DIGITAL VNswitch 900 Series

Router Management

Part Number: AA-R87CC-TE

June 1998

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

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

Contents

Preface

1 Introduction

2 Configuring VSDs and VLAN Interfaces

3 Configuring and Monitoring the IP Interface

4 Configuring and Monitoring the RIP Interface

5 Configuring and Monitoring the OSPF Interface

6 Configuring and Monitoring the ARP Interface

7 Configuring and Monitoring the BGP4 Interface

A DIGITAL Trace Facility

B Command Line Interface Quick Reference

C VNswitch Counters

D Configuration Examples

Index

Figures

Tables

Preface

Overview

Purpose of This Manual

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

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 routers.

Organization

This manual is organized as follows:

Section	Description
Chapter 1	Provides general information about the VNswitch 900 series of router products and an overview of operational basics that are common to many of the configuration and management tasks described in the manual.
Chapter 2	Describes how to configure a VLAN Inteface (VI).
Chapter 3	Discusses how to configure and monitor your router using IP protocol.
Chapter 4	Discusses how to configure your router using RIP protocol.
Chapter 5	Discusses how to configure and monitor your router using OSPF protocol.
Chapter 6	Discusses how to configure and monitor your router using ARP protocol.
Chapter 7	Discusses how to configure and monitor your router using BGP4 protocol.
Appendix A	Discusses how to configure and monitor your router using the DIGITAL Trace Facility.
Appendix B	Provides a list of all router CLI commands.
Appendix C	Provides an overview of the VNswitch counters and the effect of packets on counters as packets flow through the router.
Appendix D	Provides examples of a VNswitch configuration.

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.
Boldface Italics	Boldface type in italics indicates variables for which the user or the system supplies a value
Boldface **u**nderscore	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 **list users**, and can be entered as **l u**. Similarly, the command for viewing error statistics is **error** and can be entered as **er**.
Return	Indicates that you should press the Return key.
Ctrl/*keystroke*	Indicates you should press the key specified by *keystroke* 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 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 an "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
DIGITAL VNswitch 900 Series Switch Management AA-R2LD-TE*	Describes how to configure, monitor, and manage a VNswitch 900 series module.
DIGITAL VNswitch 900 Series Router Management AA-R87C-TE*	Describes how to use the VNswitch routing software to configure, monitor, and manage VNswitch routing functions.
DIGITAL VNswitch 900EE Installation and Configuration EK-DVNEE-IN	Describes the VNswitch 900EE module, including features, installation, and configuration information.
DIGITAL VNswitch 900EX Installation and Configuration EK-DVNEX-IN	Describes the VNswitch 900EX module, including features, installation, and configuration information.
DIGITAL VNswitch 900EF Installation and Configuration EK-DVNEF-IN	Describes the VNswitch 900EF module, including features, installation, and configuration information.
DIGITAL VNswitch 900EA Installation and Configuration EK-DVNEA-IN	Describes the VNswitch 900EA module, including features, installation, and configuration information.
DIGITAL VNswitch 900LL Installation and Configuration EK-DVNLL-IN	Describes the VNswitch 900LL module, including features, installation, and configuration information.
DIGITAL VNswitch 900XX Installation and Configuration EK-DVNXX-IN	Describes the VNswitch 900XX module, including features, installation, and configuration information.

Title and Order Number	Description
DIGITAL VNswitch 900XA Installation and Configuration EK-DVNXA-IN	Describes the VNswitch 900XA module, including features, installation, and configuration information.
DIGITAL VNswitch 900FA Installation and Configuration EK-DVNFA-IN	Describes the VNswitch 900FA module, including features, installation, and configuration information.
DIGITAL ATM Modular PHY Cards Installation EK-DAGGM-IN	Provides installation and operating guidelines for installing, verifying, and removing modular PHY cards. Describes cabling and LED information.
DIGITAL MultiSwitch 900 Owner's Manual EK-DH2MS-OM	Provides installation, use, security, and troubleshooting information for the DIGITAL MultiSwitch 900.
DEChub ONE Installation EK-DEHU2-IN	Provides installation and operation guidelines for standalone module configuration, including mounting options and cabling.
DEChub ONE-MX Installation EK-DEF1H-IN	Provides installation and operation guidelines for standalone module configuration, including mounting options and cabling.
clearVISN Installation AA-QX86*-TK	Provides pre- and post-installation information, as well as actual installation procedures for each application.
clearVISN Overview 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.
clearVISN RMON Manager User's Guide 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.
clearVISN User's Guide AA-QX88*-TK	Provides information for starting each application, configuring them, and general use information.

Associated Documents

Title and Order Number	Description
OPEN DECconnect Applications Guide EC-G6387-42	Provides information to help plan and install networking systems based on DIGITAL OPEN DECconnect System and networking products.
Event Logging System Messages Guide AA-QL2A*-TE	Describes messages logged by the Event Logging System.
Bridge and Extended LAN Reference 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.
A Primer on FDDI: Fiber Distributed Data Interface	Provides general introductory information about the features, topologies, and components of the FDDI local area network standard.
DEChub Network Modules 900-Series Concentrator Reference EK-CONTR-HR	Provides detailed reference information about the various DIGITAL 900-series concentrators supported by the DIGITAL MultiSwitch 900.

Note: The * symbol represents the revision of the manual.

Correspondence

Documentation Comments

If you have comments or suggestions about this manual, send them to the 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 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 describes the features of the DIGITAL VNswitch 900 series router, and the basics that are common to many of the tasks described in the following chapters.

In This Chapter

The following topics are covered in this chapter.

These topics are common to most of the procedures described throughout this manual and are frequently referenced by those procedures.

Router Features and Protocols

The VNswitch 900 series router is available as either a preloaded factory-installed feature or as a software upgrade for existing level 2 VNswitch modules.

The VNswitch is designed for LAN backbones and high-performance work groups supporting bridging, VLAN, and routing capabilities. When installed in a DIGITAL MultiSwitch 900 (formerly DEChub 900 MultiSwitch), the VNswitch has access to the VNbus, providing high-speed 400 Mb/s interoperability between Ethernet, Fast Ethernet, FDDI, and ATM technologies.

The VNswitch router option provides the following features:

- Supports level 3 IP routing while maintaining unicast bridging performance.
- Provides a web management interface.
- Supports dynamic IP interface management.
- Supports IP routing for a maximum of 32 VLANs.
- Provides routing between ATM emulated LANs.
- Eliminates the need for expensive external routers.
- Enables routing without using the module's external ports when connected to the MultiSwitch 900 VNbus backplane.
- Supports multiple filtering options, including filtering based on network, host address, or application (for example, File Transfer Protocol [FTP]).

The VNswitch router supports the following routing protocols:

  — Routing Information Protocol (RIP)

  — Open Shortest Path First (OSPF)

  — Border Gateway Protocol (BGP)

Understanding Network Interfaces and Ports

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

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 the VNswitch 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 imprinted 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.

Figure 1-1: Physical Interfaces

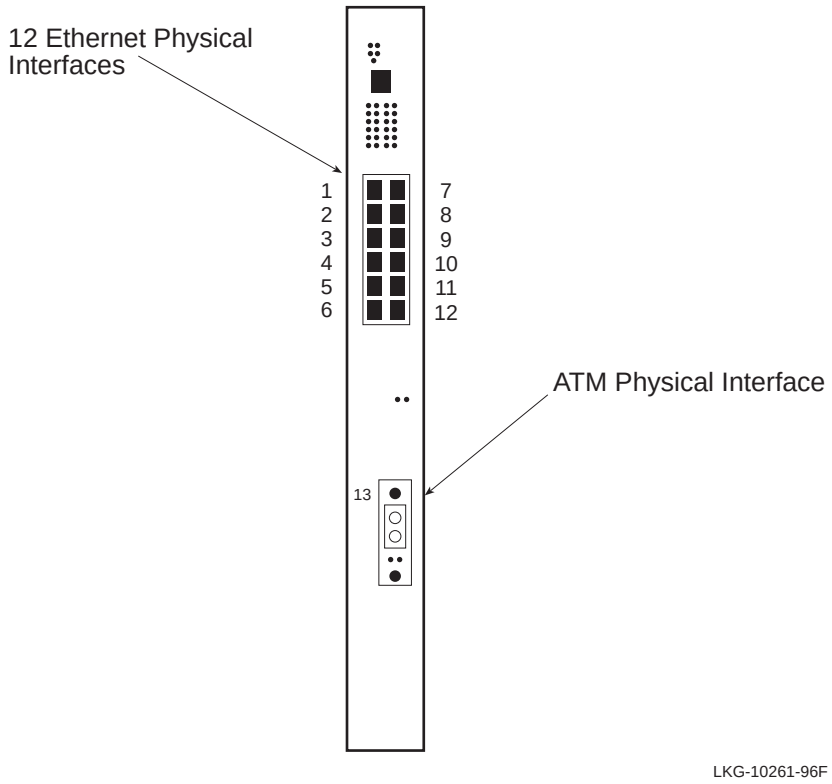

Logical Interface

A logical interface is an abstract connection point 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 to a transparent bridge. The transparent bridge forwards data to, or receives data from, another bridge port, based on the MAC address associated with the data. Each bridge port on a switch is identified by a unique number. The port number has a value that is the same as the logical interface with which it is associated. Figure 1-2 shows a VNswitch 900EA module that includes two ports from 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

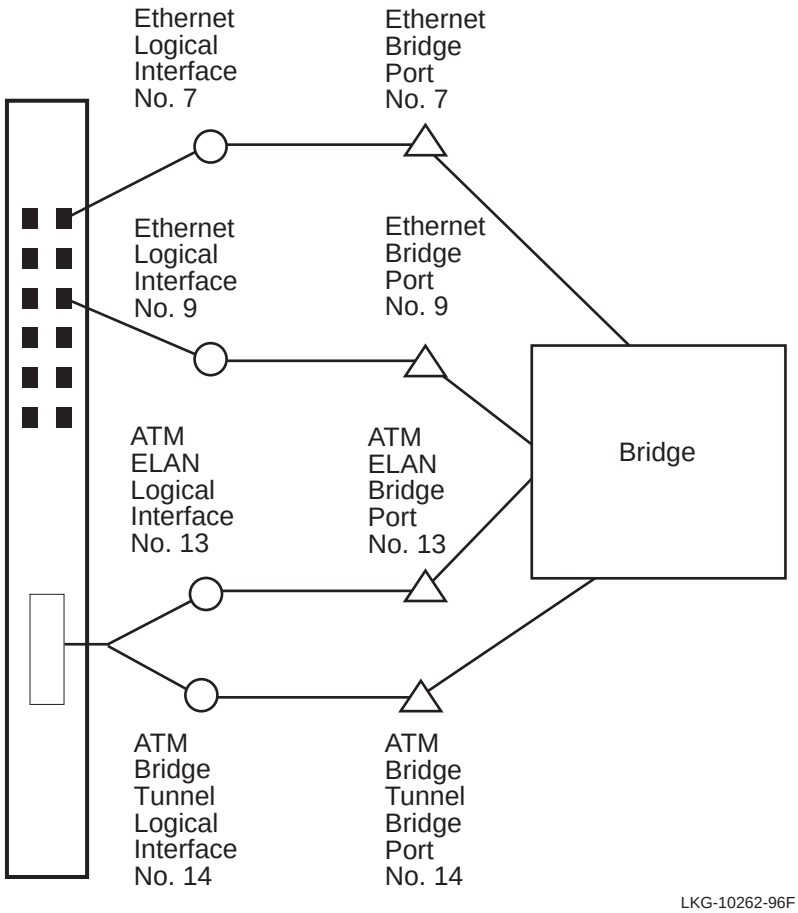

Interface and Bridge Port Numbering Scheme

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.

VLANs and VLAN Secure Domains

A VLAN is a group of bridge ports logically linked to define a LAN on one or more VNswitch modules. 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.

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

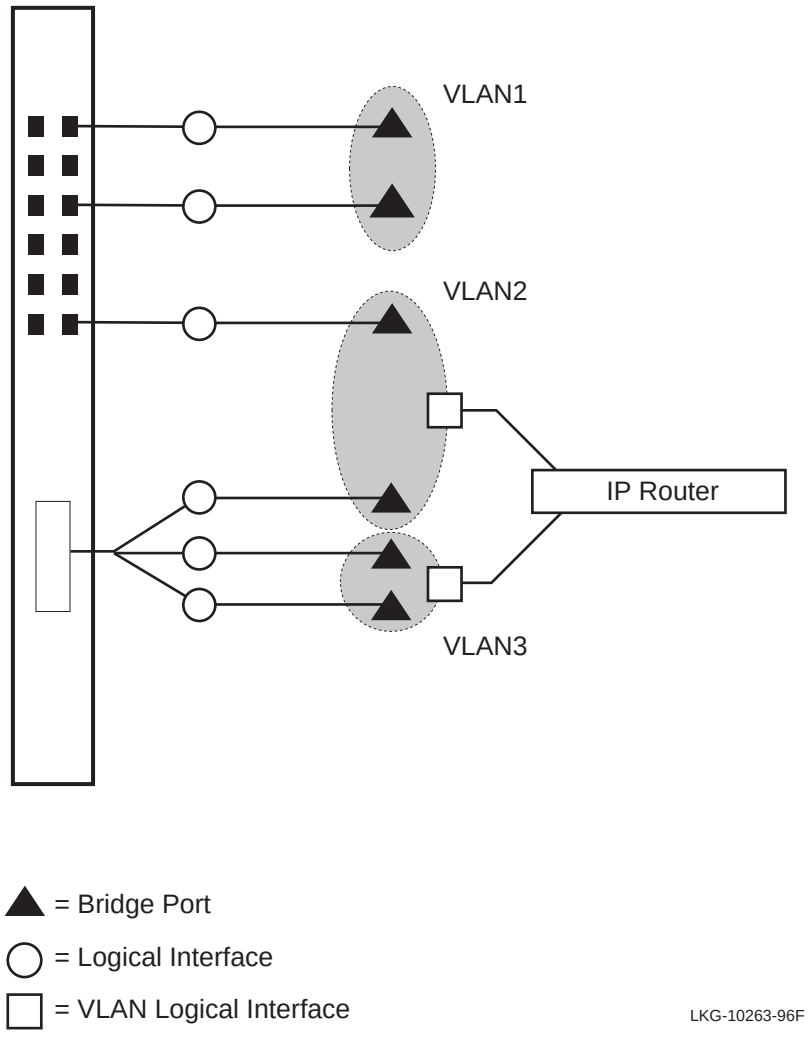

Configuring Routing on VLANs

To configure the VNswitch router to operate on a VLAN, you enable routing globally, then enable routing on a VLAN, and then configure IP to run on a VLAN. The following sections provide an overview on how to configure routing on VLANs.

Enabling Routing Globally

To operate a VNswitch router on a VLAN, you must first enable routing globally. The router is initially configured with routing globally disabled. When you issue the **enable routing** command and answer **yes** to the prompt, the router automatically invokes a restart. Upon completing the **enable routing** command, the router firmware adds an associated 32 routing interfaces called VLAN Interfaces (VIs).

Enabling Routing on a VLAN

Once the router is enabled, 32 VIs are created on the router. You can then select any one of the 32 VIs and associate that VI with a VLAN. Each VI is associated with one VLAN. The default is to associate a newly created VLAN with the first available VI.

IP routing protocols for the VNswitch operate only on VLAN interfaces, but they operate the same as they do on all other LAN interfaces, such as Ethernet, FDDI, or ATM. VLAN interfaces, just like other LAN interfaces, have unique network interface numbers that are automatically generated. A VNswitch 900EE, for example, with 24 Ethernet interfaces assigns the 32 VIs as interfaces 25 through 56.

Configuring IP on a VLAN Interface

To configure IP to run on a VLAN interface, you only need to specify the interface number that corresponds to the VI. You can configure IP on a VI although the VI is not associated with a VLAN; however, IP is not active until you associate the VI with a VLAN.

When specifying the network interface, you supply an IP address, subnet mask, and the IP protocols you want to enable. The IP address and subnet mask you specify for a VI must be valid for the IP subnet connected to the bridge ports in the VI's associated VSD.

Starting and Terminating Console Sessions

The switch console is the terminal from which you configure, monitor, and manage a VNswitch module. The terminal operates as either a local or a remote console. Local consoles access only those modules on the hub to which the console is attached. Remote consoles can access modules located on both the hub to which the console is attached, as well as modules installed in other hubs on the network.

NOTE

If you do not have a BootP server, a local console is required for initial configuration of the switch software, and may be required when upgrading to a new version. A remote console can be used after initial configuration, if either IP routing or TCP/IP Host Services is enabled.

A maximum of two remote consoles can establish a session with a switch at the same time.

Consoles are connected to either a DEChub ONE or a DIGITAL MultiSwitch 900. Whether the session is a local or a remote session depends on the type of port through which you are establishing the connection. Table 1-1 lists the connection methods available, and the type of session (local or remote) you establish through the connection.

Table 1-1: Console Connection Methods

Method of Connection	Type of Session Established
Setup port	Local console session
OBM port	Remote console session
AUI port	Remote console session (DEChub ONE docking station only)
Network server using Telnet	Remote console session

All methods of connection, except the OBM port and the setup port, provide in-band management. The OBM port provides out-of-band management.

The following instructions assume your terminal is already connected and configured for access to VNswitch modules. Refer to the VNswitch installation and configuration documentation for information about how to connect a console to a setup, OBM, or AUI port. Refer to the *DIGITAL VNswitch 900 Series Switch Management* guide for information on how to configure the VNswitch 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, the menu shown in Figure 1-4 is displayed. If your terminal is attached to a DIGITAL MultiSwitch 900, the menu shown in Figure 1-5 is displayed.
2	If your setup port is connected to a DEChub ONE, go to step 4. If your setup port is connected to a 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. The menu shown in Figure 1-4 is displayed.
4	Enter **5** (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 is displayed.

To terminate a console session from a DIGITAL MultiSwitch 900, enter **Ctrl/C** or **logout** at the Main prompt (`Main>`) and press Return. A message is displayed indicating your terminal is disconnected from the module. Press Return to redisplay the DIGITAL MultiSwitch 900 main installation menu (Figure 1-5).

To terminate a console session from a DEChub ONE, enter **logout** at the Main prompt (`Main>`) and press Return.

Figure 1-4: VNswitch 900EA Installation Menu

```
VNswitch 900EA - slot 3
=======================================================
          VNswitch 900EA INSTALLATION MENU


            [1] Reset with Factory Defaults
            [2] Reset with Current Settings
            [3] Show Current Settings
            [4] Configure IP ...
            [5] Go to Local Console

   [Ctrl/C] Return to Chassis Manager Installation Menu
=======================================================

          Enter selection:
```

Figure 1-5: MultiSwitch 900 Installation Menu

```
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:

              Press Return for Main Menu ...
```

Starting and Terminating Remote Sessions

Remote console sessions can be established only after configuring the appropriate network connections. You can start and terminate remote sessions by accessing the OBM port or by BootP, as described in the following section. Also refer to the *DIGITAL VNswitch 900 Series Switch Management* guide for additional information about how to do so.

To start a console session through an OBM port, AUI port, or through a network server, perform the following steps:

Step	Action
1	Access the network operating system prompt from your terminal.
2	Enter **telnet** ***ip-address***, where ***ip-address*** is one of the following addresses. • 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.
3	Press Return. The VNswitch 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 Process prompt is displayed.

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 Remotely Using BootP

If you have a BootP or DHCP server configured 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.

NOTE

The BootP client conforms to RFC 1542, which does not use the DHCP message type in the BootP request. Therefore, DHCP servers that conform to RFC 1534 must be configured to support BootP requests.

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 and the Events Log

The initial steps for most of the tasks discussed throughout this manual involve accessing the command line interface (CLI) prompts (Main, Config, and Monitor) and the events log. 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 Config or Monitor prompts, or the events 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 previous prompt you were using. 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 process is automatically initiated, and the Main prompt (`Main>`) is displayed, each time you start a console session.

Accessing the Config Prompt

To access the Config process, perform the following steps:

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

Accessing the Monitor Prompt

To access the Monitor process, perform the following steps.

Step	Action
1	At the Main prompt (`Main>`), enter **monitor**.
2	Press Return. The Monitor prompt (`Monitor>`) is displayed.
3	If the prompt is not displayed, press Return a second time.

Accessing the Event Log

To access the event log, perform the following steps:

Step	Action
1	At the Main process prompt (`Main>`), enter **events**.
2	Press Return. Output from the Event Log is displayed, if it is configured to do so. (No prompt is displayed.) Refer to the *DIGITAL VNswitch 900 Series Switch Management* guide for more information about configuring and monitoring the Event Logging System.

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 1-20.

Using Command Line Editing

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

Table 1-2: 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 charactor to the left	Ctrl/t

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 1-3 lists the command line recall keys.

Table 1-3: 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.

Table 1-4: 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|`

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 |

CLC ignores spaces entered after the last word of a command, treating the space as a normal space. CLC does not respond to unrecognized entries.

Disabling and Enabling CLC

Command line completion is initially enabled by default and can be disabled. You can change the setting for new sessions dynamically from the `Config>` process, or you can change your current session from the `Monitor>` process. Changing settings using either method affects all processes.

To change the setting for new sessions, use the following commands:

`Config>`**enable command-line-completion**
`Config>`**disable command-line-completion**

To change the setting for current sessions, use the following commands:

`Monitor>`**enable command-line-completion**
`Monitor>`**disable command-line-completion**

To determine the state of CLC for new sessions and the current session, perform the following commands:

`Config>`**list all**
`Monitor>`**list all**

Switching Between Processes

You can switch from one process to another without exiting. For example, you can switch from the `Monitor>` process to the `Config>` process by entering:

`Monitor>` **config**

However, to exit from any process, enter **exit** or **Ctrl/P**.

The Ctrl/P key combination is called the intercept character. You can change the intercept character, if necessary.

Changing the Intercept Character

The intercept character is used to return to the Main process from another process. 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: `Enter character [ ]:`
3	Enter the desired character (**x**, for example).
4	Press Return. The intercept character is changed.

Example `Main>`**intercept x**

Entering Commands and Command Shortcuts

You perform tasks by entering commands at a process prompt. For example, if you want to view a list of all users, you must access the Config process and enter **list users** at the Config process prompt. Similarly, if you want to view error statistics for the network, you access the Monitor process and enter **error** at the Monitor prompt.

Most tasks can also be initiated by entering only part of a command as a shortcut, rather than entering the entire 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 users**, and can be entered at the Config prompt as **l u**. Similarly, the command for viewing error statistics is shown as **error** and can be entered at the Monitor prompt as **err**.

You can obtain help at any of the process (Main, Config, or Monitor) prompts and at any of the lower-level prompts (`IP Config>`, `VSD Config>`, and so on) by typing **?**, followed by Return. Refer to Displaying CLI Help for a detailed description of displaying shortcuts using help.

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 ARP hardware configuration information without going first to the ARP console prompt (`ARP>`). At the Monitor prompt, enter:

`Monitor>` **arp h**

The Monitor prompt remains, but the data displayed is from the ARP 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 subsystems without leaving the current subsystem. For example, while at the IP Config> prompt, you can find out about ARP contents by entering:

```
IP Config> config arp list all
```

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.

Displaying CLI Help

You can obtain help at the Main, Config, or Monitor prompts and at any of the lower-level prompts (`IP Config>, OSPF Config>`, and so on) by typing **?**, followed by 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 a **?** after a specific command name to list its options.

Dynamic Command Configuration

The VNswitch router supports some dynamic configuration commands. However, only a subset of commands for the VNswitch are dynamic. A dynamic command takes effect immediately and does not require you to restart a router after the command is issued.

The VNswitch provides full support of dynamic command configuration for IP and RIP protocols and no support for OSPF protocol.

The following limitation applies to dynamic command configuration:

Once IP is enabled, it cannot be dynamically disabled. However, you can remove all the IP addresses or disable all the VIs.

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 the *DIGITAL VNswitch 900 Series Switch Management* guide for the procedure.

The management application displays with the VNswitch Management window on the right, and an application tree on the left (Figure 1-6). The VNswitch Management window is the first (or top) item in the java application tree, which expands with a menu of parameters that you can use to manage the module. For browsers with java disabled, select non-java to view the application tree.

