2 to 2-5

Static address entries. *See* Address entries

Static address filters

- creating and modifying 8-19 to 8-21
- deleting 8-21

Static addresses

- and destination address filtering 8-18
- and source address filtering 8-18
- defined 1-14

Station types

- relation to PHY ports 5-19 to 5-20
- supported on VNswitch 900EF modules 5-18

Status Report Frames. *See* SRFs

STP. *See* Spanning tree protocol

Subsystem, defined 11-3 to 11-4

Switch console sessions, defined 1-2, 2-2

Switchwide factory defaults B-3

Switchwide parameters, defined 4-1

Switchwide rate limiting

- defined 7-3
- disabling 7-4
- effect on rate limiting for multicast address filter and protocol filters 7-4
- enabling 7-4
- setting and enabling 7-3 to 7-4
- setting maximum frames per second 7-3

Synchronous Digital Hierarchy. *See* SDH

Synchronous Optical Network. *See* SONET

System security. *See* Security

T

TCP/IP Host Services

- configuring 12-20

- configuring a default gateway 12-22
- displaying current settings 12-25
- enabling and disabling 12-23
- enabling and disabling RIP listening 12-22 to 12-23
- enabling and disabling router discovery 12-22 to 12-23
- logical interface, defined 1-10
- managing 12-26 to 12-32
- monitoring 12-26 to 12-32
- selecting the VSD on which an IP end node resides 12-21
- VSD on which available 12-31 to 12-32
- Telnet, using 12-33
- Terminating
 - console sessions 2-5
 - process output to all console sessions A-7
- Testing
 - a network connection using the ping command 12-27 to 12-28
 - an interface 6-42
 - ATM ModPHY card 5-35
 - dumps 10-28 to 10-29
 - ports inactive for extended periods 7-39
- TFTP. *See* Trivial File Transfer Protocol
- Tmax-lower-bound 5-24 to 5-25
- Transmission power
 - setting for DS3 interfaces 5-43
- Transmission type
 - setting 5-33
- Transparent bridge
 - configuring 7-1 to 7-44
 - defined 1-3
 - monitoring 8-1 to 8-21
- Trap port, specifying 12-11 to 12-12
- T-req 5-24 to 5-25
- Trivial File Transfer Protocol 10-14
- Tvx-timer 5-25 to 5-26

U

- UDP. *See* User Datagram Protocol
- Unknown address, defined 1-14
- Upgrading VNswitch software 10-4 to 10-12
 - configuring installation file locations 10-8 to

- 10-11
- location of installation files 10-4
- using a preconfigured network file location 10-4 to 10-5
- using an unconfigured network file location 10-5 to 10-7
- User Datagram Protocol, specifying the trap port 12-11 to 12-12
- User interface
 - CLI 1-11, 12-1
 - graphical 1-11, 12-1
- Users
 - adding 3-2 to 3-3
 - changing IDs and passwords 3-5 to 3-7
 - deleting 3-9 to 3-10
 - displaying a list of 3-4
- using BootP to automatically obtain for the switch 2-5

V

- Valid transmission time 5-25 to 5-26
 - setting FDDI 5-26
- Views, assigning to a community 12-10
- Virtual LAN. *See* VLAN
- Virtual Network bus. *See* VNbus
- VLAN
 - configuring 9-1 to 9-21
 - defined 1-8
 - logical interface, defined 1-9
 - maximum number per VNswitch module 9-1
- VLAN Secure Domain. *See* VSD
- VNbus
 - configuring to support VLANs 9-18 to 9-19
 - defined 9-3
 - establishing a LAN segment on 9-18
 - factory defaults B-11
 - port, enabling and disabling 9-19
 - tags 9-3
- VNswitch module
 - 900EF, station types supported 5-18
 - accessing over the Web 2-18
 - as an IP end node 12-20
 - assigning a hostname 4-6
 - concepts and terminology 1-2

- configuring a default gateway 12-22
- configuring for remote sessions 12-2 to 12-5
- configuring remote management 12-1 to 12-33
- displaying current TCP/IP settings 12-25
- displaying general information about 4-7 to 4-10
- enabling and disabling RIP listening 12-22 to 12-23
- enabling and disabling router discovery 12-22 to 12-23
- enabling and disabling TCP/IP Host Services 12-23
- factory defaults B-1 to B-17
- IP protocols supported 12-20
- maintenance 10-1 to 10-34
- managing over the Web 2-20
- managing TCP/IP Host Services 12-26 to 12-32
- monitoring TCP/IP Host Services 12-26 to 12-32
- restarting 10-3
- upgrading and reinstalling switch software 10-4 to 10-12
- VLANs per VSD supported 1-8
- VSD on which TCP/IP Host Services is available 12-31 to 12-32

- VNswitch network manager
 - providing access to 3-2 to 3-3
 - responsibilities of 1-16

- VNswitch software components
 - described 1-11
 - prompts 1-11

- VNswitch web server
 - enabling and disabling 2-20

- Volatile RAM, defined 1-13

VSD

- across emulated LANs or bridge tunnels 9-9 to 9-11
- across modules in a single hub 9-8
- and the spanning tree protocol 7-21
- assigning VNswitch IP end node to 9-20 to 9-21, 12-21
- configuration parameters, displaying 9-16 to 9-17

- configuration prompt, accessing 9-5
- configuration prompt, exiting 9-5
- creating 9-6 to 9-11
- default 9-3, B-13
- defined 1-8
- deleting 9-15
- modifying 9-12 to 9-14
- on which TCP/IP Host Services is available 12-31 to 12-32
- reserving 9-11
- when you may want to modify 9-12
- within a single module 9-8

W

- Web server

- enabling and disabling 2-20

- Web-based management application
 - accessing help 2-20