Figure 1-6: Web Management Display

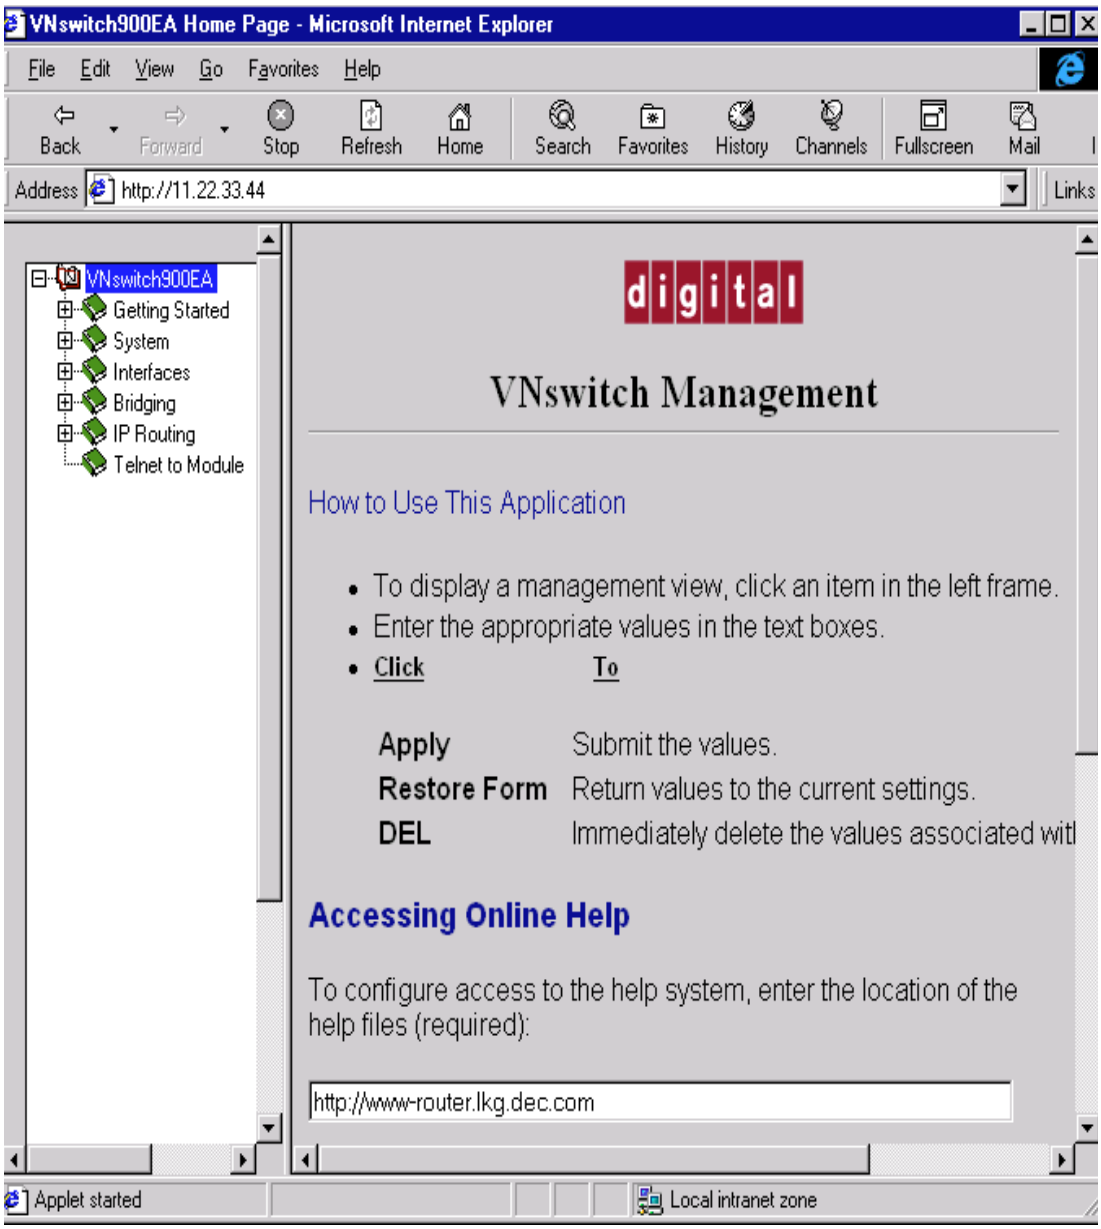

Managing VNswitch Modules Over the Web

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 intalling the help on a web server. For access to the help, you must specify its location in the VNswitch Management window (Figure 1-6).

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
```

Chapter 2

Configuring VSDs and VLAN Interfaces

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.

This chapter focuses on configuring VSDs for routing. For a detailed description of how to create, modify, and delete VSDs, refer to the *DIGITAL VNswitch 900 Series Switch Management* guide.

In This Chapter

This chapter discusses the following topics:

VLAN Secure Domains

A VLAN Secure Domain (VSD) is a logical set of one or more VLANs that operate 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.

VNbus

The Virtual Network bus (VNbus) is a set of three dynamically configured buses on the backplane of the DIGITAL 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 DIGITAL 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 Routing Between 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.

Routing Between VSDs

Routing on a VNswitch is accomplished using VSDs. Since there is currently only one VLAN in each VSD, these terms are identical and can be used interchangeably. The **routing interface** qualifier associated with the **create vsd** command establishes routing on a VI. The routing interface qualifier is available after routing is globally enabled.

Table 2-1 provides a list of VSD commands with the router qualifiers.

Table 2-1: Create VSD Command Options

VSD Command	Parameters	Description
`VSD Config>` **create vsd**	**name** *vsd-name* **ports** *port-list* **tag** *VNbus-tag* **routing-interface** (*any* | *interface-number*)	Enables routing on a VSD as it is created. The *any* keyword associates the next available VSD with the newly created VLAN interface.
`VSD Config>` **modify vsd** *(vsd-number | vsd-name)*	**name** *new-name* **ports** *port-list* **tag** *VNbus-tag* **routing-interface** (*any* | *none* | *interface-number)*	Enables or disables routing on an existing VSD. The *none* keyword disassociates the VSD from any VLAN interface.
`VSD Config>` **delete vsd** *(vsd-number | vsd-name)*		Associated VSD ports return to the default VSD, the VSD name is deleted, and the VNbus tag is available. Disassociates the VSD from the VLAN Interface (VI), but does not delete the VI.
`VSD Config>` **list vsd** *(vsd-number | vsd-name)* **list all**		Displays the specified VSD number, name, ports, VNbus tag, and the network interface number.

Enabling Routing Globally

To operate a VNswitch router on a VLAN, you must first enable routing globally. The router is initially configured with routing globally disabled. Issuing the *enable routing* command and answering **yes** to the prompt, the router automatically invokes a restart. To enable routing globally, perform the following steps:

Step	Action
1	At the Main prompt (`Main>`), enter: **config**
2	Press Return. The Config prompt (`Config>`) is displayed.
3	At the `Config>` prompt, enter: **enable routing**

4	Press Return. The following is displayed and requires action:
	`Enable RIP listening after restart [No]?`
	`Default Gateway [0.0.0.0]?`
	
	`When the box reboots the MAC address assigned to the interface associated with the HST address may be different from the one currently being used. Therefore you may need to flush the ARP cache on your host before you can reconnect via Telnet.`
	
	`***WARNING*** This will invoke an automatic RESTART`
	`Are you sure you want to do this (Yes or No):` **Yes**
	
	`System Restart ...`
	
	After the system is restarted, the VNswitch Installation Menu appears.
	Routing on your VNswitch is now enabled.

Accessing the VSD Config Prompt

To configure a VI for IP routing, you must access the `VSD Config>` prompt with routing enabled. To access the VSD prompt with routing enabled, perform the following steps:

Step	Action
1	At the `Main>` prompt, enter: **config**
2	Press Return. The Config prompt `(Config>)` is displayed.
3	At the `Config>` prompt, enter: **vlans**
4	Press Return. The VSD Config prompt `(VSD Config>)` is displayed.

Creating a VSD

Once you have enabled routing and entered the VSD Config process, you are now ready to create a VSD. The following example creates a VSD named **test**, assigning bridge ports **1-12**, assigning a VNbus tag of ***66*** (VNbus tag 65 is reserved for the default VSD), and assigning the next available VLAN Interface (VI) using the **any** keyword.

Step	Action
1	At the `VSD Config>` prompt, enter: **cr**eate **v**sd

2	Press Return. The following is displayed:
	`VSD Name: [ ]` **test**
	`Bridge Ports (range 1-24): [ ]?` **1-12**
	`VNbus tag (range 66-128): [ ]` **66**
	`Routing over VI (none, any, or one of 27-56): [any]?`
	`VSD 2 created.`

3	At the `VSD Config>` prompt, enter: **e**xit
4	Press Return. The `Config>` prompt is displayed and you have successfully created a VI.

NOTE

The `Routing over VI` *(none, any, or one of 27-56): [any}?* parameter is displayed only when routing is enabled.

Modifying a VSD

You can modify a VSD name, list of assigned ports, or VNbus tag. To modify a VSD name, perform the following steps:

Step	Action
1	To modify a VSD name, at the VSD `Config>` prompt, enter: **modify vsd *number* name *new-name***
2	To modify the ports assigned to a VSD, at the VSD `Config>` prompt, enter: **modify vsd *number* ports *new-port-list***
3	To modify the VNbus tag assigned to a VSD, at the VSD `Config>` prompt, enter: **modify vsd *number* tag *new-vnbus-tag***
4	To modify the VSD using more than one command option (for example, the name, port list, and VNbus tag), enter: **modify vsd *number* name *new-name* ports *new-port-list* tag *new-vnbus-tag***
5	Press Return. The VSD is modified according to the changes you entered.

Deleting a VSD

To delete a VSD, perform the following steps:

Step	Action
1	At the `VSD Config>` prompt, enter: **delete vsd *number*** or **delete vsd *name***
2	Press Return. The VSD is deleted.

Listing VSD Information

To list information about VSDs that are created, perform the following steps:

Step	Action
1	At the `VSD Config>` prompt, enter: **list all**
2	Press Return. All the VSDs and their associated number, name, ports, VNbus tag and IFC are displayed.

Exiting the VSD Config Prompt

You exit the VSDConfig prompt to return to the router configuration prompt. For example, to return to the router configuration prompt (`Config>`) from the `VSD Config>` prompt, enter **exit** and then press Return.

Chapter 3

Configuring and Monitoring the IP Interface

Overview

Introduction

This chapter provides instructions on how to configure and monitor the IP protocol for a VNswitch logical interface. Refer to Appendix D, Configuration Examples, for examples on configuring IP.

In This Chapter

This chapter discusses the following topics:

Enabling IP

To enable the internet protocol (IP), you must start a console session, access the Main and Configuration processes (as described in Chapter 1), enable routing (which initiates a restart), access the IP Configuration process, then assign an IP address.

Once you start a console session, the Main process is automatically initiated, and the Main prompt (`Main>`) is displayed.

Enabling Routing

To access and configure IP, you must first enable routing. To enable routing, perform the following steps:

Step	Action
1	At the Main prompt (`Main>`), enter: **config**
2	Press Return. The Config prompt (`Config>`) is displayed.
3	At the `Config>` prompt, enter: **enable routing**

4	Press Return. The following is displayed and requires action:
	`Enable RIP listening after restart [No]?`
	`Default Gateway 0.0.0.0?`
	
	`After restart, the MAC address assigned to the interface associated with the HST IP address may be different from the one currently being used. Therefore you may need to flush the ARP cache on your host before you can reconnect via Telnet.`
	
	`***WARNING*** This will invoke an automatic RESTART`
	`Are you sure you want to do this (Yes or [No]):` **Yes**
	`System Restart ...`
	
	After the system is restarted, the VNswitch Installation Menu appears.
	Routing on your VNswitch is now enabled.

Accessing the IP Configuration Process

To access the IP Configuration process, perform the following steps:

Step	Action
1	At the `Main>` prompt, enter: **config**
2	Press Return. The `Config>` prompt is displayed.
3	At the `Config>` prompt, enter: **ip**
4	Press Return. The `IP Config>` prompt is displayed.

Configuring Addresses

Adding an IP Address

To enable the IP protocol, you must assign at least one IP address to any of the router's VLAN interfaces (VI). Use the **add address** command to assign an IP address to a VI. To add interface addresses, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **add address** ***interface-number ip-address address-mask***
2	Press Return. The specified IP address is assigned to the specified VIs.

A VI does not receive or transmit IP packets until it has at least one IP address and has an associated VLAN.

You must specify an IP address together with its subnet mask. For example, if the address is on a class B network, using the third byte for subnetting, the mask is 255.255.255.0. Use the list devices command to obtain the appropriate interface-number.

Example: `IP Config>` **add address 31 128.185.123.22 255.255.255.0**

Multiple IP addresses can be added to the same interface as long as the restrictions listed below are met.

Restrictions

The limitations to assigning an IP address to subnets are:

- Only one address is allowed in each subnet.
- Each subnet can appear on only one interface.

Changing an IP Address

To change an IP address, at the `IP Config>` prompt, enter:

change address ***old-ip-address new-ip-address new-subnet-mask***

Deleting an IP Address

To delete an IP address, at the `IP Config>` prompt, enter:

delete address ***ip-address***

Listing an IP Address

To list an IP address, at the `IP Config>` prompt, enter:

list address

This command prints the IP interface addresses that were assigned to the router, along with their configured broadcast formats.

Configuring the Internal IP Address

Setting the Internal IP Address

This command sets the internal IP address that belongs to the router as a whole, and not any particular interface. DIGITAL recommends setting the internal IP address because ping, traceroute and tftp packets sent by the router use the internal IP address. To set the internal IP address, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set internal-ip-address** *ip-address*
2	Press Return.

This address is always reachable regardless of the state of the interfaces. When the internal IP address is not configured, the default router IP address is the router ID. If the router ID is not configured, the router uses the first IP address in the router's configuration.

Example: `IP Config>` **set internal-ip-address 142.82.10.1**

Listing the Internal IP Address

To list the internal IP address, at the `IP Config>` prompt, enter:

list address

Deleting the Internal IP Address

To delete the internal IP address, at the `IP Config>` prompt, enter:

set internal-ip-address 0.0.0.0

Configuring a Router ID

Setting the Router ID Default IP Address

This command sets the default IP address used by the router when sourcing various kinds of IP traffic. To set the router ID, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set router-id** ***ip-address***
2	Press Return.

This address is of particular importance. For example, the source address in pings, traceroute, and TFTP packets sent by the router are set to the router ID. In addition, the OSPF router ID is set to the configured router ID.

The router ID must match one of the configured IP interface addresses of the router. If not, it is ignored. When ignored, or just not configured, the default IP address of the router (and its OSPF router ID) is set to the first IP address in the router's configuration.

NOTE

Configuring a router ID may cause the router's OSPF router ID to change. If this happens, link state advertisements originated by the router before the router ID change persist until they age out (possibly as long as 30 minutes). This may cause an increase in link state database size.

Example: `IP Config>` **set router-id 128.185.120.209**

Listing the Router ID Default IP Address

To list the router ID, at the `IP Config>` prompt, enter:

list address

Deleting the Router ID Default IP Address

To delete the router ID, at the `IP Config>` prompt, enter:

set router-id 0.0.0.0

Configuring a Static Route

Adding a Route

This command adds static network/subnet routes to the router's IP configuration. To add a route, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **add route *ip-network/subnet ip-mask next-hop-address cost***
2	Press Return.

When dynamic routing information is not available for a particular destination, static routes are used. The destination is specified by an IP address (*ip-network/subnet*) with an address mask (ip-mask). For example, if the destination is a subnet of a class B network, and the third byte of the IP address is used as the subnet portion, the address mask is set to 255.255.255.0.

The route to the destination is specified by the IP address of the *next hop* and the *cost* of routing the packet to the destination. The next hop must be on the same subnet as one of the router's directly connected interfaces.

Example: `IP Config>` **add route 17.0.0.0 255.0.0.0 128.185.123.22 6**

Changing a Route

To change a route, at the `IP Config>` prompt, enter:

change route *destination new-mask new-first-hop-address new-cost*

Deleting a Route

To delete a route, at the `IP Config>` prompt, enter:

delete route *destination-mask*

Listing a Route

To list a route, at the `IP Config>` prompt, enter:

list route

This command displays the list of static network/subnet routes that were configured. It also lists any configured default gateways.

Configuring Routing

Setting the Routing Table Size

This command sets the size of the router's IP routing table. The default size is 768 entries. To set the routing table size, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set routing table-size** *number*
2	Press Return.

Setting the routing table size too small causes dynamic routing information to be discarded. Setting the routing table size too large wastes router memory resources.

Example: `IP Config>` **set routing table-size 1000**

Listing the Routing Table Size

To list the routing table size, at the `IP Config>` prompt, enter:

list size

This command displays the routing table size, reassembly buffer size, and the route cache size.

Configuring Access Controls

Adding Access Controls

To add access controls, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **add access-control type *ip-source source-mask ip-dest dest-mask first-protocol last-protocol first-port last-port***
2	Press Return. Before add access control takes effect, you must: • Enter the **set access-control on** command. • Add at least one access control. • Reboot the module. All subsequent add access-control commands will take effect immediately.

This command adds an IP access control entry to the end of the global IP access control list. Each entry must be assigned the following: type, IP source, source mask, IP destination, and destination mask. The type must be either inclusive (I) or exclusive (E). The *ip-source* and *ip-destination* fields are in the form of IP addresses in dotted decimal notation. Optionally, you may specify an IP protocol number range with the *first-protocol* and *last-protocol* fields, which are an inclusive range of IP protocols that match this entry. If a range of protocols was specified which include TCP and UDP protocol numbers, you may specify a TCP and UDP destination port number range with the *first-port* and *last-port* fields, which are an inclusive range of TCP and UDP ports that matches this entry.

Table 3-1 and Appendix D provide examples of setting up IP access controls.

Table 3-1: IP Access Control Examples

Command	Description
IP Config>**add access-control inclusive 0.0.0.0 0.0.0.0 192.67.67.20 255.255.255.255 6 6 25 25**	Allows any host to send packets to the SMTP TCP socket on 192.67.67.20.
IP Config>**add access-control exclusive 150.150.1.0 255.255.255.0 150.150.2.0 255.255.255.0 0 255 0 65535**	Prevents any host on subnet 1 of Class B network 150.150.0.0 from sending packets to hosts on subnet 2 of Class B network 150.150.0.0 (assuming a 1-byte subnet mask).
IP Config>**add access-control inclusive 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 17 17 520 520**	Allows the router to send and receive all RIP packets.
IP Config>**add access-control inclusive 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 89 89**	Allows the router to send and receive all OSPF packets.
IP Config>**add access-control inclusive 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0 255 0 65535**	This is the wildcard inclusive entry which allows the router to send and receive all IP packets.

If IP access control is enabled, you must be careful with packets that the router originates and receives. Be sure not to filter out the RIP or OSPF packets being sent or received by the router. The easiest way to do this is to add a wildcard inclusive entry as the last in the access control list. Alternately, you can add specific entries for RIP or OSPF, or both, perhaps with restrictive addresses and masks. Note that some OSPF packets are sent to the class D multicast addresses 224.0.0.5 and 224.0.0.6, which is important if address checking is being done for routing protocols.

Deleting Access Controls

To delete an access control, at the `IP Config>` prompt, enter:

delete access-control ***record-number***

Moving Access Controls

To move an access control, at the `IP Config>` prompt, enter:

move access-control ***from# to#***

This command places record number *from#* immediately after record number *to#*. After you move the records, they are immediately renumbered to reflect the new order.

Listing Access Controls

To list access controls, at the `IP Config>` prompt, enter:

list access-controls

This command prints the configured access control mode (enabled or disabled) and the list of configured access control records. Each record is listed with its record number. This record number can be used to reorder the list with the IP **move access-control** command.

Enabling Access Controls

The IP access control system allows the IP forwarder to control packet forwarding based on source and destination IP addresses, IP protocol number, and destination port number for the TCP and UDP protocols. This can control access to particular classes of IP addresses and services. To set access control, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set access-control on**
2	Press Return.

The IP access control system is based on one global ordered list of inclusive and exclusive access control entries. If access control is enabled, each IP packet being forwarded, or received is subject to the access control list. Each entry in the list may be inclusive or exclusive, permitting or denying forwarding.

If access controls are enabled and the access control entry list is empty, all packets will be included. If access controls are enabled and the entry list is not empty, all packets not included by an entry in the list will be dropped (excluded).

For each received packet, the headers are compared to all specified fields in each entry in the list in turn. If the entry matches the packet and the entry is inclusive, the packet is forwarded. If the entry is exclusive, the packet is dropped. If no entry matches after going through the entry list the packet is dropped.

Each entry has an IP address mask and result pair for both the source and destination IP addresses. An address is logically "AND-ed" with the mask, and compared to the result. For example, a mask of 255.0.0.0 with a result of 26.0.0.0 matches any address with 26 in the first byte. A mask of 255.255.255.255 with a result 192.67.67.20 matches only the IP host 192.67.67.20. A mask of 0.0.0.0 with a result of 0.0.0.0 is a wildcard, and matches any IP address.

Each entry may also have an optional IP protocol number range. This applies to the protocol byte in the IP header. Any IP packet with a protocol value within the specified range matches. A range of 0 to 255 matches all IP packets. The commonly used protocol numbers are 1 for ICMP, 6 for TCP, 17 for UDP, and 89 for OSPF.

Each entry may also have an optional port number range. This applies only to TCP and UDP packets because the port number is part of the TCP and UDP headers. Any TCP or UDP packet with a destination port number within the specified range matches (TCP and UDP use the same port numbers). A range of 0 to 65535 disables port filtering. Some commonly used port numbers are 21 for FTP, 23 for Telnet, 25 for SMTP, 513 for rlogin, and 520 for RIP.

Configuring Enhanced Proxy ARP

Overview

Enhanced proxy ARP is a method of communicating IP packets between two hosts in different subnets on the same LAN without requiring a router to forward the packets. Enhanced proxy ARP routers use a modified form of RFC 1027 to overcome the restrictions associated with RFC 1027.

Enhanced proxy ARP does not require the *sender IP address* field to be in a directly connected subnet, and the *sender IP address* and *target IP address* fields do not have to be in the same natural network.

Enhanced proxy ARP provides modified RFC 1027 functionality that includes communicating on a LAN, communicating on an extended LAN, and communicating on a VLAN.

Communicating on a LAN

Enhanced proxy ARP allows hosts on the same LAN to communicate directly, regardless of the network or subnet the hosts are assigned. This is accomplished by configuring the hosts to ARP for all destination addresses. In this configuration, hosts can communicate directly with other hosts without a router. Enhanced proxy ARP allows hosts to communicate without a router in subnets that are not in the same natural network.

Communicating on a Routed LAN

Separate LANs connected by a router form an extended LAN. When hosts on different LANs communicate over an extended LAN, then the router connecting them will proxy for the sender host and respond to the ARP request regardless of the destination network.

Enhanced proxy ARP on a routed LAN also provides indirect proxy (proxy-on-behalf) and compliments ICMP redirects with ARP-response redirects.

Indirect Proxy

A router that has ARP routing enabled and proxy-on-behalf enabled (refer to the section titled (Setting Enhanced Proxy ARP on page 3-20) can optionally issue ARP responses on behalf of routers that are not running proxy ARP. When an ARP request is received and the output interface is the same as the input interface and the next hop is a router, then the receiving router issues an ARP response with the next hop router's MAC address in the *sender hardware address* field instead of its own address.

ARP Redirect

When enhanced proxy ARP is enabled, an unsolicited ARP response is sent whenever an Internet Control Message Protocol (ICMP) redirect is sent. ICMP redirect is ignored because hosts check to determine the gateway issuing the redirect is the gateway the host is using for forwarding, but enhanced proxy ARP does not use a gateway and ICMP redirect is ignored.

The unsolicited ARP response maps the destination IP address to the MAC address of the redirected router. This eliminates the problem of ICMP redirects being ignored by hosts with ARP enabled (RFC 1027).

Communicating on a VLAN

Enhanced proxy ARP can be used on VLAN environments to ensure that packets flowing between two hosts in different subnets or networks on different ports of a VNswitch but in the same VLAN will be switched instead of routed.

For example, host A and host B are configured on different ports in different IP subnets of the same VLAN and both have ARP enabled. Host A is connected to host B through an enhanced proxy ARP routing VNswitch. Normally, IP packets between host A and host B are routed by the VNswitch, but with ARP routing enabled, the hosts send packets directly to each other's MAC address using the switch capabilities of the VNswitch, thus improving performance.

Configuring Hosts for Enhanced Proxy ARP

To communicate directly over a LAN environment using enhanced proxy ARP, you must configure hosts to ARP for all destination addresses. The following sections describe how to configure common operating systems for ARP routing.

Windows NT Hosts

To configure a Windows NT host, perform the following steps:

Step	Action
1	Log in as the administrator.
2	From the Control Panel, select the Network icon.
3	Choose TCP/IP network from the list box and click the **Configure** button.
4	Change the gateway address to be your local IP address.
5	Click the **OK** button. The configuration is complete.

Windows 95 Hosts

To configure a Windows 95 host, perform the following steps:

Step	Action
1	From the Control Panel, select the Network icon.
2	Highlight the TCP/IP icon from the list box in the Configuration tab and click the Properties button. The TCP/IP Properties window appears.
3	From the TCP/IP Properties window, click the Gateway tab. Remove any existing gateway addresses in the Installed gatways box.
4	Enter the gateway address to be your local IP address and click the Add button.
5	Click the **OK** button. The configuration is complete.

DIGITAL UNIX and LINUX Hosts

Modify the `/etc/routes` file and add the following entry where *n* is the IP address of the host.

/etc/route add default *n.n.n.n* 0

MIPS ULTRIX Hosts

Use the `/etc/route` program to add a default route into the routing table (insert the line in `/etc/rc.local` to ensure it is executed at startup) where *n* is the IP address.

/etc/route add default *n.n.n.n* 0

Windows 3.1 PATHWORKS Hosts

To configure Windows 3.1 PATHWORKS hosts, perform the following steps:

Step	Action
1	Edit the `cfg*.tpl` file currently being used. To determine the correct cfg file, access the `pathworks` directory and open the `select.ini` file.
2	In the `cfg*.tpl` file, replace "`TCPIPGATEWAY=`" with `0.0.0.0`.
3	Save the file and restart the system.

Windows NT 4.0 DHCP Servers

The DHCP server, NT 4.0 Service Pack 2, supports configuring hosts with their own address as the default gateway. To enable this feature for all clients in a scope, add the following value to the registry, where *n.n.n.n* is the subnet for the scope.

```
HKEY_LOCAL_MACHINE\
 SYSTEM\
  CurrentControlSet\
    Services\
      DHCPServer\
        Configuration\
          Subnets\
            n.n.n.n\
              SwitchedNetworkFlag=1
       (REG_DWORD)
```

NOTE

Enable the DHCP server function before you install the NT Service Pack 2 to ensure that the DHCP server DLL's and images are upgraded from the Service Pack.

Enabling Enhanced Proxy ARP

Enabling enhanced proxy ARP globally enables enhanced proxy ARP on all LAN interfaces running IP that have not been disabled with the **set enhanced-proxy-arp off** command.

To enable enhanced proxy ARP, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable enhanced-proxy-arp**
2	Press Return. Enhanced proxy ARP is enabled.

Disabling ARP Routing

To disable enhanced proxy ARP on all LAN interfaces running IP, at the `IP Config>` prompt enter:

disable enhanced-proxy-arp

Setting Enhanced Proxy ARP

When enhanced proxy ARP is enabled or disabled, it enables or disables all LAN interfaces running IP. You can also selectively set a specific interface to enable or disable enhanced proxy ARP.

To set an interface to disable enhanced proxy ARP, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **set enhanced-proxy-arp off** `Interface number [0]?` **25**
2	Press return. Enhanced proxy ARP is disabled on interface 25.

Setting a specific interface to re-enable enhanced proxy ARP sets enhanced proxy ARP and sets the router's operating parameters, including whether proxy-on-behalf should be enabled on the interface. Normally, proxy-on-behalf is disabled by default.

To set an interface to enable enhanced proxy ARP, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **set enhanced-proxy-arp on** `Interface number [0]?` **25** `Proxy ARP on behalf of other Routers [No]?`
2	Press Return. Enhanced proxy ARP is enabled on interface 25.

Adding Enhanced Proxy ARP Subnets

This command is used to prevent enhanced proxy ARP routing for destinations in certain subnets. For example, a LAN configured with subnets A, B, and C has a router in subnet A that does not have an address in subnet C. When an ARP request is sent by a host in subnet B to a host in subnet C, the router will initiate an ARP request destined to the host in subnet C.

NOTE

Adding enhanced proxy ARP subnets can be accomplished on VLAN interfaces only.

To add enhanced proxy ARP subnets to the router, perform the following steps:

Step	Action
1	At the IP `Config>` prompt, enter: **add enhanced-proxy-arp subnet** `Which net is this subnet for [0]?` **25** `Subnet number [0.0.0.0]?` **16.39.180.0** `Subnet mask [255.0.0.0]?` **255.255.255.0**
2	Press Return. An enhanced proxy ARP subnet is added to interface 25, as an example.

Deleting Enhanced Proxy ARP Subnets

To delete enhanced proxy ARP subnets, at the `IP Config>` prompt enter:

delete enhanced-proxy-arp subnet

Configuring BootP Forwarding

Enabling BootP Forwarding

This command turns on BootP packet forwarding. To use BootP forwarding, you must also add one or more BootP servers with the **add BootP-server** command. To enable BootP forwarding, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable bootp-forwarding** `Maximum number of forwarding hops, ranges 0-16, default [4]?` Enter the maximum number of application hops you want the BootP request to go. This is the maximum number of BootP relay agents that can forward the packet. This is *not* the maximum number of IP hops to the BootP server. A typical value for this parameter is 1.
2	Press Return.
3	`Minimum seconds before forwarding, ranges 0-65535, default [0]?` Enter the number of seconds you want the client to retry before the BootP request is forwarded. A typical value for this parameter is 0.
4	Press Return.
5	`Relay in same VLAN [No]?`
6	Press Return.

Disabling BootP Forwarding

To disable BootP forwarding, at the `IP Config>` prompt, enter:

disable bootp-forwarding

Listing BootP Forwarding

To list BootP forwarding, at the `IP Config>` prompt, enter:

list bootp

This command indicates whether BootP forwarding is enabled or disabled.

Configuring a BootP Server

Adding a BootP Server

BootP is a bootstrap protocol used by a diskless workstation to learn its IP address and the location of its boot file and boot server. This command adds a BootP server to a network configuration. To add a BootP server, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **add bootp-server** ***server-ip-address***
2	Press Return.

Acting as a bootp relay agent, your router accepts and forwards BootP requests to the BootP server.

NOTE

Before the **list all** command can display the BootP server address, you must enable BootP forwarding with the **enable BootP forwarding** command.

Example: `IP Config>` **add bootp-server 128.185.123.22**

Deleting a BootP Server

To delete a BootP server, at the `IP Config>` prompt, enter:

delete bootp-server *server-ip-address*

Listing a BootP Server

To list BootP servers, at the `IP Config>` prompt, enter:

list bootp

This command displays a list of configured BootP servers.

Configuring Broadcast Addresses

Setting Broadcast Addresses

This command specifies the IP broadcast format that the router uses when broadcasting packets out a particular interface. IP broadcasts are most commonly used by the router when sending RIP update packets. To set a broadcast address, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set broadcast-address** *ip-address style fill-pattern*
2	Press Return.

The *style* parameter can take either the value **local-wire** or the value **network**. Local-wire broadcast addresses are either all ones (255.255.255.255) or all zeros (0.0.0.0). Network style broadcasts begin with the network and subnet portion of the IP interface address.

You can set the *fill-pattern* parameter to either 1 or 0. This indicates whether the rest of the broadcast address (other than the network and subnet portions, if any) is set to all ones or all zeros.

When receiving, the router recognizes all forms of the IP broadcast address.

Examples: `IP Config>` **set broadcast-address 192.9.1.11 local-wire 1**

`IP Config>` **set broadcast-address 192.9.1.11 network 0**

Listing Broadcast Addresses

To list broadcast addresses, at the `IP Config>` prompt, enter:

list all

This command displays a list of broadcast addresses.

Configuring a Default Gateway

Setting Default Network Gateway

Routers send packets having unknown destinations (destinations not present in the routing table) toward the default gateway. A default gateway is configured in the router by specifying the next hop to use to get to the default gateway and the cost of sending packets to the default gateway. To set the default network gateway, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set default network-gateway** ***gateway-ip-address cost***
2	Press Return.

Deleting Default Network Gateway

To delete the default network gateway, at the `IP Config>` prompt, enter:

delete default network-gateway

Configuring a Default Subnet Gateway

Setting Default Subnet Gateways

There can be a default subnet gateway configured for each subnetted network that the router knows. When the router attempts to forward a packet to a destination belonging to the subnetted network, but that destination cannot be found in the routing table, the packet is forwarded instead to the default subnet gateway.

Configuring default subnet gateways is the same as configuring the preceding default network gateway. The only difference is that you must specify the subnetted network on the command line.

To set the default subnet gateway, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set default subnet-gateway subnetted-network** *subnet-gateway-ip-address* ***cost*** **Example:** `IP Config>` **set default subnet-gateway 128.0.0.0 128.185.123.22 2**
2	Press Return.

Deleting Default Subnet Gateways

To delete the default subnet gateway, at the `IP Config>` prompt, enter:

delete default subnet-gateway *subnetted-network*

Configuring Directed Broadcast

Enabling Directed Broadcast

This command enables the forwarding of IP packets whose destination is a nonlocal (for example, remote LAN) broadcast address. To enable directed broadcast, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable directed-broadcast**
2	Press Return.

With directed broadcast, the packet is originated by the source host as a unicast where it is then forwarded as a unicast to a destination subnet and exploded into a broadcast. These packets can be used to locate network servers. This command enables both the forwarding and exploding of directed broadcasts. The IP packet forwarder never forwards link-level broadcasts, unless they correspond to class D IP addresses. The default setting for this feature is enabled.

NOTE

Forwarding and exploding cannot be implemented separately. Also, the router does not forward subnet-wide IP broadcasts.

Disabling Directed Broadcast

To disable directed broadcast, at the `IP Config>` prompt, enter:

disable directed-broadcast

Configuring a Filtered Route

Adding a Filter

This command designates an IP network/subnet to be filtered. To add a filter, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **add filter** ***ip-address ip-mask***
2	Press Return.

IP packets are not forwarded to filtered networks/subnets, nor is routing information disseminated concerning such destinations. Packets destined for filtered network/subnets are simply discarded.

You must specify a filtered network/subnet together with its subnet mask. For example, to filter a subnet of a class B network, using the third byte for subnetting, the mask is 255.255.255.0.

Using the filter mechanism is more efficient than IP access controls, although not as flexible. Filters also affect the operation of the IP routing protocols, unlike access controls.

Example: `IP Config>` **add filter 127.0.0.0 255.0.0.0**

Deleting a Filter

To delete a filter, at the `IP Config>` prompt, enter:

delete filter ***ip-address ip-mask***

Listing a Filter

To list a filter, at the `IP Config>` prompt, enter:

list filter

This command displays the networks/subnets that are filtered.

Configuring Path Splitting

Enabling Path Splitting

Path splitting allows packets to be routed to their destination address through any one of four equal-cost paths. When the VNswitch receives a packet with a new destination address, it assigns a path to the packet based on the destination address. All the packets with the same destination address are then forwarded using the assigned path.

If IP access controls are enabled, then the source address, destination address, protocol type and destination part number (for UDP and TCP protocol types) are all used when determining the path. Packets with the same values for all of these fields will use the same path. The default setting for path splitting is set to disable.

To enable path splitting, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable path-splitting**
2	Press Return.

Disabling Path Splitting

To disable path splitting, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **disable path-splitting**
2	Press Return.

Configuring Reassembly Size

Setting Reassembly Size

This command configures the size of the buffers that are used for the reassembly of fragmented IP packets. The default size is 12000. To set the reassembly size, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set reassembly-size** ***number***
2	Press Return.

Example: `IP Config>` **set reassembly-size 12000**

Listing Reassembly Size

To list the reassembly buffer size, at the `IP Config>` prompt, enter:

list size

This command displays the routing table size, and the reassembly buffer size.

Configuring UDP Broadcast Forwarding

User datagram protocol (UDP) broadcast forwarding provides controlled application level functionality that forwards local broadcasts to other specified nodes or networks by using UDP ports. Controlled forwarding is accomplished by forwarding specific broadcasts to specified networks or hosts.

Advantages of UDP Forwarding

Without UDP forwarding, an IP router cannot broadcast outside of the current broadcast domain. With UDP forwarding, you can selectively broadcast to other networks outside of the broadcast domain.

Some client applications search for servers by broadcasting the request. Most of these broadcasts are link layer broadcasts in the same network as the server. A typical example is a NetBIOS Name Server running over TCP/IP services. As the network grows and network components such as bridges and routers are being added, users no longer need to be constrained to local servers. Users can now reach the server outside the broadcast domain using UDP forwarding.

The following sections describe how to add and delete UDP broadcast forwarding and how to disable and enable UDP broadcast forwarding.

Adding a UDP Broadcast Server

Adding a UDP broadcast server automatically enables the udp port being added. To add UDP broadcast forwarding, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **add broadcast-forwarder udp** ***udp-port-number interface-number destination-ip-address***
2	Press Return. UDP broadcast forwarding for the associated port is now added to the list.

Specifying **all** for the interface number includes all interface numbers to the specified UDP port and destination IP address.

Example: `IP Config>` **add br u 25 all 124.24.10.255**

Deleting a UDP Broadcast Server

To delete UDP broadcast server, at the `IP Config>` prompt enter:

delete broadcast-forwarder udp ***udp-port-number interface-number destination-ip-address***

Specifying **all** for the destination IP address deletes broadcast forwarding for all IP addresses with the specified UDP port number and interface number.

Example: `IP Config>` **del br u 25** ***9*** **all**

Specifying *all* for the interface number deletes *only* the add command entry and specified UDP port and destination IP address that used **all** as the interface number. Previous UDP add command entries for that port prior to the delete command entry are not affected.

Example:

In the following example, CLI entry 1 adds broadcast forwarding to UDP port 25, interface 9, with a destination IP address of 124.35.13.255. CLI entry 2 uses a different UDP port and interface number, but the destination address is the same. CLI entry 3 adds broadcasting on all interfaces on UDP port 25 to destination IP address 124.35.13.255. CLI entry 4 deletes only CLI entry 3 parameters. CLI entry 1 parameters are not affected and broadcast forwarding on UDP port 25, interface 9 remains active.

CLI Entry	`IP Config>` Command
1	**add br u 25 9 124.35.13.255**
2	**add br u 26 7 124.35.13.255**
3	**add br u 25 all 124.35.13.255**
4	**del br u 25 all 124.35.13.255**

Disabling UDP Broadcast Forwarding

To disable UDP broadcast forwarding, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **dis**able **br**oadcast-forwarding **u**dp *udp-port interface-number* **Note:** the default interface is **all**.
2	Press Return. UDP broadcast forwarding is now disabled.

Enabling UDP Broadcast Forwarding

To enable UDP broadcast forwarding, at the `IP Config>` prompt enter:

enable **br**oadcast-forwarding **u**dp *udp-port interface-number*

Examples of Configuring UDP Broadcast Forwarding

In the following two examples, IP networks N1 through N5 are configured on VLAN Interfaces (VIs). You want networks N1, N2, and N3 to forward broadcast messages to network N4, but not to network N5.

Both examples provide two methods of configuring a UDP forwarder with the same results.

Example 1

To configure UDP forwarding, issue an add command to broadcast from all networks to network N4, using UDP port 200. Next, disable network N5 from broadcasting messages to other networks. The following steps and Figure 3-1 provide an example of this configuration:

Step	Action
1	At the `IP Config>` prompt enter: **add br**oadcast-forwarder **u**dp 200 all *network-4 ip-address* This command adds and enables all interfaces. (See Figure 3-1A.)
2	At the `IP Config>` prompt enter: **dis**able **br**oadcast-forwarding **u**dp 200 5 This command disables interface N5. (See Figure 3-1B.)

Figure 3-1: UDP Broadcast Forwarder Example

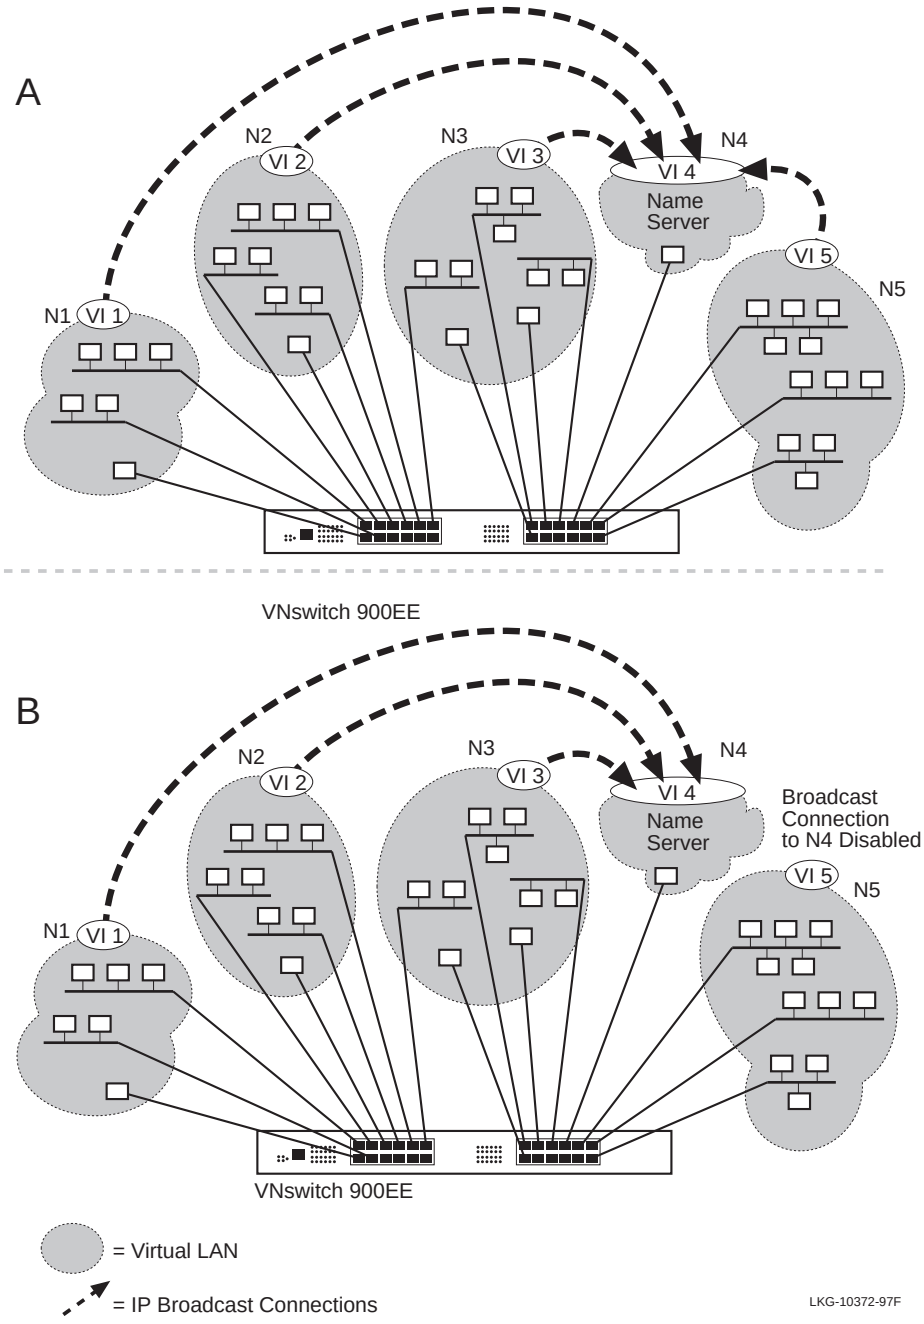

Example 2

To configure UDP forwarding, issue an add command to broadcast from networks N1, N2, and N3 to network N4, using UDP port 200. The following steps provide an example of this configuration:

Step	Action
1	At the `IP Config>` prompt enter: **add broadcast-forwarder udp 200 1 *network-4 ip-address*** This command adds and enables N1 interface.
2	At the `IP Config>` prompt enter: **add broadcast-forwarder udp 200 2 *network-4 ip-address*** This command adds and enables N2 interface.
3	At the `IP Config>` prompt enter: **add broadcast-forwarder udp 200 3 *network-4 ip-address*** This command adds and enables N3 interface.

Listing UDP Broadcast Forwarding

To list broadcast forwarding, at the `IP Config>` prompt, enter:

list broadcast-forwarding

This command displays the protocol, port, interface, state, and destination of the broadcast forwarder.

Configuring New Software

When you configure your router, you can configure it to accommodate future software upgrades. The DIGITAL VNswitch 900 series modules allow you to upgrade your module remotely to the latest release using TFTP. To accomplish an upgrade, the module may require a gateway IP address, which can be configured at any time.

When routing is enabled, the primary method to define the default gateway when performing a load or reload command is by using the `IP Config>` **set default network-gateway** command (preferred method), or as arguments to the load or reload commands. However, when using the DIGITAL clearVISN Flashloader or Recovery Manager applications, the default gateway cannot be specified on clearVISN.

If you cannot use the `IP Config>` **set default network-gateway** command, the `IP Config>` **set ip-host-only-default network-gateway** command can be used to define a gateway address. This method can only used when performing a software upgrade or a restore of configuration information. During the upgrade or restore process, routing protocols such as, RIP and OSPF are not enabled.

For additional information regarding software upgrade procedures, refer to the *DIGITAL VNswitch 900 Series Switch Management* guide.

Setting the IP Host-Only Default Network Gateway

To load new software to your router, the IP host-only default network gateway and subnet gateway commands may be required in conjunction with installing new software. To set the IP host-only default network gateway, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set ip-host-only-default network-gateway *ip-address***
2	Press Return.

Deleting the IP Host-Only Default Network Gateway

To delete the IP host-only default network gateway, at the `IP Config>` prompt, enter:

delete ip-host-only-default network-gateway

Setting the IP Host-Only Default Subnet Gateway

To set the IP host-only default subnet gateway, perform the following steps:

Step	Action
1	At the `IP Config>` prompt enter: **set ip-host-only-default subnet-gateway *ip-address***
2	Press Return.

Deleting the IP Host Only Default Subnet Gateway

To delete IP host-only default subnet gateway, at the `IP Config>` prompt, enter:

delete ip-host-only-default subnet-gateway

Listing IP Protocols

This command displays the configured state of the IP routing protocols (OSPF and RIP) along with whether ARP subnet routing is enabled or disabled. To list the configured routing protocols, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **list protocols**
2	Press Return.

Monitoring IP

This section describes tasks you can perform to monitor your router's IP protocol. To access the IP Monitor process, perform the following steps:

Step	Action
1	At the Main prompt (`Main>`) enter: **monitor**
2	Press Return. The Monitor prompt (`Monitor>`) is displayed.
3	If the prompt is not displayed, press Return a second time.
4	At the `Monitor>` prompt, enter: **ip**
5	Press Return. The `IP>` prompt is displayed.

From the `IP>` prompt, you can perform specific tasks to determine IP parameters, statistics, routing destinations, routing table contents, interface addresses, and routing paths.

Monitoring IP Access Control

You can monitor the router's access control system to determine which IP addresses and services your router is configured to handle. To display the access control list, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **access**
2	Press Return. The access control list is displayed.

Example

```
IP> access
Access control currently enabled
Access control run 8 times, 7 cache hits

List of access control records:

    Ty   Source          Mask       Destination    Mask       Beg Pro  End Pro  Beg Prt  End Prt  Invoc
 0  I    0.0.0.0         00000000   111.67.67.21   FFFFFFFF   6        6        25       25       0
 1  E    151.151.1.0     FFFFFF00   150.150.2.0    FFFFFFFF   0        255      0        655      0
 2  I    0.0.0.0         00000000   0.0.0.0        00000000   89       89       0        655      27
```

The `Invoc` column above contains the number of times an access control lookup match occurred for the associated filter. The invoc counter increments when the access control list is searched and the associated filter is the first matching entry.

The invoc counter can increment more than once per packet or not at all (when the result has been cached). Therefore, the entry does not necessarily represent the number of packets that were forwarded or terminated on a match for the associated filter.

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 `IP>` prompt, enter: **icmp-counters**
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
```

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 IP Interface Addresses

You can display a list of your router's IP interface addresses. Each address in the list contains an interface type, interface mask, and status. To display IP interface addresses, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **interface**
2	Press Return. IP interface addresses are displayed.

Example

`IP>` **interface**

```
 Ifc   Name        IP Addresses(es)     Mask(s)         Status
  25   VLAN/25     128.185.123.22       255.255.255.0   Down
  56   VLAN/56     128.185.124.23       255.255.255.0   Up
```

Monitoring IP Routing Table Contents

You can display the contents of the IP routing table and determine the protocol type, destination network, destination mask, cost, age, and next hop for each entry, as well as the routing table size. To display the IP routing table, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **dump**
2	Press Return. IP routing table information is displayed.

Example

```
IP> dump
 Type    Destination Net   Mask       Cost   Age   Next hop
 Stat*   0.0.0.0           00000000      0     0   16.20.48.254
 Sbnt    16.0.0.0          FF000000      1     0   16.20.48.254
  SPF*   16.20.0.0         FFFF0000      1     1   VLAN/26


Default gateway in use.
 Type    Cost   Age   Next hop
 Stat    0      0     16.20.48.254

Routing table size: 768 nets (67584 bytes), 4 nets known
```

Field	Description
`Type` (route type)	Indicates how the route was derived. • `Sbnt` — Indicates that the network is subnetted; such an entry is a placeholder only. • `Dir` — Indicates a directly connected network or subnet. • `RIP` — Indicates the route was learned through the RIP protocol. • `Del` — Indicates the route was deleted. • `Locl` — Indicates the local interface address. • `Stat` — Indicates a statically configured route. • `Fltr` — Indicates a routing filter. • `SPF` — Indicates that the route is an OSPF intra-area route. • `SPIA` — Indicates that the route is an OSPF inter-area route. • `SPE1, SPE2` — Indicates OSPF external routes (types 1 and 2, respectively). • `Rnge` — Indicates a route type that is an active OSPF area address range and is not used in forwarding packets.
`Destination Net`	IP destination network/subnet.
`Masks`	IP address mask.
`Cost`	Route cost.
`Age`	For RIP routes, the time that has elapsed since the routing table entry was last refreshed.
`Next hop`	Displays either the VLAN transmit interface, indicating the next hop exists, or the IP address of the next hop for static routes.

An asterisk (*) after the route type indicates that the route has a static or directly connected backup. A percent sign (%) after the route type indicates that RIP updates are always accepted for this network/subnet.

A number in parentheses at the end of the column indicates the number of equal-cost routes to the destination. The first hops belonging to these routes can be displayed with the IP route command.

Monitoring IP Routing Destinations

You can determine if IP routing destinations exist for a specific IP destination. To determine if a specific IP destination exists, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **route *ip-destination-address***
2	Press Return. IP routing table information is displayed.

Example

```
IP> route 128.185.232.0
 Destination:   128.185.232.0
 Mask:          255.255.255.0
 Route type:    RIP
 Distance:      3
 Age:           1
 Tag            0
 Next hop(s):   128.185.146.4   (VLAN/15)
```

Monitoring IP Routing Paths

You can display the entire IP routing path to a given destination, hop by hop, using traceroute. Traceroute sends out three probes and prints the IP address of the responder, together with the round-trip time associated with the response. To display IP routing paths, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **traceroute *ip-destination-address***
2	Press return. IP destination area address, packet size and the results of the traceroute are displayed

Example

```
IP> traceroute 128.185.142.239
TRACEROUTE 128.185.124.110:    56 data bytes
 1  128.185.142.7  16 ms  0  ms  0  ms
 2  128.185.123.22  16 ms  0  ms  16  ms
 3  * * *
 4  * * *
 5  128.185.124.110  16 ms  !  0  ms  !  0  ms
```

Field	Description
`Traceroute`	Displays the destination area address and the size of the packet being sent to that address.
`1`	Displays the first trace showing the destination's NSAP and the amount of time it took the packet to arrive at the destination. The packet is traced three times.
`Destination unreachable`	Indicates that no route to the destination is available.
`3 * * *` `4 * * *`	Indicates that the router is expecting some form of response from the destination, but the destination is not responding.

The traceroute is done whenever the destination is reached, an ICMP Destination Unreachable is received, or the path length reaches 32 router hops.

When a probe receives an unexpected result, several indications can be displayed. They are:

Probe Result	Description
!N	Indicates that an ICMP Destination Unreachable (net unreachable) was received.
!H	Indicates that an ICMP Destination Unreachable (host unreachable) was received.
!P	Indicates that 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.
!	Indicates that the destination was reached, but the reply sent by the destination was received with a TTL of 1. This usually indicates an error in the destination, prevalent in some versions of UNIX. The destination is inserting the probe's TTL in its replies. This unfortunately leads to a number of lines consisting solely of asterisks before the destination is finally reached.

Monitoring IP Static Routes

You can display a list of configured static routes, including filtered routes, default gateways and default subnet gateways using static routes. To display static routes, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **static**
2	Press return. Static routes are displayed.

Example

```
IP> static

Net                 Mask                Cost     Next hop
0.0.0.0             0.0.0.0             1        128.185.123.18
128.185.0.0         255.255.0.0         1        128.185.123.22
192.9.10.0          255.255.255.0       10       128.185.123.22
```

Field	Description
`Net`	Indicates the network address of the route.
`Mask`	Indicates the subnet mask of the IP address.
`Cost`	Indicates the cost of using this route.
`Next hop`	Indicates the next router a packet passes through that uses this route.

Each static route's destination is specified by an address-mask pair. Default gateways appear as static routes to destination 0.0.0.0. Default subnet gateways also appear as static routes to the entire IP subnetted network.

Monitoring IP Parameters

You can display a list of configured sizes of specific IP parameters. To display IP parameter sizes, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **sizes**
2	Press Return. IP parameter sizes are displayed.

Example

```
IP> sizes

 Routing table size         (#entries):            768

 Table entries used         (#entries):              3

 Reassembly buffer size        (bytes):          12000

 Largest reassembled pkt       (bytes):              0
```

Field	Description
`Routing table size`	Indicates the configured number of entries that the routing table maintains.
`Table entries in use`	Indicates the number entries used from the routing table.
`Reassembly buffer size`	Indicates the configured size of the reassembly buffer that is used to reassemble fragmented IP packets.
`Largest reassembled packet`	Indicates the largest IP packet that this router reassembled.

Monitoring IP Forwarding Statistics

You can display the statistics related to the IP forwarding process, including a count of routing errors, along with the number of packets that were dropped due to congestion. To display IP forwarding statistics, perform the following steps:

Step	Action
1	At the `IP>` prompt, enter: **counters**
2	Press Return. IP forwarding statistics are displayed.

Example

```
IP> counters
Receive:
 Total                          (ipInReceives)            0
 Dropped, Header Error          (ipInHdrErrors)           0
 Dropped, Unknown Protocol      (ipInUnknownProtos)       0
 Dropped, Internal Error        (ipInDiscards)            0
 To be Forwarded                (ipForwDatagrams)         0
 Terminated                     (ipInDelivers)            0


Transmit:
 Generated                      (ipOutRequests)           0
 Dropped, No Route              (ipOutNoRoutes)           0
 Total                                                    0


Reassembly:
 Requests                       (ipReasmReqds)            0
 Successful                     (ipReasmOKs)              0
 Failed                         (ipReasmFails)            0
```

(Continued on next page)

Monitoring IP Forwarding Statistics

```
Fragmentation:
 Successful                          (ipFragOKs)                0
 Failed                              (ipFragFails)              0
 Created                             (ipFragCreates)            0


Routing errors
 Routing table overflow                                         0
 Net unreachable                                             2539
 Bad subnet number                                              0
 Bad net number                                                 0
 Unhandled broadcast                                            0
 Unhandled multicast                                            0
 Unhandled directed broadcast                                   0
 Attempted forward of LL broadcast                           4048


Packets discarded through filter    0
IP multicasts accepted:             0


 Ifc     Name           IP input-packet overflow errors
 25      VLAN/25                                      0
 26      VLAN/26                                      0
 ...                                                  0
 56      VLAN/56                                      0
```

Field	Description
`Receive Group`	Received packets that were dropped, terminated or forwarded.
`Transmit Group`	Transmitted packets that were generated, dropped or forwarded.
`Reassembly Group`	Packets that were reassembled successfully, failed or requested.
`Fragmentation Group`	Packets that were fragmented successfully, failed or created.
`Routing table overflow`	Lists the number of routes that were discarded due to the routing table being full.
`Net unreachable`	Indicates the number of packets that were not forwarded due to unknown destinations. This does not count the number of packets that were forwarded to the authoritative router (default gateway).
`Bad subnet number`	Counts the number of packets or routes that were received for illegal subnets (all ones or all zeros).
`Bad net number`	Counts the number of packets or routes that were received for illegal IP destinations (for example, class E addresses).
`Unhandled broadcast`	Counts the number of (nonlocal) IP broadcasts received (these are not forwarded).
`Unhandled multicast`	Counts the number of IP multicasts that were received but whose address was not recognized by the router (these are discarded).
`Unhandled directed broadcast`	Counts the number of directed (nonlocal) IP broadcasts received when forwarding of these packets is disabled.
`Attempted forward of LL broadcast`	Counts the number of packets that are received having nonlocal IP addresses but were sent to a link level broadcast address. These are discarded.
`Packets discarded through filter`	Counts the number of received packets that were addressed to filtered networks/subnets. These are discarded silently.
`IP multicast accepted`	Counts the number of IP multicasts that were received and successfully processed by the router.
`IP input packet overflows`	Counts the number of packets that were discarded due to congestion at the forwarder's input queue. These counts are sorted by the receiving interface.

Chapter 4

Configuring and Monitoring the RIP Interface

Overview

Introduction

The Routing Information Protocol (RIP) is a distance-vector protocol (based on the Bellman-Ford technology) that allows routers to exchange information about destinations for computing routes throughout the network. Destinations may be networks or a special destination used to convey a default route.

Bellman-Ford algorithms make each router periodically broadcast its routing tables to all its neighbors. Then a router knowing its neighbors' tables can decide to which neighbor to forward a packet. Refer to Appendix D for examples on configuring RIP.

In This Chapter

This chapter discusses the following topics:

Configuring RIP

Enabling RIP

When configuring RIP, you can specify which set of routes the router advertises or accepts on each IP interface, or both. You can also specify how RIP information affects static routing. Since RIP uses broadcast messages for its routing updates, the format of the IP broadcast address must also be specified when using RIP. To enable RIP, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable rip**
2	Press Return. RIP is now enabled.

When RIP is enabled, use the **enable/disable sending** commands to configure its routing update sending behavior. Its routing update receiving behavior is defined by the **enable/disable receiving** and **enable/disable override** commands.

Disabling RIP

To disable RIP, at the `IP Config>` prompt, enter:

disable rip

RIP Limitations

RIP is primarily intended for use in homogeneous networks of moderate size. Because of this, RIP has the following limitations:

- Autonomous system (AS) diameter limitation of 15 hops.
- RIP's metric (hop count) cannot adequately describe variations in a path's characteristics, sometimes resulting in less than optimal routing.
- Slow to find new routes when the network changes. This search consumes considerable bandwidth, and in extreme cases, exhibits a slow convergence behavior referred to as a *count to infinity*.

NOTE

All bridging router interfaces running RIP must have the same subnet mask.

Enabling RIP Flags

By default, RIP advertises all network and subnet routes on all interfaces of the router. Once RIP is enabled, you can configure what it listens to and what it advertises by setting the various RIP flags. These flags are configured on a per-IP-interface basis. The following commands can enable or disable the various flags:

Flag Type	Command Syntax
Sending net routes	**enable sending net-routes** **disable sending net-routes**
Sending subnet routes	**enable sending subnet-routes** **disable sending subnet-route**s
Sending static routes	**enable sending static-routes** **disable sending static-routes**
Sending default routes	**enable sending default-routes** **disable sending default-routes**
Receiving RIP	**enable receiving rip** **disable receiving rip**
Receiving dynamic nets	**enable receiving dynamic nets** **disable receiving dynamic nets**
Receiving dynamic subnets	**enable receiving dynamic subnets** **disable receiving dynamic subnets**
Overriding default	**enable override default** **disable override default**
Overriding static routes	**enable override static-routes** **disable override static-routes**
Sending Poisoned Reversed Routes	**enable sending poisoned reversed routes** **disable sending poisoned reversed routes**

Customizing RIP

In IP, you can customize RIP with a number of configurable flags. Most flags take effect on a specified IP interface address. These flags control sending and receiving RIP information about each router interface.

The set of routes sent out from a particular address is the union of the routes selected by setting any of the following four flags. Subnet-level routes are sent only when the destination subnet is a member of the same IP network as the sending address.

Send Flag Name	Description
Send Net Routes	Sends all network-level routes.
Send Subnet Routes	Sends appropriate subnet-level routes.
Send Default Routes	Advertises a default route if the router itself has a default route.
Send Static and Direct Routes	Advertises all directly connected networks and statically configured routes.

The following flags control how information received by RIP is incorporated into the router's routing tables. Certain flag settings allow RIP routes to override static routing information, but only if the RIP metric is better than the static route's metric.

Receive Flag Name	Description
Override Default	RIP packets received on this IP interface may override the router's default gateway.
Override Static Routes	RIP packets received on this IP interface may override any of the router's statically configured routing information.
Disable RIP Receive	RIP packets received on this IP interface are ignored.
Receive Dynamic Net Routes	When not set, the router accepts RIP updates only for those networks that are specified in an **add accept-rip-route** command.
Receive Dynamic Subnets	When not set, the router accepts RIP updates only for those subnets that are specified in an **add accept-rip-route** command.

Setting RIP Broadcasts

The RIP protocol uses IP broadcast when sending its routing updates. Since there are different formats of IP broadcast in use, you must specify which broadcast format to use. You specify IP broadcast format on a per-interface basis by using the following command procedures:

Step	Action
1	At the `IP Config>` prompt, enter: **set broadcast-address *ip-interface-address***
2	Select either `local-wire` or `network` from the following prompt: `Use a NET or LOCAL-WIRE style address NETWORK?`
3	Select whether you want the rest of the broadcast address filled with either ones or zeros from the following prompt: `Fill pattern for wildcard part of address (0 or 1) 0?`
4	Press Return. The RIP broadcast is now enabled.

Converting from RIP to OSPF

To convert your autonomous system (AS) from RIP to OSPF, install OSPF on one router at a time, leaving RIP running. Gradually, all your internal routes shift from being learned through RIP to being learned by OSPF (OSPF routes have precedence over RIP routes). If you want to have your routes look exactly as they did under RIP (to check that the conversion is working properly), use hop count as your OSPF metric. This is done by assigning the cost of each OSPF interface to 1.

The size of your OSPF system must be estimated when the protocol is enabled. This size estimate should reflect the final size of the OSPF routing domain.

After installing OSPF on your routers, turn on AS boundary routing in all those routers that still need to learn routes through other protocols (RIP, and statically configured routes). Keep the number of these AS boundary routers to a minimum.

Finally, you can disable the receiving of RIP information about all those routers that are not AS boundary routers.

Configuring Accept RIP Routes

Adding Accept RIP Route

Adding accept RIP route allows an interface to accept a RIP route when RIP route filtering is enabled for an interface. To add accept RIP route, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **add accept-rip-route *ip-network/subnet***
2	Press Return.

You can list networks/subnets that are already entered using the **list rip-routes-accept** command. You can enable the input filtering of RIP routes on a per-IP-interface basis. This is done separately for network-level routes (for example, a route to 10.0.0.0) and for subnet-level routes (for example, a route to 128.185.0.0). To enable input filtering of network-level routes on an IP interface, use the **disable dynamic nets** command. To enable input filtering of subnet-level routes, use the **disable dynamic subnets** command.

Example: `IP Config>` **add accept-rip-route 10.0.0.0**

Deleting Accept RIP Route

To delete an accept RIP route, at the `IP Config>` prompt, enter:

delete accept-rip-route *net-number*

Listing Accept RIP Route

To list an accept RIP route, at the `IP Config>` prompt, enter:

list rip-routes-accept

This command displays the set of routes that the RIP routing protocol always accepts. See the IP configuration commands Enabling Receiving Dynamic Nets and Enabling Receiving Dynamic Subnets for more information.

Configuring RIP to Override Default Routes

This command configures the conditions under which the router originates a RIP default route, and the cost that will be used when originating the default. To set the originate RIP default, perform the following steps:

Step	Action

1	At the `IP Config>` prompt enter:
	set originate-rip-default
	`Always originate default route? [No]:` **yes**
	`Originate default if OSPF routes available? [No]:`
	`Originate default of cost [1]?`

2	Press Return.

Prompt	Description
`Always originate default route? [No]:`	With the **send default routes** flag enabled, **Yes** sets the router to always advertise a RIP default route whether there is a route in the routing table or not. **No** enables the following prompt.
`Originate default if OSPF routes available? [No]:`	With the **send default routes** flag enabled, **Yes** sets the router to always advertise a RIP default route when OSPF routes are available. **No** only advertises a default route if the **send default routes** flag is enabled.
`Originate default of cost [1]?`	Selects the metric for the RIP default route. The range is from 0 to 16 with a default of 1.

Configuring RIP to Override Default and Static Routes

Enabling RIP Override Routes

Enabling the RIP Override Default

This command enables received RIP information to override the router's default gateway. It is invoked on a per-IP-interface basis. To enable the override default, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable override default** ***gateway-ip-interface-address***
2	Press Return. When the enable override default command is invoked, default RIP routes received on the interface with a specified IP interface address overwrite the router's current default gateway, providing the cost of the new default is cheaper.

Example: `IP Config>` **enable override default 128.185.123.22**

Disabling RIP Override Default

To disable the RIP override default, at the `IP Config>` prompt, enter:

disable override default ***gateway-ip-interface-address***

Enabling RIP Override Static Routes

This command enables received RIP information to override some of the router's statically configured routing information. It is invoked on a per-IP-interface basis. To enable overriding of static routes, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable override static-routes *ip-interface-address***
2	Press Return. When the enable override static routes command is invoked, RIP routing information received on the interface with a specified IP interface address overwrite statically configured network/ subnet routes, providing the cost of the RIP information is cheaper.

Example: `IP Config>` **enable override static-routes 128.185.123.22**

Disabling RIP Override Static Routes

To disable RIP overriding of static routes, at the `IP Config>` prompt, enter:

disable override static-routes *ip-interface-address*

Configuring Receiving RIP, Dynamic Nets/Subnets

Enabling RIP Reception on an Interface

This command modifies the processing of RIP updates that are received on a particular interface. To enable receiving RIP, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable receiving rip *ip-interface-address***
2	Press Return.

This command is enabled by default.

If you invoke the disable receiving RIP command, no RIP updates are accepted on interface *ip-interface-address*.

Example: `IP Config>` **enable receiving rip 128.185.123.22**

Disabling RIP Reception on an Interface

To disable RIP reception, at the `IP Config>` prompt, enter:

disable receiving rip *ip-interface-address*

Enabling Receiving Dynamic Nets

This command modifies the processing of RIP updates that are received on a particular interface. To enable receiving dynamic nets, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable receiving dynamic net *ip-interface-address***
2	Press Return.

This command is enabled by default.

If you invoke the disable receiving dynamic nets command, RIP updates received on interface *ip-interface-address* cannot accept any network-level routes unless they were previously specified in an **add accept-rip-route** command.

Example: **enable receiving dynamic nets 128.185.123.22**

Disabling Receiving Dynamic Nets

To disable receiving dynamic nets, at the `IP Config>` prompt, enter:

disable receiving dynamic net ***ip-interface-address***

Enabling Receiving Dynamic Subnets

This command modifies the processing of RIP updates that are received on a particular interface. To enable receiving dynamic subnets, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable receiving dynamic subnets** ***ip-interface-address***
2	Press Return.

This command is enabled by default.

Example: `IP Config>` **enable receiving dynamic subnets 128.185.123.22**

Disabling Receiving Dynamic Subnets

To disable receiving dynamic nets, at the `IP Config>` prompt, enter:

disable receiving dynamic subnet ***ip-interface-address***

If you invoke the disable receiving dynamic subnets command, RIP updates received on interface *ip-interface-address* cannot accept any subnet-level routes unless they were previously specified in an **add accept-rip-route** command.

Configuring Sending of Routes in RIP

Enabling Sending Default Routes in RIP

This command determines the contents of RIP updates that are sent out a particular interface. To enable sending default routes, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable sending default-routes *ip-interface-address***
2	Press Return.

The effect of the enable sending command is cumulative. Each separate enable sending command specifies that a certain set of routes is advertised from a particular interface. A route is included in a RIP update only if it was included by at least one of the enable sending commands. The enable sending default-routes command specifies that the default route (if one exists) is included in RIP updates sent out interface *ip-interface-address*.

Example: `IP Config>` **enable sending default-routes 128.185.123.22**

NOTE

Some settings of the enable sending commands are redundant. For example, if you invoke enable sending net routes and enable sending subnet routes for a particular interface, there is no need to also specify enable sending static routes (because each static route is either a network-level or subnet route). By default, when you first enable RIP, sending net routes and sending subnet routes are enabled for each interface, while sending static routes and sending default are disabled.

Disabling Sending Default Routes in RIP

To disable sending default routes, at the `IP Config>` prompt, enter:

disable sending default-routes *ip-interface-address*

Enabling Sending Net Routes

This command determines the contents of RIP updates that are sent out a particular interface. To enable sending net routes, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable sending net-routes *ip-interface-address***
2	Press Return.

The effect of the enable sending command is cumulative. Each separate enable sending command specifies that a certain set of routes is advertised from a particular interface. A route is included in a RIP update only if it was included by at least one of the enable sending commands. The enable sending network-routes command specifies that all network-level routes are included in RIP updates sent out interface *ip-interface-address*. A network-level route is a route to a single class A, B, or C IP network.

Example: `IP Config>` **enable sending net-routes 128.185.123.22**

Disabling Sending Net Routes

To disable sending net routes, at the `IP Config>` prompt, enter:

disable sending net-routes *ip-interface-address*

Enabling Sending Poisoned Reverse Routes

This command determines the contents of RIP updates that are sent out a particular interface. To enable sending poisoned reverse routes, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable sending poisoned-reverse-routes *ip-interface-address***
2	Press Return.

The effect of the enable sending command is cumulative. Each separate enable sending command specifies that a certain set of routes is advertised from a particular interface. A route is included in a RIP update only if it was included by at least one of the enable sending commands. The enable sending poisoned-reverse-routes command specifies that all routes learned on this interface are sent out at a cost of 16. A network-level route is a route to a single class A, B, or C IP network.

Example: `IP Config>` **enable sending poisoned-reverse-routes 128.185.123.22**

Disabling Sending Poisoned Reverse Routes

To disable sending poisoned-reverse-routes, at the `IP Config>` prompt, enter:

disable sending poisoned-reverse-routes ip-interface-*address*

Enabling Sending Subnet Routes

This command determines the contents of RIP updates that are sent out a particular interface. To enable sending subnet routes, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable sending subnet-routes *ip-interface-address***
2	Press Return.

The effect of the enable sending command is cumulative. Each separate enable sending command specifies that a certain set of routes is advertised through a particular interface. A route is included in a RIP update only if it was included by at least one of the enable sending commands. The enable sending subnet routes command specifies that all subnet routes are included in RIP updates sent out interface ip-interface-address. However, a subnet route is included only if *ip-interface-address* connects directly to a subnet of the same IP subnetted network.

Example: `IP Config>` **enable sending subnet-routes 128.185.123.22**

Disabling Sending Subnet Routes

To disable sending subnet routes, at the `IP Config>` prompt, enter:

disable sending subnet-routes *ip-interface-address*

Enabling Sending Static Routes

This command determines the contents of RIP updates that are sent out a particular interface. To enable sending static routes, perform the following steps:

Step	Action
1	At the `IP Config>` prompt, enter: **enable sending static-routes *ip-interface-address***
2	Press Return.

The effect of the enable sending command is cumulative. Each separate enable sending command specifies that a certain set of routes is advertised through a particular interface. A route is included in a RIP update only if it was included by at least one of the enable sending commands. The enable sending static routes command specifies that all statically configured and directly connected routes are included in RIP updates sent out interface *ip-interface-address*.

Example: `IP Config>` **enable sending static-routes 128.185.123.22**

Disabling Sending Static Routes

To disable sending static routes, at the `IP Config>` prompt, enter:

disable sending static-routes *ip-interface-address*

Chapter 5

Configuring and Monitoring the OSPF Interface

Overview

Introduction

The Open Shortest Path First (OSPF) protocol is a link state dynamic routing protocol that detects and learns the best routes to (reachable) destinations. OSPF can quickly perceive changes in the topology of an autonomous system, and after a short convergence period, calculate new routes.

Each router running the OSPF protocol has a database describing a map of the routing domain. This database is identical in all participating routers. From this database, the IP routing table is built through the construction of a shortest-path tree, with the router itself as root. The routing domain refers to an autonomous system (AS) running the OSPF protocol, which allows you to split the AS into regions called *areas*. These areas are a collection of contiguous networks, and the topology of any one area is hidden from that of the other areas, significantly reducing routing traffic. Refer to Appendix D for examples on configuring OSPF.

OSPF commands are not dynamic. You must restart the module before any commands take effect.

In This Chapter

This chapter discusses the following topics:

Configuring the OSPF Protocol

Enabling the OSPF Protocol

OSPF configuration is done through the protocol's own configuration console. The OSPF routing protocol is enabled on an interface-by-interface basis. Each OSPF interface is assigned a cost. Also, an estimate of the OSPF database's size must be given, and the interaction between OSPF and RIP must be defined. Size estimates for the OSPF link state database tells the router software approximately how much memory to reserve for OSPF. Use the following procedures to initially configure OSPF.

To enable OSPF, perform the following steps:

Step	Action
1	At the `Config>` prompt, enter: **ospf**
2	Press Return. The `OSPF Config>` prompt is displayed.
3	At the `OSPF Config>` prompt, enter: **enable ospf routing protocol**
4	Enter the total number of AS external routes imported into the OSPF routing domain. **Example:** `Estimated # external routes[0]?` **200**
5	Enter the total number of OSPF routers in the routing domain. **Example:** `Estimated # OSPF routers [0]?` **60** `OSPF Config>`
6	Reboot the module to enable OSPF.

Once the OSPF protocol is enabled, proceed to Configuring Attached OSPF Areas.

Disabling the OSPF Routing Protocol

To disable the OSPF routing protocol, at the `OSPF Config>` prompt, enter:

disable ospf routing protocol

Reboot the module to disable OSPF.

Configuring Attached OSPF Areas

Setting OSPF Areas

The next step in the configuration process is setting the parameters that define the OSPF areas that are directly attached to the router. If no areas are defined, the router software assumes that all the router's directly attached networks belong to the backbone area (area ID 0.0.0.0).

To set the parameters for an OSPF area, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt enter: **set area**
2	Enter the area number. **Example:** `Area number [0.0.0.0]?` **1.1.2.2**
3	Enter the authentication type. **Example:** `Authentication type [0]?` **1**
4	Enter whether the area is a stub area. **Example:** `Is this a stub area? (No):` **y**
5	Enter a stub default cost. **Example:** `Stub default cost [0]:`
6	Enter import summaries. **Example:** `Import summaries?` **y** `OSPF Config>`

Where:

- **Area number** is the OSPF area ID number, which is a contiguous group of networks defined by a list of address ranges, each indicated by a combination of the IP address and address mask. A network belongs to an area if its address is in the list.
- **Authentication type** (security scheme) to be used in the area. A 1 indicates a simple password; a 0 indicates that no authentication is necessary.
- **Stub area** designation. If you designate YES:
  - — Disables external routes from being flooded into the router. The area does not receive any AS external link advertisements, reducing the size of the area's OSPF database and decreasing memory usage for external routers in the stub area.
  - — You cannot configure virtual links through a stub area.
  - — You cannot configure a router within the stub area as an AS boundary router.
- **Import summaries** (inter-area routes) allows summary routes to be flooded into the router.

External Routing in Stub Areas. You cannot configure the backbone as a stub area. External routing in stub areas is based on a default route. Each border area router attaching to a stub area originates a default route for this purpose. The cost of this default route is also configurable in the OSPF set area command.

Deleting OSPF Areas

To delete an area, at the `OSPF Config>` prompt, enter:

delete area ***area#***

Listing OSPF Areas

To list an area, at the `OSPF Config>` prompt, enter:

list area

Adding Ranges to OSPF Areas

Once you have defined attached OSPF areas, you can add a range of addresses to the area. OSPF areas are defined in terms of address ranges. External to the area, a single route is advertised for each address range. For example, if an OSPF area consists of all subnets of the class B network 128.185.0.0, it is defined as consisting of a single address range. The address range is specified as an address of 128.185.0.0 with a mask of 255.255.0.0. Outside of the area, the entire subnetted network is advertised as a single route to network 128.185.0.0, but this can be inhibited. To add range addresses to defined OSPF areas using the above example, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **add range** ***area# ip-address mask*** `Inhibit advertisement? [No]:`
2	Press Return.

Example: `OSPF Config>` **add range 0.0.0.2 128.185.0.0 255.255.0.0 inhibit**

Deleting Ranges from OSPF Areas

To delete a range of areas, at the `OSPF Config>` prompt, enter:

delete range ***area# ip-address***

Configuring Routing Interfaces

Setting OSPF Interfaces

The **set interface** command adds an OSPF interface or changes an existing one. This command can be used only when the interface is disabled. It is used with the **enable interface** command.

When setting OSPF interfaces to routers attached to common network segments, you must enter the same values for the following parameters:

- Hello interval
- Dead router interval
- Authentication key (if an authentication type of 1 [simple password] is used)

To set the OSPF parameters for the router's interfaces, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **set interface**
2	Enter the IP address. **Example:** `Interface IP address [0.0.0.0]?` **16.24.11.251**

3	Enter the remaining values when prompted.
	`Attaches to area [0.0.0.0]?`
	`Retransmission Interval (in seconds) [5]?`
	`Transmission Delay (in seconds) [1]?`
	`Router Priority [1]?`
	`Hello Interval (in seconds) [10]?`
	`Dead Router Interval (in seconds) [40]?`
	`Type Of Service 0 cost [1]?`
	`Authentication Key []?` **auth_key**
	`Retype Auth. Key []?` **auth_key**
	`OSPF Config>`

Deleting OSPF Interfaces

To delete an interface, at the `OSPF Config>` prompt, enter:

delete interface ***ip-address***

This command can be used only when the interface is disabled or when dynamic management is disabled.

Listing OSPF Interfaces

To list an interface, at the `OSPF Config>` prompt, enter:

list interface

Example:

`OSPF Config>`**list interface**

```
                        - - Interface configuration - -

IP address        Sta     Area            Cost   Rtrns   TrnsDly   Pri   Hello   Dead
16.24.11.251      Ena     1.1.2.2            1       5         1     1      10     40


                           - - Authentication Keys - -

IP Address         AuType                   Key (Hex/Ascii)
128.185.138.21                              0x617574685F6B6579 “Auth_key”
```

Field	Description
IP address	An IP address is displayed for each interface, together with configured parameters.
Sta	Indicates the status of the interface is enabled or disabled.
Area	The OSPF area to which the interface attaches.
Cost	The TOS 0 cost (or metric) associated with the interface.
Rtrns	The retransmission interval, which is the number of seconds between retransmissions of unacknowledged routing information.
TrnsDly	The transmission delay, which is an estimate of the number of seconds it takes to transmit routing information over the interface (it must be greater than 0).
Pri	The interface's router priority, which is used when selecting the designated router.
Hello	The number of seconds between Hello Packets sent out the interface.
Dead	The number of seconds after Hellos cease to be heard that the router is declared down.

Enabling OSPF Interfaces

The **enable interface** command is a dynamic command that activates an OSPF interface and starts sending and receiving packets over it. This command is used with the **set interface** command, and once the interface is set, you can dynamically enable it using the enable command.

To enable an OSPF interface, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **enable interface *ip-address***
2	Press Return. The specified OSPF interface is enabled.

Example: `OSPF Config>` **enable interface 16.24.11.251**

Disabling OSPF Interfaces

To disable the interface, at the `OSPF Config>` prompt, enter:

disable interface *ip-address*

Configuring Nonbroadcast Interface Parameters

Setting Nonbroadcast Network Interface Parameters

If the router is connected to a nonbroadcast, multiaccess network, such as an X.25 PDN, you have to configure parameters to help the router discover its OSPF neighbors. This configuration is necessary only if the router is eligible to become the designated router of the nonbroadcast network.

To configure a router for a nonbroadcast network, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **set non-broadcast**
2	Enter the IP address. **Example:** `Interface IP address [0.0.0.0]?` **128.185.138.19**
3	Enter the Poll Interval. **Example:** `Poll Interval [120]?` `OSPF Config>`

Then configure the IP addresses of all other OSPF routers that are attached to the nonbroadcast network. For each router configured, you must also specify its eligibility to become the designated router.

Deleting Nonbroadcast Network Interface Parameters

To delete a nonbroadcast network interface, at the `OSPF Config>` prompt, enter:

delete non-broadcast ***ip-address***

Listing Nonbroadcast Network Interface Parameters

To list a nonbroadcast network description, at the `OSPF Config>` prompt, enter:

list non-broadcast network description

This command lists all information related to interfaces connected to nonbroadcast networks. For each nonbroadcast interface, as long as the router is eligible to become the designated router on the attached network, the polling interval is displayed with a list of the router's neighbors on the nonbroadcast network.

Adding Neighbors to Nonbroadcast Networks

You can add neighbors to nonbroadcast networks. If the router is connected to a nonbroadcast, multiaccess network, such as an X.25 PDN, you have to use this command to help the router discover its OSPF neighbors. This configuration is only necessary if the router is eligible to become the designated router of the nonbroadcast network. Configure the IP addresses of all other OSPF routers that are attached to the nonbroadcast network. For each router configured, you must also specify its eligibility to become designated router. To add neighbors to OSPF areas, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **add neighbor**
2	Press Return.
3	You will then be prompted for: `Interface IP address [0.0.0.0]?` **128.185.138.19** `IP address of neighbor [0.0.0.0]?` **128.185.138.21** `Can that router become Designated Router on the net [Yes]?`

Deleting Neighbors from Nonbroadcast Networks

To delete a neighbor, at the `OSPF Config>` prompt, enter:

delete neighbor ***ip-address ip-address-of-neighbor***

Listing Neighbors from Nonbroadcast Networks

To list a neighbor, at the `OSPF Config>` prompt, enter:

list neighbor

This command lists all configured OSPF IP addresses, their configured neighbors, and whether the neighbor is eligible to be a designated router (DR).

Configuring AS Boundary Routing

Enabling AS Boundry Routing

To import routes learned from other protocols (RIP, and statically configured information) into the OSPF domain, enable AS boundary routing. You must do this even if the only route you want to import is the default route (destination 0.0.0.0).

When enabling AS boundary routing, you are asked which external routes you want to import. You can choose to import, or not to import, routes belonging to several categories. Independent of the external categories, you can also configure whether or not to import subnet routes into the OSPF domain. This configuration item defaults to OFF (subnets not imported).

The metric type used in importing routes determines how the imported cost is viewed by the OSPF domain. When comparing two type 2 metrics, only the external cost is considered in picking the best route. When comparing two type 1 metrics, the external and internal costs of the route are combined before making the comparison.

To enable AS boundary, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **enable as boundary**

2	Enter the remaining values when prompted.
	`Import RIP routes? [No]:`
	`Import static routes? [No]:`
	`Import direct routes? [No]:` **y**
	`Import subnet routes? [No]:`
	`Always originate default route? [No]:`**y**
	`Originate as type 1 or 2 [2]?`
	`Default route cost [1]?`
	`Default forwarding address [0.0.0.0]?`
	`OSPF Config>`

Disabling AS Boundary Routing

To disable as boundary routing, at the `OSPF Config>` prompt, enter:

disable as boundary routing

Configuring For Routing Protocol Comparisons

If you use a routing protocol in addition to OSPF, or when you change your routing protocol to OSPF, you must set the routing protocol comparison.

OSPF has a 4-level route hierarchy. The **set comparison** command tells the router where the RIP/static routes fit in the OSPF hierarchy. The two lower levels consist of the OSPF internal routes. OSPF intra-area and inter-area routes take precedence over information obtained from any other sources, all of which are located on a single level.

To set the routing protocol comparison, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **set comparison**
2	Enter the external comparison type. **Example:** `Compare to type 1 or 2 externals [2]?` **1** `OSPF Config>`

To put the RIP/static routes on the same level as OSPF external type 1 routes, set the comparison to 1. To put the RIP/static routes on the same level as OSPF external type 2 routes, set the comparison to 2. The default setting is 2.

Example: Suppose the comparison is set to 2. In this case, when RIP routes are imported into the OSPF domain, they are imported as type 2 externals. All OSPF external type 1 routes override received RIP routes, regardless of metric. However, if the RIP routes have a smaller cost, the RIP routes override OSPF external type 2 routes.

Configuring OSPF Virtual Links

Setting OSPF Virtual Links

To maintain backbone connectivity you must have all of your backbone routers interconnected either by permanent or virtual links. Virtual links may be configured between any two area border routers that share a common nonbackbone and nonstub area. Virtual links must be configured in each of the link's two endpoints.

NOTE

If you are configuring an area border router (ABR) that does not directly attach to the backbone area, you must create a virtual link and an area 0.0.0.0.

The **set virtual-link** command adds an OSPF virtual link interface or changes an existing one. This command is used when the interface is disabled. It is used with the **enable interface** command.

To set a virtual link, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **set virtual-link**

2	Enter the remaining values when prompted.
	`Virtual endpt. (Router ID) [0.0.0.0]?` **128.185.138.21**
	`Link's transit area [0.0.0.1]?`
	`Retransmission Interval (in seconds) [10]?`
	`Transmission Delay (in seconds) [5]?`
	`Hello Interval (in seconds) [30]?`
	`Hello Interval (in seconds) [30]?`
	`Dead Router Interval (in seconds) [180]?`
	`Authentication Key []?` **3-14159**
	`OSPF Config>`

Deleting OSPF Virtual Links

To delete a virtual link, at the `OSPF Config>` prompt, enter:

delete virtual-link

This command can be used only when the interface is disabled or when dynamic management is disabled.

Listing OSPF Virtual Links

To list a virtual link, at the `OSPF Config>` prompt, enter:

list virtual-link

This command lists all virtual links that were configured with this router as endpoint. **Virtual endpoint** indicates the OSPF router ID of the other endpoint. **Transit area** indicates the nonbackbone area through which the virtual link is configured. Virtual links are considered treated by the OSPF protocol similarly to point-to-point networks. The other parameters listed in the command (`Rtrns, TrnsDly, Hello and Dead`) are maintained for all interfaces.

Enabling OSPF Virtual Links

The enable OSPF virtual-link command is a dynamic command that activates an OSPF interface and starts sending and receiving packets over it. This command is used with the **set virtual-link** command, and once the virtual link is set, you can dynamically enable it using the enable command.

Figure 5-1 illustrates an OSPF routing domain with areas configured with virtual links. Although the backbone area within the domain must be contiguous, you can configure areas that are not physically contiguous using virtual links, as shown.

To enable an OSPF interface, at the `OSPF Config>` prompt, enter:

enable virtual-link *ip-address*

Example: `OSPF Config>` **enable virtual-link 16.24.11.251**

Disabling OSPF Virtual Links

To disable the virtual link, at the `OSPF Config>` prompt, enter:

disable virtual-link *ip-address*

Figure 5-1: OSPF Areas Configured Using Virtual Links

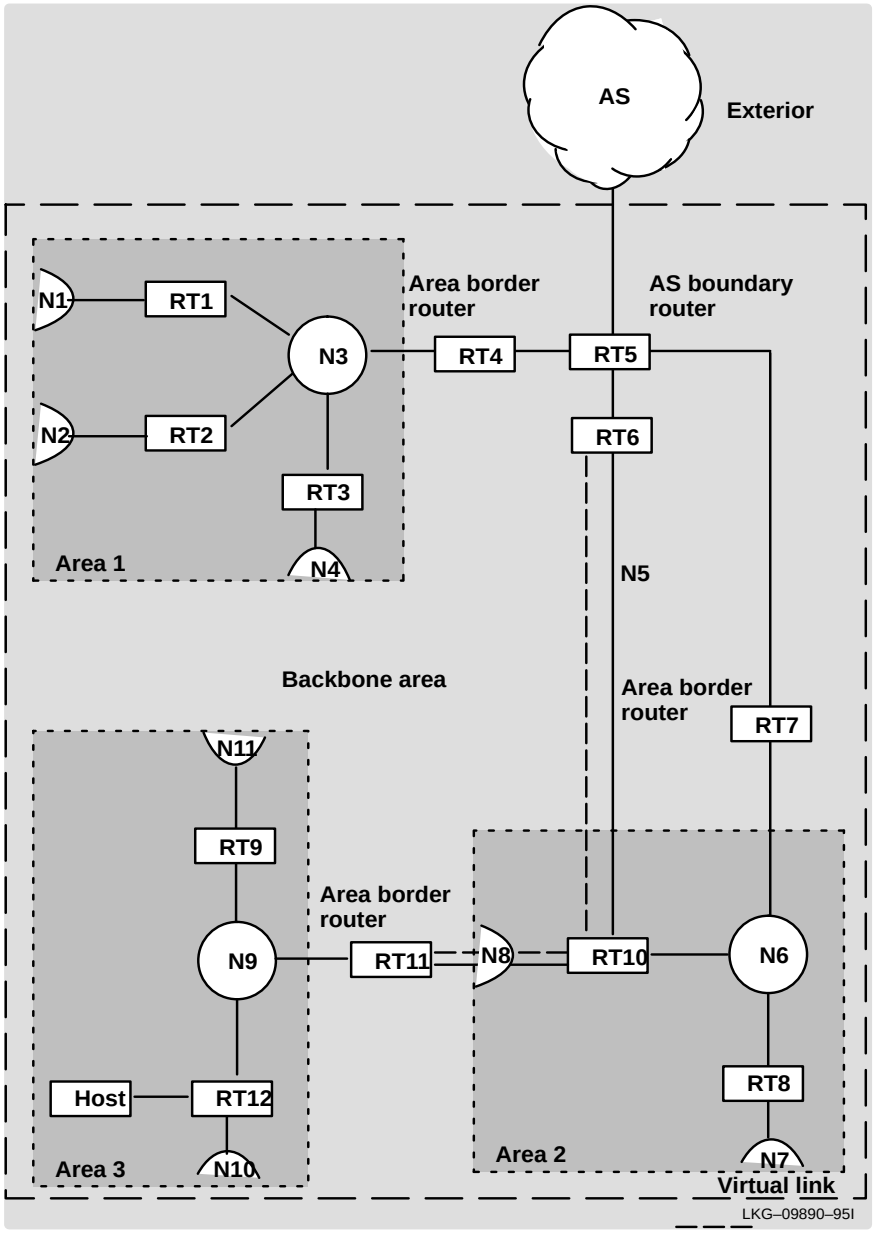

Configuring OSPF Router IDs

Every router in an OSPF routing domain must be assigned a 32-bit router ID. The current OSPF implementation sets the OSPF router ID to be the address of the first OSPF interface appearing in the router's configuration.

The OSPF router ID can also be explicitly set by the using the `Config>` **ip set router id** command. In this case, the router ID must still be one of the router's IP interface addresses.

To set an OSPF router ID using the **set router id** command, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **exit** The `Config>` prompt appears.
2	At the `Config>` prompt, enter: **ip set router-id *ip-interface-address*** **Example: set router-id 128.185.138.21**
3	To return to OSPF, at the `Config>` prompt, enter: **ospf**

Example Configuration Procedure for OSPF

The following example demonstrates the configuration of OSPF on a router.

In this example, each parameter that has a default value is given that value. For parameters that have no default (for example, addresses), the procedure uses an arbitrary value.

The following steps summarize the procedure you can use to configure the OSPF protocol on router interfaces. You may have to change the values or parameters given here according to your network setup.

NOTE

You must restart the router for these values to take effect.

To configure OSPF on the router, perform the following steps:

Step	Action	Configuration Script
1	**Enable the OSPF protocol.**	Config> **ospf** Open SPF-Based Routing Protocol configuration console OSPF Config> **enable ospf** Estimated - external routes [0]? **200** Estimated - OSPF routers [0]? **60**
2	**Define the attached OSPF areas.**	OSPF Config> **set area** Area number [0.0.0.0]? **1.1.2.2** Authentication Type [0]? **1** Is this a stub area [No]?
3	**Set the OSPF interface.**	OSPF Config> **set interface** Interface IP address [0.0.0.0]? **16.24.11.251** Attaches to area [0.0.0.0]? Retransmission Interval (in seconds) [5]? Transmission Delay (in Seconds) [1]? Router Priority [1]? Hello Interval (in seconds) [10]? Dead Router Interval (in seconds) [40]? Type of Service 0 cost [1]? Authentication key []? **auth_key** Retype Auth. Key []? **auth_key**
4	**Enable the interface.**	OSPF Config> **enable interface 16.24.11.251**
5	**Set nonbroadcast interface parameters.**	OSPF Config> **set non-broadcast** Interface IP address [0.0.0.0]? **128.185.138.19** Poll interval [120]?
6	**Add a neighbor.**	OSPF Config> **add neighbor** Interface IP Address [0.0.0.0]? **128.185.138.19** IP Address of Neighbor [0.0.0.0]? **128.185.138.21** Can the router become Designated Router on this net [yes]?

Example Configuration Procedure for OSPF

Step	Action	Configuration Script
7	**Enable AS boundary routing.**	OSPF Config> **enable as boundary** Import RIP routes? [No]: Import static routes? [No]: Import direct routes? [No]: **yes** Import subnet routes? [No]: Always originate default route? [No]: Originate as type 1 or 2 [2]? Default route cost [1]? Default forwarding address [0.0.0.0]?
8	**Configure for routing protocol comparisons.**	OSPF Config> **set comparison** Compare to type 1 or 2 externals [2]?
9	**Set up virtual links.**	OSPF Config> **set virtual** Virtual endpoint (Router ID) [0.0.0.0]? **128.185.138.21** Link's transit area [0.0.0.1]? Retransmission Interval (in seconds) [10]? Transmission Delay (in seconds) [5]? Hello Interval (in Seconds) [30]? Dead Router Interval (in seconds) [180]? Authentication Key []? **3-14159**
10	**Set the OSPF router ID.**	OSPF Config> **exit** Config> **ip set router-id 128.185.138.21**
11	**Return to OSPF.**	Config> **ospf**

Listing OSPF Configuration Information

Once your router has been configured for OSPF, you can display your router's configuration information by using the **list all** command.

To determine the type of configuration information that can be displayed, perform the following steps:

Step	Action
1	At the `OSPF Config>` prompt, enter: **list all**
2	Press Return.

Listing OSPF Configuration Information

Example

```
OSPF Config> list all
                            - - Global Configuration - -

                OSPF Protocol:                Enabled
                # AS ext. routes:             200
                Estimated # routers:          60
                External comparison:          Type 2
                AS boundary capability:       Enabled
                Import external routes:       DIR
                Orig. default route:          No (0,0.0.0.0)
                Default route cost:           (1 type 2)
                Default forward addr:         0.0.0.0


                           - - Area Configuration - -

 Area ID        AuType          Stub?    Default-cost   Import-summaries?
 1.1.2.2        1=Simple-pass     No          N/A              N/A


                              - - Area Ranges - -

 Area ID        Address        Mask               Advertise?
 1.1.2.2        128.185.0.0    255.255.0.0        Yes


                           - - Interface configuration - -

 IP address        Sta     Area            Cost   Rtrns   TrnsDly   Pri   Hello   Dead
 16.24.11.251      Ena     1.1.2.2            1       5         1     1      10     40
```

(Continued on next page)

```
                        - - Authentication Keys - -

IP Address         AuType                   Key (Hex/Ascii)
128.185.138.21                              0x617574685F6B6579 “Auth_key”


                       - - Virtual link configuration - -

Virtual endpoint   Sta     Transit Area            Rtrns   TrnsDly   Hello   Dead
128.185.138.21     Dis     0.0.0.1                    10         5      30    180


                          - - NBMA configuration - -

                     Interface Addr          Poll Interval
                     128.185.138.19          120


                           - - Neighbor configuration - -

                Neighbor Addr           Interface Address        DR eligible?
                128.185.138.19          128.185.138.21           yes
                128.185.138.17          128.185.138.21           no
                138.185.139.19          128.185.139.21           no
```

Monitoring OSPF

This section describes tasks you can perform to monitor your router's OSPF protocol.

To access the Monitor process, perform the following steps:

Step	Action
1	At the Main prompt (`Main>`) enter: **monitor**
2	Press Return. The Monitor prompt (`Monitor>`) is displayed.
3	If the prompt is not displayed, press Return a second time.
4	At the `Monitor>` prompt, enter: **protocol ospf** The `OSPF>` prompt is displayed.

From the `OSPF>` prompt, you can perform specific tasks to determine OSPF parameters, routes, advertisements, database summaries, routing tables, and interface addresses.

Monitoring OSPF Advertisements

You can display the contents of a link state advertisement contained in the OSPF database. For a summary of the router's advertisements, use the **database** command.

To display OSPF advertisements, perform the following steps:

Step	Action
1	At the OSPF> prompt, enter: **ad**vertisement *ls-type link-state-id advertising-router area-id*
2	Press Return. The router's link state advertisements are displayed.

A link state advertisement is defined by its link state type, link state ID, and its advertising router. There is a separate link state database for each OSPF area. Providing an area ID on the command line tells the software which database you want to search. Listed below are different kinds of advertisements that depend on the value given for link state type:

- **Router links** — Contain descriptions of a single router's interface.
- **Network links** — Contain the list of routers attached to a particular interface.
- **Summary nets** — Contain descriptions of a single inter-area route.
- **Summary AS boundary routers** — Contain descriptions of the route to an AS boundary router in another area.
- **AS external nets** — Contain descriptions of a single route.

NOTE

Link state IDs, advertising routers (specified by their router IDs), and area IDs take the same format as IP addresses. For example, the backbone area can be entered as 0.0.0.0.

Example of Router Links Advertisement

The following example shows an expansion of a router links advertisement. The router's ID is 128.185.184.11. It is an AS boundary router and has three interfaces to the backbone area (all of cost 1). Multicast routing was enabled. Detailed field descriptions are provided with the example shown.

This command was also enhanced in two ways. First, when displaying router-LSAs and network-LSAs, the reverse cost of each router-to-router link and router-to-transit-network link is displayed, as well as the previously displayed forward cost. This is done because routing of multicast datagrams whose source lies in different areas/ASs is based on reverse cost instead of forward cost. In those cases where there is no reverse link (which means that the link is never used by the Dijkstra), the reverse cost is shown as 1-way.

In addition, the LSA's OSPF options are displayed in the same manner as they were displayed in the detailed OSPF neighbor command.

New group-membership-LSAs can also be displayed. An example follows. The ''LS destination" of each group-membership-LSA is a group address. A router originates a group-membership-LSA for each group that has members on one or more of the router's attached networks. The group-membership-LSA for the group lists those attached transit networks having group members (the type ''2" vertices), and when there are members belonging to one or more attached stub networks, or if the router itself is a member of the multicast group, a type ''1" vertex whose ID is the router's OSPF router ID is included.

```
OSPF> advertisement 1 128.185.184.11 0.0.0.0
LS age:     173
LS options:  E
 LS type:     1
 LS destination (ID): 128.185.184.11
 LS originator:   128.185.184.11
 LS sequence no:  0x80000047
LS checksum:     0x122
 LS length:       60
Router type:  ASBR,W
# router ifcs:   3
               Link ID:          128.185.177.31
               Link Data:        128.185.177.11
                Interface type:   2
                   No. of metrics: 0
                   TOS 0 metric:   3 (0)
               Link ID:          128.185.142.40
               Link Data:        128.185.142.11
```

```
Interface type:   2
     No. of metrics: 0
    TOS 0 metric:   4 (0)
Link ID:          128.185.184.0
Link Data:        255.255.255.0
Interface type:   3
    No. of metrics: 0
    TOS 0 metric:   1
```

Field	Description
`LS age`	Indicates the age of the advertisement in seconds. This field is common to all advertisements.
`LS options`	Indicates the optional OSPF capabilities supported by the piece of the routing domain described by the advertisement, denoted by any combination of: **E** — processes type 5 externals **T** — can route based on TOS **MC** — can forward IP multicast datagrams This field is common to all advertisements.
`LS type`	Classifies the advertisement and dictates contents: 1 (router links advertisement), 2 (network link advertisement), 3 (summary link advertisement), 4 (summary ASBR advertisement), and 5 (AS external link).
`LS destination`	Identifies what is being described by the advertisement. Depends on the advertisement type. For router links and ASBR summaries, it is the OSPF router ID. For network links, it is the IP address of the network's designated router. For summary links and AS external links, it is a network/subnet number. For group-membership advertisements, it is a particular multicast group. This field is common to all advertisements.
`LS originator`	OSPF router ID of the originating router. This field is common to all advertisements.
`LS sequence number`	Distinguishes separate instances of the same advertisement. A signed 32-bit integer, starting at 0x80000001, and increments by 1each time the advertisement is updated. This field is common to all advertisements.
`LS checksum`	A checksum of advertisement contents, used to detect data corruption. This field is common to all advertisements.

`LS length`	The size of the advertisement in bytes. This field is common to all advertisements.
`Router type`	Indicates the level of functionality of the router. ASBR is an AS boundary router, ABR is an area border router, and W indicates that the router is a wildcard multicast receiver. This field is used only in router link advertisements.
`# Router ifcs`	The number of router interface described in the advertisement. This field is used only in router link advertisements.
`Link ID`	Indicates what the interface connects to. Depends on interface type. For interfaces to routers (point-to-point links), the link ID is the neighbor's router ID. For interfaces to transit networks, it is the IP address of the network designated router. For interfaces to stub networks, it is the network's network/subnet number. Each link in the router advertisement is described by the Link ID, Link Data, and Interface type fields.
`Link Data`	Indicates 4 bytes of extra information concerning the link. It is either the IP address of the interface (for interfaces to point-to-point networks and transit networks), or the subnet mask (for interfaces to stub networks). Each link in the router advertisement is described by the Link ID, Link Data, and Interface type fields.
`Interface type`	One of the following: 1 (point-to-point connection to another router), 2 (connection to transit network), 3 (connection to stub network), or 4 (virtual link). Each link in the router advertisement is described by the Link ID, Link Data, and Interface type fields.
`No of metrics`	The number of nonzero TOS values for the metrics are provided for this interface.
`TOS 0 metric`	Cost of the interface. The reverse cost of the link is given in parentheses (derived from another advertisement). If there is no reverse link, 1-way is displayed. Each link can also be assigned a separate cost for each IP Type of Service (TOS). This is described by the No. of metrics and TOS 0 metric fields (the router currently does not route based on TOS, and only looks at the TOS 0 cost).

Monitoring OSPF Areas

You can display the statistics and parameters for all OSPF areas attached to the router, by performing the following steps:

Step	Action
1	At the `OSPF>` prompt, enter: **area**
2	Press Return. The router's attached areas are displayed.

Example

In the following example, the router attaches to a single area (the backbone area). A simple password scheme is being used for the area's authentication. The router has three interfaces attaching to the area, and has found four transit networks, seven routers, and no area border routers when doing the SPF tree calculation for the backbone.

```
OSPF> area

 Area ID        Authentication          #ifcs      #nets      #rtrs      #brdrs

 0.0.0.0           Simple-pass            3          4          7          0
```

Field	Description
`#ifcs`	Indicates the number of router interfaces attached to the particular area. These interfaces are not necessarily functional.
`#nets`	Indicates the number of transit networks that were found while doing the SPF tree calculation for this area.
`#rtrs`	Indicates the number of routers that were found when doing the SPF tree calculation for this area.
`#brdrs`	Indicates the number of area border routers that were found when doing the SPF tree calculation for this area.

Monitoring AS External Advertisements

You can display the AS external advertisements belonging to the OSPF routing domain. One line is printed for each advertisement. Each advertisement is defined by the following three parameters: its link state type (always 5 for AS external advertisements), its link state ID (called the LS destination), and the advertising router (called the LS originator).

To display AS external advertisements, perform the following steps:

Step	Action
1	At the OSPF> prompt, enter: **as-external**
2	Press Return. The router's AS external advertisements are displayed.

Example

```
OSPF> as-external

  Type    LS destination   LS originator    Seqno        Age     Xsum

    5     0.0.0.0          128.185.123.22   0x80000084   430     0x41C7

    5     128.185.131.0    128.185.123.22   0x80000080   450     0x71DC

    5     129.9.0.0        128.185.126.24   0x80000080   676     0x324A

    5     134.216.0.0      128.185.123.22   0x80000082   451     0x505A

    5     192.26.100.0     128.185.126.24   0x80000080   21      0xDEE8


               # advertisements:    133

                Checksum total:       0x43CC41
```

Field	Description
`Type`	Always 5 for AS external advertisements.
`LS destination`	Indicates an IP network/subnet number. These network numbers belong to other autonomous systems.
`LS originator`	Advertising router.
`Seqno, Age, Xsum`	It is possible for several instances of an advertisement to be present in the OSPF routing domain at any one time. However, only the most recent instance is kept in the OSPF link state database (and printed by this command). The LS sequence number (Seqno), LS age (Age), and LS checksum fields (Xsum) are compared to see which instance is most recent. The LS age field is expressed in seconds. Its maximum value is 3600.
`# advertisements`	Indicates the total number of AS external advertisements.
`Checksum total`	The 32-bit sum (carries discarded) of the individual advertisement's LS checksum fields. This information can be used to quickly determine whether two OSPF routers have synchronized databases.

Monitoring OSPF Databases

You can display a description of the contents of a particular OSPF area's link state database. AS external advertisements are omitted from the display. A single line is printed for each advertisement. Each advertisement is defined by the following three parameters: its link state type (called Type), its link state ID (called the LS destination), and the advertising router (called the LS originator).

To display a particular database, perform the following steps:

Step	Action
1	At the `OSPF>` prompt, enter: **database** ***area-id***
2	Press Return. The area's link state database is displayed.

Example

```
OSPF> database 0.0.0.0

  Type    LS destination   LS originator    Seqno        Age     Xsum

    1     128.185.123.22   128.185.123.22   0x80000084   442     0xCE2D

    1     128.185.136.39   128.185.126.39   0x80000082   469     0x5045

    2     128.285.125.40   128.185.129.40   0x80000049   457     0xA31

    2     128.185.129.40   128.185.129.40   0x80000001   1623    0x12C9

    6     224.0.2.6        128.185.184.11   0x80000003   376     0x2250

                # advertisements:    14

                Checksum total:       0x4BBC2
```

Field	Description
`Type`	Separate LS types are numerically displayed: type 1 (router links advertisements), type 2 (network links advertisements), type 3 (network summaries), type 4 (AS boundary router summaries), and type 6 (group-membership-LSAs).
`LS destination`	Indicates what is being described by the advertisement.
`LS originator`	Advertising router.
`Seqno, Age, Xsum`	It is possible for several instances of an advertisement to be present in the OSPF routing domain at any one time. However, only the most recent instance is kept in the OSPF link state database (and printed by this command). The LS sequence number (Seqno), LS age (Age) and LS checksum fields (Xsum) are compared to see which instance is most recent. The LS age field is expressed in seconds. Its maximum value is 3600.
`# advertisements`	Indicates the total number of advertisements in the area database.
`Checksum total`	The 32-bit sum (carries discarded) of the individual advertisement's LS checksum fields. This information can be used to quickly determine whether two OSPF routers have synchronized databases.

NOTE

When comparing multicast-capable to nonmulticast routers, the above database checksum (and also # advertisements) will not necessarily match, because nonmulticast routers do not handle or store group-membership LSAs.

Monitoring OSPF Dump Routing Tables

You can display all the routes that were calculated by OSPF and are now present in the routing table. Its output is similar in format to the IP console's dump routing tables command.

To display the routing tables, perform the following steps:

Step	Action
1	At the OSPF> prompt, enter: **dump**
2	Press Return. The OSPF routing tables are displayed.

Example

```
OSPF> dump
  Type   Dest net       Mask           Cost   Age   Next hop (s)
  Sbnt   16.0.0.0       FF000000         1     0    None
   SPF   16.24.8.0      FFFFFF00         2     2    20.24.12.230
  SPF*   20.24.12.0     FFFFFF00         1     1    Eth/1
  Sbnt   21.0.0.0       FF000000         1     0    None
  Dir*   21.24.166.0    FFFFFF00         1     0    Eth/5

Routing table size: 76 nets (55296 bytes), 5 nets known
```

Field	Description
`Type`	Indicates destination type. Net indicates that the destination is a network. All other destinations are covered by the OSPF routers command. • **Sbnt** — Indicates that the network is subnetted; such an entry is a placeholder only. • **Dir** — Indicates a directly connected network or subnet. • **RIP** — Indicates the route was learned through the RIP protocol. • **Del** — Indicates the route was deleted. • **Stat** — Indicates a statically configured route. • **SPF** — Indicates that the route is an OSPF intra-area route. • **SPIA** — Indicates that it is an OSPF inter-area routes. • **SPE1, SPE2** — Indicates OSPF external routes (types 1 and 2 respectively). • **Rnge** — Indicates a route type that is an active OSPF area address range and is not used in forwarding packets.
`Dest net`	Destination host or network.
`Mast`	Displays the entry's subnet mask.
`Cost Age`	Displays the route cost.
`Next hop (s)`	Address of the next router on the path toward the destination host. A number in parentheses at the end of the column indicates the number of equal-cost routes to the destination. The first hops belonging to these routes can be displayed with the IP console's route command.

Monitoring OSPF Interface Statistics and Parameters

You can display statistics and parameters related to OSPF interfaces. If no arguments are given, a single line is printed summarizing each interface. If an interface's IP address is given, detailed statistics for that interface are displayed.

To display OSPF interface statistics and parameters, perform the following steps:

Step	Action
1	At the `OSPF>` prompt, enter: **interface** or **interface** ***ip-address***
2	Press Return. The OSPF interface statistics and parameters are displayed.

Example

```
OSPF> interface

 Ifc Address     Phys    assoc. Area   Type     State   #nbrs   #adjs

 16.24.8.251     Eth/1   0.0.0.0       Brdcst    64       1         1

 16.24.11.251    Eth/1   0.0.0.0       Brdcst    64       1         1

 17.1.1.251      Eth/2   0.0.0.0       Brdcst    64       0         0

 25.24.13.251    Eth/3   0.0.0.0       Brdcst    64       0         0

 18.1.1.251      Eth41   0.0.0.0       Brdcst    64       0         0

 16.24.10.251    Eth/0   0.0.0.0       Brdcst    64       0         0

 135.24.10.251   Eth/0   0.0.0.0       Brdcst    64       0         0



 Ifc Address      assoc. Area   Type     State   #nbrs   #adjs

 128.185.123.22   0.0.0.0       Brdcst    64       0         0

 128.185.124.22   0.0.0.0       Brdcst    64       0         0

 128.185.125.22   0.0.0.0       Brdcst    64       6         2
```

Field	Description
`Ifc Address`	Interface IP address.
`Phys`	The physical interface.
`Assoc Area`	Attached area ID.
`Type`	Can be either Brdcst (broadcast, for example, an Ethernet interface), P-P (a point-to-point network, for example, a synchronous serial line), Multi (nonbroadcast, multi-access, for example, an X.25 connection), or VLink (an OSPF virtual link).
`State`	Can be one of the following: 1 (down), 2 (looped back), 4 (waiting), 8 (point-to-point), 16 (DR other), 32 (backup DR), or 64 (designated router).
`#nbrs`	Number of neighbors. This is the number of routers whose hellos were received, plus those that were configured.
`#adjs`	Number of adjacencies. This is the number of neighbors in state Exchange or greater. These are the neighbors with whom the router has synchronized or is in the process of synchronization.

Monitoring OSPF Neighbors

You can display statistics and parameters related to OSPF neighbors. If no arguments are given, a single line is printed summarizing each neighbor. If a neighbor's IP address is given, detailed statistics for that neighbor are displayed.

To display OSPF neighbor statistics and parameters, perform the following steps:

Step	Action
1	At the `OSPF>` prompt, enter: **neighbor** or **neighbor** ***ip-address***
2	Press Return. OSPF neighbor statistics are displayed.

Example

```
OSPF> neighbor
 Neighbor addr      Neighbor ID       State   LSrxl    DBsum    LSreq    Ifc
  128.185.125.39   128.185.136.39    128        0        0        0      Eth/1
```

Field	Description
`Neighbor addr`	Displays the neighbor's address.
`Neighbor ID`	Displays the neighbor's OSPF router ID.
`Neighbor State`	Can be one of the following: 1 (Down), 2 (Attempt), 4 (Init), 8 (2-Way), 16 (ExStart), 32 (Exchange), 64 (Loading), or 128 (Full).
`LSrxl`	Displays the size of the current link state retransmission list for this neighbor.
`DBsum`	Displays the size of the database summary list waiting to be sent to the neighbor.
`LSreq`	Displays the number of more recent advertisements that are being requested from the neighbor.
`Ifc`	Displays the interface shared by the router and the neighbor.

Monitoring OSPF Router Routes

You can display all router routes that were calculated by OSPF and are now present in the routing table. With the dump routing tables command, the Net field indicates that the destination is a network. The routers command covers all other destinations.

To display OSPF router routes, perform the following steps:

Step	Action
1	At the `OSPF>` prompt, enter: **routers**
2	Press Return. OSPF multicast routing statistics are displayed.

Example

```
OSPF> routers
 DType      RType      Destination    Area          Cost        Next hop(s)


 BR         SPF        20.24.12.230    0.0.0.0      1           20.24.12.230
 Fadd       SPF        20.24.12.230    0.0.0.0      1           0.0.0.2
 BR         SPF        16.24.8.251     0.0.0.0      2           20.24.12.230
 ASBR       SPIA       19.24.9.252     0.0.0.0      3           20.24.12.230
```

Field	Description
DType	Indicates destination type. Net indicates that the destination is a network, ASBR indicates that the destination is an AS boundary router, and BR indicates that the destination is an area border router, and Fadd indicates a forwarding address (for external routes).
RType	Indicates route type and how the route was derived. SPF indicates that the route is an intra-area route (comes from the Dijkstra calculation); SPIA indicates that it is an inter-area route (comes from considering summary link advertisements).
Destination	Destination router's OSPF ID. For type D entries, one of the router's IP addresses is displayed (which corresponds to a router in another AS).
Area	Always displayed as 0.0.0.0.
Cost	Displays the route cost.
Next hop	Address of the next router on the path toward the destination host. A number in parentheses at the end of the column indicates the number of equal-cost routes to the destination.

Monitoring OSPF Link State Advertisement Size

You can display the number of link state advertisements (LSA) currently in the link state database, categorized by type.

To display the number of LSAs in the database, perform the following steps:

Step	Action
1	At the OSPF> prompt, enter: **size**
2	Press Return. OSPF LSAs are displayed.

Example

OSPF> **size**

```
# Router-LSAs:                7
# Network-LSAs:               6
# Summary-LSAs:               14
# Summary Router-LSAs:        2
# AS External-LSAs:           44
# Group-membership-LSAs:      21

# Intra-area routes:          2
# Inter-area routes:          1
# Type 1 external routes      3
# Type 2 external routes      1
```

Monitoring OSPF Statistics

You can display statistics generated by the OSPF routing protocol. The statistics indicate how well the implementation is performing, including its memory and network utilization. Many of the fields displayed are confirmation of the OSPF configuration.

To display OSPF statistics, perform the following steps:

Step	Action
1	At the `OSPF>` prompt, enter: **statistics**
2	Press Return. OSPF statistics are displayed.

Example

`OSPF>` **statistics**

```
S/W version:              2.0
OSPF Router ID:           128.185.184.11
External comparison:      Type 2
AS boundary capability:   Yes
Import external routes:   EGP RIP STA DIR SUB
Orig. default route:       No (0,0.0.0.0)
Default route cost:        (1, Type 2)
Default forward. addr:     0.0.0.0
```

(continued on next page)

```
Attached areas:                  2     Estimated # external routes:     300
Estimated # OSPF routers:      100     Estimated heap usage:          76000
OSPF packets rcvd:           60822     OSPF packets rcvd w/ errs:     28305
Transit nodes allocated:      1728     Transit nodes freed:            1715
LS adv. allocated:            7394     LS adv. freed:                  7313
Queue headers alloc:           224     Queue headers avail:             224

# Dijkstra runs:              391     Incremental summ. updates:         0
Incremental VL updates:          0     Buffer alloc failures:             0
Multicast pkts sent:         49487     Unicast pkts sent:               557
LS adv. aged out:                0     LS adv. flushed:                 521
Incremental ext. updates:        0
```

Field	Description
`S/W version`	Displays the current OSPF software revision level.
`OSPF Router ID`	Displays the router's OSPF ID.
`External comparison`	Displays the external route type used by the router when importing external routes.
`AS boundary capability`	Displays whether external routes are imported.
`Import external routes`	Displays which external routes are imported.
`Orig default route`	Displays whether the router will advertise an OSPF default route. If the value is Yes and a nonzero number is displayed in parentheses, then a default route is advertised only when a route to the network exists.
`Default route cost`	Displays the cost and type of the default route (if advertised).
`Default forward addr`	Displays the forwarding address specified in the default route (if advertised).
`Attached areas`	Indicates the number of areas that the router has active interfaces to.

`Estimated heap usage`	Rough indication of the size of the OSPF link state database (in bytes).
`Transit nodes`	Allocated to store router links and network links advertisements.
`LS adv`	Allocated to store summary link and AS external link advertisements.
`Queue headers`	Form lists of link state advertisements. These lists are used in the flooding and database exchange processes; if the number of queue headers allocated is not equal to the number freed, database synchronization with some neighbor is in progress.
`# Dijkstra runs`	Indicates how many times the OSPF routing table was calculated from scratch.
`Incremental summ updates, Incremental VL updates`	Indicate that new summary link advertisements caused the routing table to be partially rebuilt.
`Buffer alloc failures`	Indicate buffer allocation failures. The OSPF system recovers from temporary lack of packet buffers.
`Multicast pkts sent`	Covers OSPF hello packets and packets sent during the flooding procedure.
`Unicast pkts sent`	Covers OSPF packet retransmissions and the database exchange procedure.
`LS adv. aged out`	Counts the number of advertisements that have hit 60 minutes. Link state advertisements are aged out after 60 minutes. Usually, they are refreshed before this time.
`LS adv. flushed`	Indicates number of advertisements removed (and not replaced) from the link state database.
`Incremental ext. updates`	Displays the number of changes to external destinations that are incrementally installed in the routing table.

Monitoring OSPF Traceroute Addresses

You can display the entire OSPF routing path to a given destination, hop by hop using traceroute. Traceroute sends out three probes and prints the IP address of the responder, with the round-trip time associated with the response.

To display OSPF routing paths, perform the following steps:

Step	Action
1	At the `OSPF>` prompt, enter: **traceroute** ***ip-destination-address***
2	Press Return. The OSPF destination area address and packet size is displayed.

Example

```
OSPF> traceroute 128.185.142.239
TRACEROUTE 128.185.124.110:    56 data bytes
 1  128.185.142.7  16 ms  0  ms  0  ms
 2  128.185.123.22  16 ms  0  ms  16  ms
 3  * * *
 4  * * *
 5  128.185.124.110  16 ms  !  0  ms  !  0  ms
```

Field	Description
`Traceroute`	Displays the destination area address and the size of the packet being sent to that address.
`1`	Displays the first trace showing the destination's NSAP and the amount of time it took the packet to arrive at the destination. The packet is traced three times.
`Destination unreachable`	Indicates that no route to the destination is available.
`3 * * *` `4 * * *`	Indicates that the router is expecting some form of response from the destination, but the destination is not responding.

The traceroute is done whenever the destination is reached, an ICMP Destination Unreachable is received, or the path length reaches 32 router hops.

When a probe receives an unexpected result, several indications can be displayed. They are:

Probe Result	Description
!N	Indicates that an ICMP Destination Unreachable (net unreachable) was received.
!H	Indicates that an ICMP Destination Unreachable (host unreachable) was received.
!P	Indicates that 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.
!	Indicates that the destination was reached, but the reply sent by the destination was received with a TTL of 1. This usually indicates an error in the destination, prevalent in some versions of UNIX. The destination is inserting the probe's TTL in its replies. This unfortunately leads to a number of lines consisting solely of asterisks before the destination is finally reached.

Chapter 6

Configuring and Monitoring the ARP Interface

Overview

Introduction

The Address Resolution Protocol (ARP) is a low-level protocol that dynamically maps network layer addresses to physical medium access control (MAC) addresses. ARP allows the source host or VNswitch router to find the MAC address of the destination host on the same network segment, given only the network layer address of the destination system.

For a detailed description of configuring ARP, see the *DIGITAL VNswitch 900 Series Switch Management* guide.

In This Chapter

This chapter discusses the following topics:

ARP Address Translation Overview

When a router translates a network layer address to a physical address, first the router accesses the ARP (translation) cache for the physical MAC address that corresponds to that network layer address. If the cache does not contain the physical MAC address, then the router broadcasts an ARP request to all hosts requesting a response from the host with the correct physical MAC address. The destination host with the correct physical MAC address sends a positive response to the router. The router sends the packet to the destination host and enters the physical MAC address into the translation cache for future use. Figure 6-1 illustrates how a router translates a network address to a physical address.

Figure 6-1: ARP Physical MAC Address Broadcast

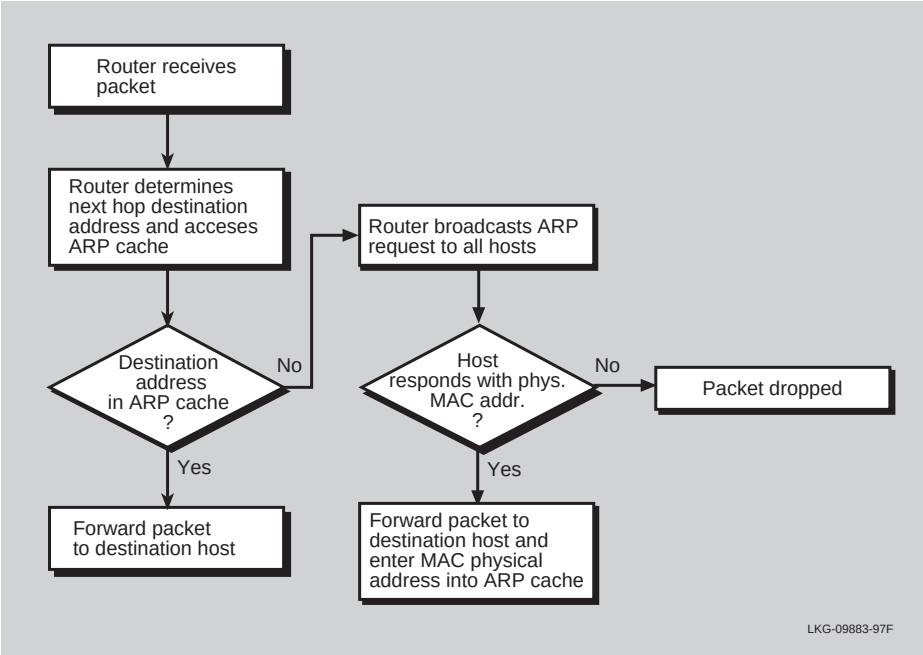

Accessing ARP

To access the ARP configuration commands, perform the following step:

Step	Action
1	At the `Config>` prompt enter: **arp**
2	Press Return. The `ARP Config>` prompt is displayed.

Configuring ARP Entries

Adding an ARP Entry

ARP cache contains a list of MAC addresses that map to network layer addresses. To add a MAC address translation entry to the ARP cache, perform the following steps:

Step	Action
1	At the `ARP Config>` prompt, enter: **add entry** ***interface-number ip-address mac-addr***
2	Press Return. The ARP updates its database with the new data.

Changing an ARP Entry

To change a MAC address translation entry in the ARP cache, at the `ARP Config>` prompt, enter:

change entry ***interface-number ip-address mac-addr***

The hardware address parameter (*mac-addr*) is the address of the node being changed.

Deleting an ARP Entry

To delete a MAC address translation entry from the ARP cache, at the `ARP Config>` prompt, enter:

delete entry ***interface-number ip-address***

The ARP deletes the entry from its database.

Configuring ARP Auto-Refresh

The auto-refresh function is the router's capability to send an ARP request based on the entry in the translation cache before the entry is deleted due to aging. An entry is deleted when its age reaches the refresh time value. The request is sent directly to the hardware address in the current translation instead of a broadcast. If auto-refresh is enabled, an ARP request is sent in this manner before the refresh timer is allowed to expire.

Enabling ARP Auto-Refresh

To enable ARP auto-refresh, perform the following steps:

Step	Action
1	At the `ARP Config>` prompt, enter: **enable auto-refresh**
2	Press Return.

Disabling Auto-Refresh

To disable ARP auto-refresh, at the `ARP Config>` prompt, enter:

disable auto-refresh

Listing ARP Configuration Data

Listing ARP

The contents of the router's ARP configuration are stored in SRAM. The list command displays the current timeout settings for the refresh timer, whether auto refresh is enabled/disabled and any configure entries.

To list the ARP configuration data, perform the following steps:

Step	Action
1	At the `ARP Config>` prompt, enter: **list all**
2	Press Return. The ARP's configuration contents are displayed.

Listing ARP Configuration

To list the configuration for the ARP refresh timer and whether auto refresh is enabled/disabled, at the `ARP Config>` prompt, enter:

list config

Listing ARP Entries

To list the ARP entries in nonvolatile memory, at the `ARP Config>` prompt, enter:

list entry

Setting the ARP Refresh Timer

Setting the ARP refresh timer changes the timeout value used for aging ARP entries. To change the timeout value for the refresh timer, enter the timeout value in minutes. A setting of zero (0) turns off (disables) aging ARP entries.

To set the ARP refresh timer, perform the following steps:

Step	Action
1	At the `ARP Config>` prompt, enter: **set refresh-timer *timeout-value***
2	Press Return. The ARP's refresh time is set.

Monitoring ARP

The following sections describe tasks you can perform to monitor your router's ARP protocol. To access the Monitor process, perform the following steps:

Step	Action
1	At the Main prompt (`Main>`) enter: **monitor**.
2	Press Return. The Monitor prompt (`Monitor>`) is displayed.
3	If the prompt is not displayed, press Return a second time.
4	At the `Monitor>` prompt, enter: **protocol arp**. The `ARP>` prompt is displayed.

From the `ARP Monitor>` prompt, you can perform specific tasks to determine ARP statistics, ARP cache, and ARP-configured networks and protocols.

Clearing the ARP Cache

You can flush the ARP cache for a given network interface. The clear command can be used to force the deletion of bad translations.

To clear a particular interface, enter the interface number as part of the command. To obtain the interface number, use the `Config>` **list devices** command.

To clear ARP cache, perform the following steps:

Step	Action
1	At the ARP> prompt, enter: **clear** ***ifc-#***
2	Press Return. The ARP cache for the specified interface is cleared.

Example

ARP> **clear 15**

Monitoring the ARP Cache

You can display the ARP cache for a given interface/protocol combination. To display the ARP cache for a particular interface, enter the interface number as part of the command. To obtain the interface number, use the `Config>` **list devices** command.

If the protocol is other than IP, HST for example, than the protocol number must also be given. This causes the console to display the hardware address to protocol address mappings stored in that database.

To display ARP cache, perform the following steps:

Step	Action
1	At the `ARP>` prompt, enter: **dump** ***ifc-# optional-protocol-#***
2	Press Return. The ARP cache for the specified interface is displayed.

Example

`ARP>` **dump 15**

```
Hardware Address       IP Address        Refresh

02-07-01-00-00-01      192.9.1.2            5
```

Monitoring ARP Interfaces

You can display the interfaces registered with ARP. The hardware command lists each ARP-registered interface, and displays each interface's hardware address space (Hardware AS) and local hardware address. Hardware addresses are displayed in hexadecimal.

To display ARP interfaces, perform the following steps:

Step	Action
1	At the ARP> prompt, enter: **hardware**
2	Press Return. ARP-registered interfaces are displayed.

Example

```
ARP> hardware
 Ifc   Interface             Hardware AS       Hardware Address
  0    VNbus/0               1                 00-00-F8-00-00-01
  1    Eth/1                 1                 02-07-01-00-00-02
  .
  .
  .
 15    VLAN/15               1                 00-00-F8-00-00-20
 16    VLAN/16               1                 00-00-F8-00-00-21
```

Monitoring ARP Protocols

You can display the protocol addresses registered with ARP. This command displays the interfaces, protocol name, protocol number, protocol address space (in hexadecimal), and local protocol addresses.

To display ARP protocol addresses, perform the following steps:

Step	Action
1	At the ARP> prompt, enter: **protocol**
2	Press Return. ARP-registered interfaces and protocol addresses are displayed.

Example

```
ARP> protocol

 Interfaces     Protocol       (num)   AS     Protocol Address (es)

 15 VLAN/15     IP              (0)    0800   192.9.1.1 18.124.0.11
```

Monitoring ARP Statistics

You can display a variety of statistics about the operation of the VNswitch router with ARP enabled.

To display ARP statistics, perform the following steps:

Step	Action
1	At the ARP> prompt, enter: **statistics**
2	Press Return. ARP statistics are displayed.

Example

```
ARP> statistics
ARP input packet overflows
    Ifc        Net        Count
     0         VNbus/0        0
     1         Eth/1          0
     .
     .
     .
    15         VLAN/15        0
    16         VLAN/16        0

ARP cache meters
 Ifc   Prot   Max   Cur   Cnt   Alloc   Refresh:tot   Failure   TMOs:Refresh
   0      4     1     1     1       1             0         0              0
   1      0     2     2    12      12             0         0              0
   2      4     1     1     1       1             0         0              0
```

Field	Description
`ARP input packet overflow`	Displays counters that represent the number of ARP packets discarded on input because the ARP layer was too busy. The counts shown are per interface.
`ARP cache meters`	Consists of a variety of meters on the operation of the ARP cache. The counts shown are all per protocol, per interface.
`Ifc`	Displays the interface numbers.
`Prot`	Displays the protocol numbers.
`Max`	Displays the all-time maximum length hash chain.
`Cur`	Displays the current maximum length hash chain.
`Cnt`	Displays the count of entries currently active.
`Alloc`	Displays the count of entries created.
`Refresh:Tot`	Displays the number of refresh requests sent for this network interface and protocol.
`Failure`	Displays the 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`	Displays the count of entries deleted due to a timeout of the refresh timer.

Chapter 7

Configuring and Monitoring the BGP4 Interface

Overview

Introduction

This chapter describes how to configure and monitor the Border Gateway Protocol (BGP) for a VNswitch logical interface.

BGP is not a routing protocol, but a reachability protocol. In essence, BGP routers selectively collect and advertise reachability information to and from BGP neighbors in their own and other autonomous systems (ASs). Reachability information consists of the sequences of AS numbers that form the paths to particular BGP speakers, and the list of IP addresses that can be reached via each advertised path. An *AS* is an administrative group of networks and routers that share reachability information using one or more Interior Gateway Protocols (IGPs), such as RIP or OSPF.

BGP commands are not dynamic. You must restart the module before any commands take effect.

In This Chapter

This chapter discusses the following topics:

Border Gateway Protocol Overview

BGP is an exterior gateway routing protocol used to exchange network reachability information among autonomous systems (ASs). An AS is essentially a collection of routers and end nodes that operate under a single administrative organization. Within each AS, routers and end nodes share routing information using an interior gateway protocol. The interior gateway protocol may be either RIP, OSPF, or Integrated IS-IS.

BGP was introduced in the Internet in the late 1980s to facilitate the loop-free exchange of routing information between autonomous systems. Based on Classless Inter-Domain Routing (CIDR), BGP has since evolved to support the *aggregation* and *reduction* of routing information.

In essence, CIDR is a strategy designed to address the following problems:

- Exhaustion of Class B address space
- Routing table growth

CIDR eliminates the concept of address classes, and provides a method for summarizing *n* different routes into single routes. This significantly reduces the amount of routing information that BGP routers must store and exchange.

NOTE

Digital Equipment Corporation supports only the latest version of BGP, BGP4, which is defined in RFC 1654. All references to BGP in this chapter and on the interface of DIGITAL routers are to BGP4, and do not apply to previous versions of BGP.

How BGP Works

Routers that run BGP are called *BGP speakers.* These routers function as servers with respect to their BGP neighbors (their clients). Each BGP router opens a passive TCP connection on port 179, and listens for incoming connections from neighbors at this well-known address. The router also opens active TCP connections to enabled BGP neighbors. This TCP connection enables BGP routers to share and update reachability information with neighbors in the same or other ASs. Connections between BGP speakers in the same AS are called *internal BGP (IBGP) connections*, while connections between BGP speakers in different ASs are *external BGP (EBGP) connections.* A single AS may have one or many BGP connections to outside ASs.

Figure 7-1 shows two ASs. The BGP speaker in AS1 attempts to establish a TCP connection with its neighbor in AS2. After this connection is established, the routers can share reachability information.

Figure 7-1: BGP Connections Between Two Autonomous Systems

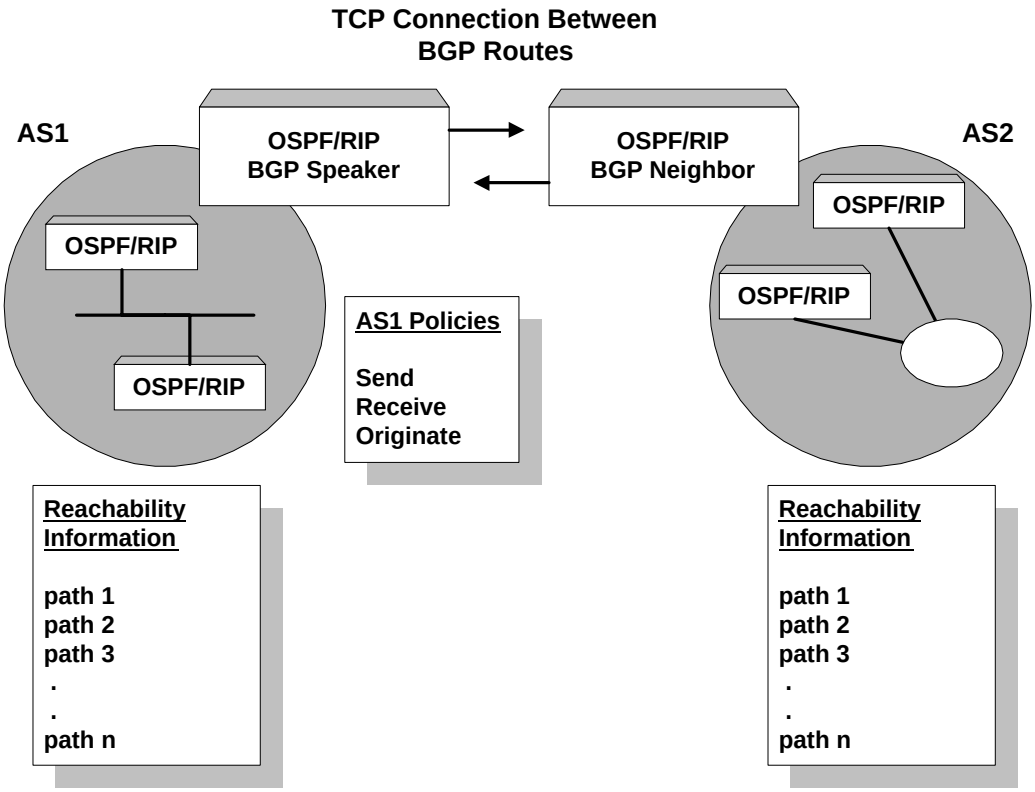

Once the BGP speaker in AS1 establishes a TCP connection with its BGP neighbor in AS2, the two routes can selectively exchange reachability information. The information each router sends or accepts is determined by policies defined for each router.

LKG-10601-97V

While the ASs shown in Figure 7-1 have only one BGP router, each may have multiple connections to other ASs. As an example of this, Figure 7-2 shows three interconnected ASs. AS1 has three BGP connections to outside ASs: one to AS2, one to AS3, and one to ASx. Similarly, AS3 has connections to AS1, AS2 and to ASy.

Figure 7-2: BGP Connections Between Three Autonomous Systems

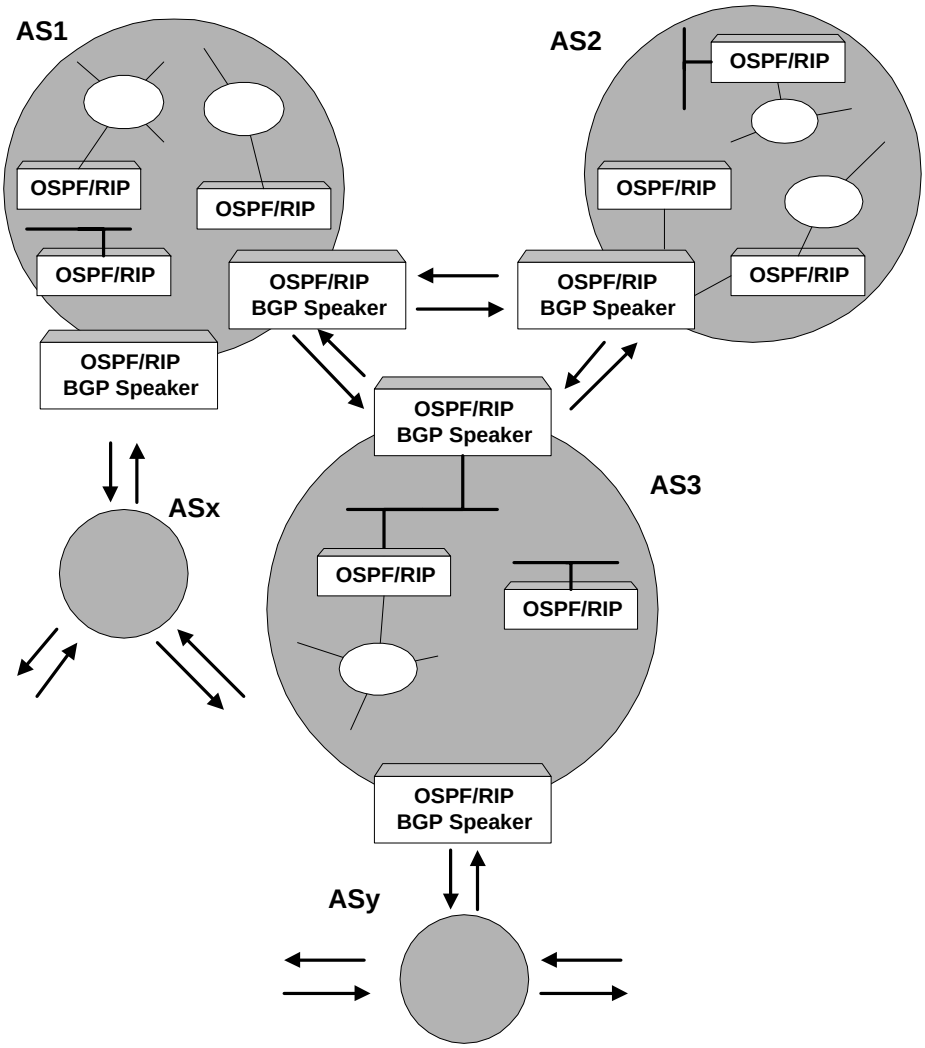

BGP relationships between three autonomous systems. Note that AS1 and AS3 have two BGP speakers.

LKG-10602-97V

Setting Up BGP

Setting up BGP involves four basic steps:

Step	Action
1	Access the BGP Configuration Prompt. Refer to Accessing BGP for details.
2	Configure a BGP speaker. Refer to Configuring a BGP Speaker for details.
3	Configure BGP Neighbors. Refer to Configuring Neighbors for details.
4	Configure Policies. Refer to Configuring Policies for details.

BGP Messages

BGP routers use four kinds of messages to communicate with their neighbors: *open, KeepAlive, update,* and *notification* messages.

Open Messages

Open messages are the first transmitted when a link to a BGP neighbor comes up and establishes a connection.

KeepAlive Messages

KeepAlive messages are used by BGP routers to inform one another that a particular connection is alive and working.

Update Messages

Update messages contain the interior routing table information. BGP speakers send update messages only when there is a change in their routing tables.

Notification Messages

Notification messages are sent whenever a BGP speaker detects a condition that forces it to terminate an existing connection. These messages are advertised before the connection is transmitted.

Accessing BGP

To access BGP, you must start a console session, access the Main process and then the Configuration process (as described in Chapter 1).

Once you start a console session, the Main process is automatically initiated, and the Main prompt (`Main>`) is displayed.

To access and configure BGP, perform the following steps:

Step	Action
1	At the Main prompt (`Main>`), enter: **config**
2	Press Return. The Config prompt `(Config>)` is displayed.
3	At the `Config>` prompt, enter: **bgp**
4	Press Return. The `BGP Config>` prompt is displayed.

Determining the BGP ID

BGP requires an ID to exchange network reachability information with its BGP neighbors. However, you do not have to configure a BGP address. BGP uses the first available address from one of the following addresses in the priority listed:

- Router ID
- Router ID selected by OSPF
- Router default IP address

Configuring a BGP Speaker

Enabling a BGP Speaker

Enabling BGP requires you to specify the BGP router's unique AS number. AS numbers are assigned by Stanford Research Institute Network Information Center. To enable a BGP speaker, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **enable bgp speaker *as#* tcp-segment-size**
2	Press Return. BGP is now enabled.

The AS number must be greater than zero, but less than 65536.

The TCP segment size must be greater than zero, but less than 65536. The default value is 1024. This number represents the maximum segment size BGP uses for passive TCP connections.

Example: `BGP Config>` **enable bgp speaker 165 2048**

Disabling a BGP Speaker

To disable a BGP speaker, at the `BGP Config>` prompt, enter:

disable bgp speaker

Configuring Neighbors

BGP neighbors are BGP routers with which a BGP speaker establishes a TCP connection. After enabling a BGP speaker, you must define its neighbors. BGP neighbors can be internal or external. Internal neighbors exist in the same AS, and do not need to have a direct connection to one another. External neighbors exist in different ASs. They must have a direct connection to one another.

Adding Neighbors

To define internal or external BGP neighbors, you must specify the IP address of the neighbor, and assign an AS number to the neighbor. Internal neighbors must have the same AS number as the BGP speaker.

To enable a BGP speaker, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **add neighbor neighbor-ip-address *as#* init-timer connect-timer hold-timer tcp-segment-size**
2	Press Return.

Adding a BGP neighbor automatically enables it, causing the BGP speaker to send out a connection request to the neighbor.

Example: `BGP Config>` **add neighbor**
`Neighbor address [0.0.0.0]?` **192.0.251.165**
`AS [0]?` **165**
```
Init timer [12]?
Connect timer [120]?
Hold timer [90]?
TCP segment size [1024]?
BGP Config>
```

Field	Description
Neighbor IP Address	Address of the neighbor you wish to peer with. It may be within your own autonomous system or in another autonomous system. If it is an external neighbor, both BGP speakers must share the same network. There is no such restriction for internal neighbors.
AS#	Your own autonomous system number for internal neighbor or neighbor's autonomous system number.
Init Timer	Specifies the amount of time the BGP speaker waits to initialize resources and reinitiate transport connection with the neighbor in case the speaker has previously transitioned to idle state due to an error. If the error persists, this timer increases exponentially. The default is 12 seconds.
Connect Timer	The amount of time the BGP speaker waits to reinitiate transport connection to its neighbor if the TCP connection fails while in either connect or active state. In the meantime, the BGP speaker continues to listen for any connection that may be initiated by its neighbor. The default is 120 seconds.
Hold Timer	The length of time the BGP speaker waits before assuming that the neighbor is unreachable. Both neighbors exchange the configured information using the open message and choose the smallest of the two timers as their negotiated hold timer value. The default is 90 seconds. Once neighbors have established BGP connection, they exchange KeepAlive messages at frequent intervals to ensure that the connection is still alive and the neighbors are reachable. The KeepAlive timer interval is calculated to be one third of the negotiated hold timer value. Hence, the hold timer value must be either zero or at least three seconds. Note that on switched lines, you may wish to have the hold timer value of zero to save bandwidth by not sending KeepAlive messages at frequent intervals.
TCP Segment Size	The maximum data size that may be exchanged on the TCP connection with a neighbor. This value is used for active TCP connection with the neighbor. It defaults to 1024, but can be set up to 65535.

Enabling Neighbors

To enable a BGP neighbor, at the `BGP Config>` prompt, enter:

enable neighbor neighbor-ip-address

Example: `BGP Config>` **enable neighbor 192.0.190.178**

Disabling Neighbors

To disable a BGP neighbor, at the `BGP Config>` prompt, enter:

disable neighbor neighbor-ip-address

Changing Neighbors

To change a BGP neighbor, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **change neighbor neighbor-ip-address as# init-timer connect-timer hold-timer tcp-segment-size**
2	Press Return.

The following example changes the value of the hold timer to zero for neighbor 192.0.251.165.

Example: `BGP Config>` **change neighbor 192.0.251.165**

```
AS [165]?
Init timer [12]?
Connect timer [60]?
Hold timer [12]? 0
TCP segment size [1024]?
```

Deleting Neighbors

To delete a BGP neighbor, at the `BGP Config>` prompt, enter:

delete neighbor neighbor-ip-address

Configuring Policies

The policies you establish determine which routes are imported and exported by the BGP speaker. Decisions about which reachability information to advertise (send) and which to accept (receive) are made on the basis of explicitly defined policy statements. DIGITAL BGP implementation supports three types of policy statements:

- **Originate Policy** — Enables you to select the interior gateway protocol (IGP) networks to export. These policies apply to routes to which the BGP speaker is directly connected; that is, routes that are local to the BGP speaker.
- **Receive Policy** — Enables you to select the route information to import from BGP peers.
- **Send Policy** — Enables you to select the route information to export to peers. Note that exportable route information can include information collected from neighboring ASs, as well as the routes that originate in the IGP.

Once a TCP connection is established, the BGP speaker shown in Figure 7-2 can send its entire routing table to its BGP neighbor in AS2. However, for security or other reasons, it may not be desirable to send reachability information on each network to AS2. Similarly, it may not be desirable for AS2 to receive reachability information on each network in AS1.

NOTE

Before you can send or receive information, you must establish policies.

Adding Policies

Adding Originate Policy

This command creates a policy that determines whether a specific address, or range of addresses, can be imported to the BGP speaker's routing table from the IGP routing table.

To add an originate policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **add originate-policy (exclusive/ inclusive) network-prefix network-mask address-match (Exact/Range) tag**
2	Press Return.

The following example includes all routes in the BGP speaker's IGP routing table to be advertised:

Example: `BGP Config>` **add originate-policy exclusive**
```
Network Prefix [0.0.0.0]?
Network Mask [0.0.0.0]?
Address Match (Exact/Range) [Exact]? range
Tag [0]?
```

Field	Description
`Exclusive`	Exclusive policies prevent route information from being included in the BGP speaker's routing table.
`Inclusive`	Inclusive policies ensure that specific routes are included in the BGP speaker's routing table.
`Network Prefix`	The network prefix for the addresses being affected.
`Address Match`	The address, or range of addresses, that is affected by the policy statement.
`Tag`	The value that was set for a particular AS. All tag values match that of the AS from which they were learned.

Adding Receive Policy

This command determines what routes are imported to the BGP speaker's routing table.

To add a receive policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **add receive-policy (exclusive/ inclusive) network-prefix network-mask address-match originating-as# adjacent-as# igp-metric (inclusive only)**
2	Press Return.

Example: `BGP Config>` **add receive-policy exclusive**
`Network Prefix [0.0.0.0]?` **10.0.0.0**
`Network Mask [0.0.0.0]?` **255.0.0.0**
`Address Match (Exact/Range) [Exact]?` **range**
`Originating AS# [0]?` **168**
`Adjacent AS# [0]?` **165**

Adding Send Policy

This command creates policies that determine which of the BGP speaker's learned routes are readvertised. These routes may be internal or external to the BGP speaker's AS.

To add a send policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **add send-policy (exclusive/ inclusive) network-prefix network-mask address-match tag adjacent-as#**
2	Press Return.

Example: `BGP Config>` **add send-policy exclusive**
`Network Prefix [0.0.0.0]?` **180.220.0.0**
`Network Mask [0.0.0.0]?` **255.255.0.0**
`Address Match (Exact/Range) [Exact]?` **range**
`Tag [0]?`
`Adjacent AS# [0]?` **25**

Changing Policies

Changing Originate Policy

This command alters an existing originate policy definition.

To change an originate policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **change originate-policy index# (exclusive/ inclusive) network-prefix network-mask address-match tag**
2	Press Return.

The following example alters the BGP speaker's originate policy. Rather than excluding networks with prefix 194.10.16.0 from the IGP routing table, the policy now includes all routes.

Example: `BGP Config>` **change originate-policy**
```
Enter index of originate-policy to be modified [1]?
Policy Type (Inclusive/Exclusive) [Exclusive]? inclusive
Network Prefix [194.10.16.0]? 0.0.0.0
Network Mask [255.255.240.0]? 0.0.0.0
Address Match (Exact/Range) [Range]?
Tag [0]?
```

Changing Receive Policy

This command alters an existing receive policy definition.

To change a receive policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **change receive-policy index# (exclusive/inclusive) network-prefix network-mask address-match originating-as# adjacent-as# igp-metric (inclusive only)**
2	Press Return.

The following example adds a restriction to the BGP speaker's receive policy. Rather than import route information from every BGP peer into its IGP routing table, it now prevents routes from AS 165 from being imported.

Example: `BGP Config>` **change receive-policy**

```
Enter index of receive-policy to be modified [1]?
Policy Type (Inclusive/Exclusive) [Inclusive]? exclusive
Network Prefix [0.0.0.0]?
Network Mask [0.0.0.0]?
Address Match (Exact/Range) [Range]?
Originating AS# [0]?
Adjacent AS# [0]? 165
```

Changing Send Policy

This command alters an existing send policy to one that is more inclusive or more exclusive.

To change a send policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **change send-policy index# (exclusive/ inclusive) network-prefix network-mask address-match tag adjacent-as#**
2	Press Return.

The following example adds a restriction to the BGP speaker's send policy. The restriction ensures that all routes in the address range 194.10.16.0 to 194.10.31.255 are excluded when advertising to autonomous system 165.

Example: `BGP Config>` **change send-policy**

```
Enter index of send-policy to be modified [1]?
Policy Type (Inclusive/Exclusive) [Inclusive]? exclusive
Network Prefix [0.0.0.0]? 194.10.16.0
Network Mask [0.0.0.0]? 255.255.240.0
Address Match (Exact/Range) [Range]?
Tag [0]?
Adjacent AS# [0]? 165
```

Deleting Policies

Deleting Originate Policy

This command deletes a specific originate policy. You must specify the index number associated with the policy.

To delete an originate policy, at the `BGP Config>` prompt, enter:

delete originate-policy index#

Example: `BGP Config>` **delete originate-policy 2**

Deleting Receive Policy

This command deletes a specific receive policy. You must specify the index number associated with the policy.

To delete a receive policy, at the `BGP Config>` prompt, enter:

delete receive-policy index#

Example: `BGP Config>` **delete receive-policy**

```
Enter index of receive-policy to be deleted [1]?
```

Deleting Send Policy

This command deletes a specific send policy. You must specify the index number associated with the policy.

To delete a send policy, at the `BGP Config>` prompt, enter:

delete send-policy index#

Example: `BGP Config>` **delete send-policy 4**

Sample Policy Definitions

This section provides a set of examples of some specific policies you can set up for a BGP speaker.

Originate Policy Examples

Including All Routes for Advertisement

The following example includes all routes in the BGP speaker's IGP routing table for advertisement. In this sense, you can view this command as the default originate policy statement for BGP.

Notice that the command specifies a range of addresses, rather than a single (exact) address.

Example: `BGP Config>` **add originate-policy inclusive**

```
Network Prefix [0.0.0.0]?
Network Mask [0.0.0.0]?
Address Match (Exact/Range) [Exact]? range
Tag [0]?
```

Excluding a Range of Routes

The following example excludes all routes in the range 194.10.16.0 to 194.10.31.255 from the BGP routing table, which in turn, prevents them from being advertised. It also specifies a range, but in this case the goal is to prevent the BGP speaker from advertising addresses in this range to its neighbors.

Example: `BGP Config>` **add originate-policy exclusive**

```
Network Prefix [0.0.0.0]? 194.10.16.0
Network Mask [0.0.0.0]? 255.255.240.0
Address Match (Exact/Range) [Exact]? range
Tag [0]?
```

Receive Policy Examples

Importing all Routes from All BGP Neighbors

The following example ensures that the BGP speaker imports all routes from all of its neighbors into its IGP routing table.

Example: `BGP Config>` **add receive-policy inclusive**

```
Network Prefix [0.0.0.0]?
Network Mask [0.0.0.0]?
Address Match (Exact/Range) [Exact]? range
Originating AS# [0]?
Adjacent AS# [0]?
IGP-metric [0]?
```

IGP-metric specifies the metric value with which the accepted routes are imported into the speaker's IGP routing table. You are prompted to enter a value for *IGP-metric* only when setting up a policy for route inclusion.

Blocking Specific Routes from a Transit AS

The following example prevents the BGP speaker from importing any routes originating at AS 168 from neighboring AS 165. You might use this command if you do not want the BGP speaker to receive any routes from AS 168 for security reasons.

Example: `BGP Config>` **add receive-policy exclusive**

```
Network Prefix [0.0.0.0]?
Network Mask [0.0.0.0]?
Address Match (Exact/Range) [Exact]? range
Originating AS# [0]? 168
Adjacent AS# [0]? 165
```

Send Policy Examples

Restricting Route Advertisement to a Specific AS

The following example restricts the BGP speaker. The speaker cannot advertise routes in the address range 143.116.0.0 to 143.116.255.255 that originate from AS 165 to autonomous system 168.

Example: `BGP Config>` **add send-policy exclusive**
`Network Prefix [0.0.0.0]?` **143.116.0.0**
`Network Mask [0.0.0.0]?` **255.255.0.0**
`Address Match (Exact/Range) [Exact]?` **range**
`Tag [0]?` **165**
`Adjacent AS# [0]?` **168**

Advertising All Known Routes

The following example ensures that the BGP speaker advertises all routes originated from its IGP, and all routes learned from its neighboring autonomous systems.

Example: `BGP Config>` **add send-policy inclusive**
```
Network Prefix [0.0.0.0]?
Network Mask [0.0.0.0]?
Address Match (Exact/Range) [Exact]? range
Tag [0]?
Adjacent AS# [0]?
```

Configuring Aggregate Addresses

Adding Aggregate Addresses

This command causes the BGP speaker to aggregate a block of addresses and advertise a single route to its BGP neighbors. You must specify the network prefix common to all the routes being aggregated and the prefix mask.

To add an aggregate address, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **add aggregate network-prefix network-mask**
2	Press Return.

The following example illustrates how to aggregate a block of addresses from 194.10.16.0 through 194.10.31.255.

Example: `BGP Config>` **add aggregate**
`Network Prefix [0.0.0.0]?` **194.10.16.0**
`Network Mask [0.0.0.0]?` **255.255.240.0**

When you add an aggregate definition, remember to define a policy to block the aggregated routes from being exported. If you do not, the router supports both the individual routes and the aggregate you have defined.

Changing Aggregate Addresses

To change an aggregate address, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **change aggregate index# network-prefix network-mask**
2	Press Return.

The following example changes the current aggregate (aggregate 1). The change causes aggregate 1 to use a different network prefix and mask to aggregate all routes in the address range from 128.185.0.0 to 128.185.255.255.

Example: `BGP Config>` **change aggregate 1**
`Network Prefix [128.185.0.0]?` **128.128.0.0**
`Network Mask [255.255.0.0]?` **255.192.0.0**

Deleting Aggregate Addresses

You must specify the index number of the aggregate you want to delete. The index number is equivalent to the AS number.

To delete an aggregate address, at the `BGP Config>` prompt, enter:

delete aggregate index#

Example: `BGP Config>` **delete aggregate 1**

Configuring No Receive Policy for Autonomous Systems

Adding No Receive Policy

This command excludes updates from a particular AS.

To add no receive policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **add no-receive as#**
2	Press Return.

Example: `BGP Config>` **add no-receive 178**

Deleting No Receive Policy

This command deletes the no-receive policy set up for a particular AS. You must specify the AS number.

To delete no receive policy, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **delete no-receive as#**
2	Press Return.

Example: `BGP Config>` **delete no-receive 168**

Clearing the BGP Configuration

This command erases the complete BGP configuration.

To erase the complete BGP configuration, perform the following steps:

Step	Action
1	At the `BGP Config>` prompt, enter: **clear**
2	Press Return.

Listing the BGP Configuration

This command displays various pieces of the IP configuration data, depending on the particular subcommand (**aggregate, all, bgp speaker, neighbor, no-receive, originate-policy, receive-policy** and **send-policy**) you invoke. The following example uses the **all** subcommand to display all the BGP configuration data.

Example: `BGP Config>` **list all**

```
            BGP Protocol:                Enabled
            AS:                          167
            TCP-Segment Size:            1024

 Neighbors and their AS:
                                  Init    Conn    Hold    TCPSEG
 Address           State     AS   Timer   Timer   Timer   Size
 128.185.250.168   ENABLED   168  12      60      12      1024
 192.0.251.165     ENABLED   165  12      60      12      1024

 Receive-Policies:
 Index   Type   Prefix    Mask      Match   OrgAS   AdjAS   IGPmetric
 1       INCL   0.0.0.0   0.0.0.0   Range   0       0       0

 Send-Policies:
 Index   Type   Prefix    Mask      Match   Tag     AdjAS
 1       INCL   0.0.0.0   0.0.0.0   Range   0       0

 Originate-Policies:
 Index   Type   Prefix        Mask            Match   Tag
 1       EXCL   194.10.16.0   255.255.240.0   Range   0

 Aggregation:
 Index   Prefix        Mask
 1       194.10.16.0   255.255.240.0

 AS-PATH with following ASs will be discarded:
 AS 178
 AS 165
```

Monitoring BGP

This section describes tasks you can perform to monitor your router's BGP protocol. To access the BGP Monitor process, perform the following steps:

Step	Action
1	At the Main prompt (`Main>`) enter: **monitor**
2	Press Return. The Monitor prompt (`Monitor>`) is displayed.
3	If the prompt is not displayed, press Return a second time.
4	At the `Monitor>` prompt, enter: **bgp**
5	Press Return. The `BGP>` prompt is displayed.

From the `BGP>` prompt, you can perform specific tasks to determine BGP destinations, neighbors, paths, and sizes.

Monitoring Destinations

This command dumps all BGP routing table entries, or display information on routes advertised to, or received from, specified BGP neighbor addresses (destinations).

Destinations

To display the BGP routing table entries, perform the following steps:

Step	Action
1	At the `BGP>` prompt, enter: **destinations**
2	Press Return.

Example: `BGP>` **destinations**

```
Network       Mask       NextHop           MED   AAG   AGRAS   ORG   AS-Path
128.185.0.0   FFFF0000   192.0.251.165     0     No    0       IGP
142.4.0.0     FFFF0000   192.0.190.178     0     No    0       IGP   seq[178]
143.116.0.0   FFFF0000   128.185.252.168   0     No    0       IGP   seq[168]
192.0.190.0   FFFFFF00   192.0.251.165     0     No    0       IGP
192.0.251.0   FFFFFF00   192.0.251.165     0     No    0       IGP
194.10.16.0   FFFFF000   192.0.251.167     0     No    167     IGP   seq[167]
```

Field	Description
`Network`	Indicates the IP addresses of the destinations in the routing table.
`Mask`	The address mask for each entry in the table.
`NextHop`	Indicates the address of the router to use as the forwarding address towards this destination.
`MED`	Specifies a multi-exit discriminator value, used to discriminate among multiple entry/exit points to the same AS.
`AAG`	Indicates whether the route is an aggregate or not. Values are `Yes` or `No`.

Field	Description
AGRAS	The number of the AS that aggregated the route.
ORG	Specifies the originator of this destination: either EGP, IGP, or Incomplete (originated by some other means not known). This field is blank if the *Network* destination is not being used for forwarding.
AS-PATH	Enumeration of ASs along the path. • **seq**: Sequence of ASs in order in the path • **set**: Set of ASs in the path This field is blank if the *Network* destination is not being used for forwarding.

Destinations Net Address

This command displays detailed information on the specified route or destination network. The command shows how a specific route was learned, the best path to a specific destination, the metric associated with the route, and other information.

To display specific information on a BGP route, perform the following steps:

Step	Action
1	At the BGP> prompt, enter: **destinations net-address**
2	Press Return.

Example: BGP> **destinations 142.4.0.0**

```
 Network       Mask       NextHop           MED   AAG   AGRAS   ORG   AS-Path

 142.4.0.0     FFFF0000   192.0.251.165     0     No    0       IGP   seq[165-178]

 Dest:142.4.0.0, Mask:FFFF0000, Age:180, Upd#:13,

 LastSent:0001:53:32 Eligible paths: 2
PathID: 8 - (Best Path)
       ASpath: seq[165-178]
       Origin: IGP, Pref: 507, LocalPref: 0
       Metric: 0, Weight: 0, MED: 0
       NextHop: 192.0.251.165, Neighbor: 192.0.251.165
       AtomicAggr: No
```

```
PathID: 21
       ASpath: seq[168-165-178]
       Origin: IGP, Pref: 505, LocalPref: 0
       Metric: 0, Weight: 0, MED: 0
       NextHop: 128.185.250.168, Neighbor: 128.185.250.168
       AtomicAggr: No
```

Field	Description
Network	Indicates the IP address of the specified destination.
Mask	The address mask for this entry.
NextHop	Indicates the address of the router to use as the forwarding address towards this destination.
MED	Specifies a multi-exit discriminator value, used to discriminate among multiple entry/exit points to the same AS. This field is blank if the *Network* destination is not being used for forwarding.
AAG	Indicates whether the route is an aggregate or not. Values are Yes or No. This field is blank if the *Network* destination is not being used for forwarding.
AGRAS	The number of the AS that aggregated the route. This field is blank if the *Network* destination is not being used for forwarding.
ORG	Specifies the originator of this destination: either EGP, IGP, or Incomplete (originated by some other means not known). This field is blank if the *Network* destination is not being used for forwarding.
AS-PATH	Enumeration of ASs along the path. • **seq**: Sequence of ASs in order in the path • **set**: Set of ASs in the path This field is blank if the *Network* destination is not being used for forwarding.
Dest	Indicates the IP address of the specified destination.
Mask	The address mask for this entry.
Age	Indicates the age of this entry in seconds.
Upd#	Indicates the sequence number of the last update message for this destination.
LastSent	Indicates the time that the last message was sent to this destination.
Eligible paths	Indicates the number of eligible paths to this destination.

Field	Description
`Path ID`	Indicates the unique identifier for each path.
`ASpath`	Enumeration of ASs along the path. • **seq:** Sequence of ASs in order in the path • **set:** Set of ASs in the path
`Origin`	Indicates the originator of the destination. This is either EGP, IGP, or Incomplete (originated by some other means not known).
`LocalPref`	Indicates the originating router's degree of preference for the destination.
`Metric`	Specifies the path metric with which the route is imported.
`Weight`	Specifies the path weight.
`MED`	Specifies a multi-exit discriminator value, used to discriminate among multiple entry/exit points to the same AS.
`NextHop`	Indicates the address of the router to use as the forwarding address for destinations reachable via the given path.
`AtomicAggr`	Indicates whether the router advertising the path has included the path in an atomic-aggregate.

Destinations Net Address Net Mask

This command is similar to the **destinations net address** command. This command is useful in cases where multiple network addresses have the same prefix and different masks. In such cases, specifying the network mask narrows the scope of the information presented.

To display specific information on a BGP route, perform the following steps:

Step	Action
1	At the `BGP>` prompt, enter: **destinations net-address net-mask**
2	Press Return.

Example: BGP> **destinations 194.10.16.0 255.255.240.0**

```
Dest:194.10.16.0, Mask:FFFFF000, Age:0, Upd#:3, LastSent:0002:00:00
Eligible paths: 1
PathID: 0 - (Best Path)
       ASpath:
       Origin: IGP, Pref: 0, LocalPref: 0
       Metric: 0, Weight: 0, MED: 0
       NextHop: 194.10.16.167, Neighbor: 194.10.16.167
       AtomicAggr: No, Aggregator AS167/194.10.16.167
```

Destinations Advertised To Net Address

This command lists all routes advertised to the specified BGP neighbor.

To display the list of routes advertised to the specified BGP neighbor, perform the following steps:

Step	Action
1	At the BGP> prompt, enter: **destinations advertised-to net-address**
2	Press Return.

Example: BGP> **destinations advertised-to 192.0.251.165**

```
Destinations advertised to BGP neighbor 192.0.251.165
 Network       Mask       NextHop           MED   AAG   AGRAS   ORG   AS-Path
 194.10.16.0   FFFFF000   194.10.16.167     0     No    167     IGP
 192.0.190.0   FFFFFF00   192.0.251.165     0     No    0       IGP   seq[165]
 142.4.0.0     FFFF0000   192.0.251.165     0     No    0       IGP   seq[165-178]
 143.116.0.0   FFFF0000   128.185.250.168   0     No    0       IGP   seq[168]
```

Destinations Received From Net Address

This command lists all routes received from the specified BGP neighbor.

To display the list of routes received from the specified BGP neighbor, perform the following steps:

Step	Action
1	At the `BGP>` prompt, enter: **destinations received-from net-address**
2	Press Return.

Example: `BGP>` **destinations received-from 128.185.250.167**

```
Destinations obtained from BGP neighbor 128.185.250.167
 Network       Mask       NextHop           MED   AAG   AGRAS   ORG   AS-Path

 194.10.16.0   FFFFF000   128.185.250.167   0     No    167     IGP   seq[167]

 192.0.190.0   FFFFFF00   128.185.250.167   0     No    0       IGP   seq[167-165]

 142.4.0.0     FFFF0000   128.185.250.167   0     No    0       IGP   seq[167-165-178]
```

Monitoring Neighbors

This command displays information on all active BGP neighbors.

To display active BGP neighbors, perform the following steps:

Step	Action
1	At the BGP> prompt, enter: **neighbors**
2	Press Return.

Example: BGP> **neighbors**

```
IP-Address        State         DAY-HH:MM:SS   BGP-ID           AS    Upd#

128.185.252.168   Established   000-00:48:52   128.185.142.168  168   16

192.0.190.178     Established   000-02:01:49   142.4.140.178    178   16

192.0.251.167     Established   000-02:01:45   194.10.16.167    167   16
```

This command displays detailed data on a particular BGP neighbor.

To display information on a particular BGP neighbor, perform the following steps:

Step	Action
1	At the BGP> prompt, enter: **neighbor ip-address**
2	Press Return.

Example: BGP> **neighbor 192.0.251.167**

```
Active Conn:Sprt:1026 Dprt:179 State:Established KeepAlive/Hold Time 4/12
Passve Conn:None
TCP connection errors: 0          TCP state transitions: 0

BGP Messages:     Sent   Received                    Sent   Received
Open:             1      1          Update:          11     11
Notification:     0      0          KeepAlive:       1828   1830
Total Messages:   1840   1842

Msg Header Errs:  Sent   Received                    Sent   Received
Conn sync err:    0      0          Bad msg length:  0      0
Bad msg type:     0      0
```

```
Open Msg Errs:    Sent   Received                    Sent   Received
Unsupp versions:  0      0          Unsupp auth code: 0      0
Bad peer AS ident:0      0          Auth failure:     0      0
Bad BGP ident:    0      0          Bad hold time:    0      0

Update Msg Errs:  Sent   Received                    Sent   Received
Bad attr list:    0      0          AS routing loop:  0      0
Bad wlkn attr:    0      0          Bad NEXT_HOP atr: 0      0
Mssng wlkn attr:  0      0          Optional atr err: 0      0
Attr flags err:   0      0          Bad netwrk field: 0      0
Attr length err:  0      0          Bad AS_PATH attr: 0      0
Bad ORIGIN attr:  0      0

Total Errors:     Sent   Received                    Sent   Received
Msg Header Errs:  0      0          Hold Timer Exprd: 0      0
Open Msg Errs:    0      0          FSM Errs:         0      0
Update Msg Errs:  0      0          Cease:            0      0
```

Monitoring Paths

This command displays the paths stored in the path description database.

To display paths, perform the following steps:

Step	Action
1	At the `BGP>` prompt, enter: **paths**
2	Press Return.

Example: `BGP>` **paths**

```
PathId       NextHop          MED   AAG   AGRAS   RefCnt   ORG   AS-Path
0            10.2.0.3         0     No    0       2        IGP
4            192.2.0.2        0     No    0       2        IGP   seq[2]
5            192.2.0.2        0     No    2       1        IGP   seq[2]
6            192.2.0.2        0     No    0       1        IGP   seq[2-1]
7            10.2.0.168       0     No    0       4        IGP
8            192.3.0.1        0     No    0       2        IGP   seq[1]
9            192.2.0.2        0     No    2       1        IGP   seq[2]
10           10.2.0.3         0     No    0       1        IGP
```

Field	Description
PathId	Path identifier.
NextHop	Specifies the address of the router to use as the forwarding address for the destinations that can be reached via the given path.
MED	Specifies the multi exit discriminator used to discriminate among multiple entry/exit points to the same AS.
AAG	Indicates whether the path was atomic-aggregated; that is, the router that is advertising the given path has selected a less specific route over the more specific one when presented with overlapping routes.
AGRAS	Indicates the AS number of the BGP speaker that aggregated the routes.
RefCnt	Indicates the number of path entities referring to the descriptor.
ORG	Specifies the originator of the advertised destinations in the given path: either EGP, IGP, or Incomplete (originated by some other means not known).
AS-Path	Enumeration of ASs along the path. • **seq**: Sequence of ASs in order in the path • **set**: Set of ASs in the path

Monitoring Sizes

This command displays the number of entries stored in the various databases.

To display the number of entries in the databases, perform the following steps:

Step	Action
1	At the `BGP>` prompt, enter: **sizes**
2	Press Return.

Example: `BGP>` **sizes**

```
 IP-Address       State         DAY-HH:MM:SS   BGP-ID           AS    Upd#

 128.185.252.168  Established   000-00:48:52   128.185.142.168  168   16

 192.0.190.178    Established   000-02:01:49   142.4.140.178    178   16

 192.0.251.167    Established   000-02:01:45   194.10.16.167    167   16

# Paths:                                11
# Path descriptors:                      7
Update sequence#:                       22
# Routing tbl entries (allocated):       6
# Current tbl entries (not imported):    0
# Current tbl entries (imported to IGP): 3
```

Field	Description	
IP-Address	Specifies the IP address of the BGP neighbor.	
State	Specifies the state of the connection. Possible states are:	
	Connect	Waiting for the TCP connection to the neighbor to be completed.
	Active	In the event of TCP connection failure, the state is changed to Active, and the attempt to acquire the neighbor continues.
	OpenSent	In this state open message was sent, and BGP waits for an open message from the neighbor.
	OpenConfirm	In this state a KeepAlive message was sent in response to the neighbor's open message, and waits for a KeepAlive notification from the neighbor.
	Established	A BGP connection was successfully established, and can now start to exchange update messages.
BGP-ID	Specifies the neighbor's BGP identification number.	
AS	Specifies the neighbor's AS number.	
Upd#	Specifies the sequence number of the last update message sent to the neighbor.	
# Paths	Total number of eligible paths for all the routes in the BGP routing table.	
# Path descriptors	Total number of path descriptors in the database used to hold common path information.	
Update sequence#	Indicates the current update sequence number.	
# Routing tbl entries (allocated)	Indicates the number of entries in BGP routing table.	
# Current tbl entries (not imported)	Indicates the number of BGP routes not imported into IGP.	
# Current tbl entries (imported to IGP)	Indicates the number of BGP routes imported into IGP.	

Appendix A

DIGITAL Trace Facility

Overview

Introduction

This appendix provides an overview of the DIGITAL Trace Facility (DTF) version 3.0 and describes what information can be traced over the interfaces in DIGITAL VNswitch routers having active tracepoints.

In This Appendix

The following topics are covered in this chapter:

DIGITAL Trace Facility

The VNswitch 900 router supports the DIGITAL Trace Facility (DTF) operating in TCP/IP networks. DTF is a host-based utility that traces packets as they pass through the protocol layers within a router and displays the decoded packet on the host or records the trace data in a file for later analysis. DTF includes facilities for:

- Filtering certain packet types at the source
- Filtering the output using regular expressions
- Performing validity check on packets
- Displaying the output in full, brief, or raw formats

The VNswitch 900 performs all packet forwarding and bridging in its fast-path processor (FP), with only terminating and multicast packets sent to the application processor (AP) for processing. There are no DTF tracepoints in the FP, which means that forwarded and bridged packets will not be traced. Router and bridge control packets (and all protocol control packets) are either terminating or are multicast and are processed by the AP and are traceable. By removing the forwarded packets from the traced data, you can capture more control packets on any one tracepoint in the VNswitch.

Tracepoints

The points within a router that can be traced are known as *tracepoints.* Each tracepoint defines a number of *events* (up to 64) and each packet that is traced through the tracepoint is marked with one of these events by the router.

Tracing occurs through tracepoints that are usually positioned in the transmit and receive routines of the protocol modules in the router. Each tracepoint has a name and state. When a packet passes through an active tracepoint, the packet is copied and queued to the DTF module in the VNswitch, which transmits it to the host where the DTF utility resides.

Each protocol module within the router may have zero or more tracepoints. For example, a tracepoint of *ETHERNET INTERFACE Eth/** traces all Ethernet interfaces whose names match the wildcard string *Eth/**. Table A-1 summarizes the DTF tracepoints for the VNswitch router.

Table A-1: Router Tracepoints

Tracepoint	Filters	Description
Ethernet interface *name*	tx, rx, aarptx, aarprx, apltx, aplrx, arptx, arprx, dntx, dnrx, iptx, iprx, ipxtx, ipxrx, ipv6tx, ipv6rx, moptx, moprx, ositx, osirx, xnstx, xnsrx	Traces all packets received by the AP on the Ethernet interface and all packets transmitted over that interface that originated from the router.
FDDI interface *name*	tx, rx, arptx, aarprx, apltx, aplrx, arptx, arprx, dntx, dnrx, iptx, iprx, ipxtx, ipxrx, ipv6tx, ipv6rx, moptx, moprx, ositx, osirx, xnstx, xnsrx	Traces all packets received by the AP on the FDDI interface and all packets transmitted over that interface that originated from the router.
VLAN interface *name*	tx, rx, aarptx, aarprx, apltx, aplrx, arptx, arprx, dntx, dnrx, iptx, iprx, ipxtx, ipxrx, ipv6tx, ipv6rx, moptx, moprx, ositx, osirx, xnstx, xnsrx	Traces all packets received by the AP on the VLAN interface and all packets transmitted over that interface that originated from the router.
IGMP interface *name*	tx, rx	Traces all IGMP packets transmitted and received by the IGMP module on the AP over the specified interface.
IP interface *name*	tx, rx, icmptx, icmprx, ospftx, ospfrx, tcptx, tcprx, udptx, udprx,	Traces all packets transmitted and received by the IP module on the AP over the specified interface. All IP packets received on the AP are traced by this tracepoint; however, not all IP packets transmitted by the AP pass through this tracepoint. In particular, most IP packets originated by the routing control protocols such as RIP and OSPF are sent directly to the datalink drivers and bypass this tracepoint. Packets originated by the UDP (other than RIP) and TCP based protocols do pass through this tracepoint.
OSPF interface *name*	tx, rx	Traces all packets transmitted and received by the OSPF module on the AP over the specified interface.
Event subsystem *subsystem-list*	trace, info, error, always	Traces all the ELS messages. The instance *name* is a list of ELS subsystems whose messages are traced. By default, all message types are traced and the filters can be used to restrict the trace to messages of a particular type. Note: Wildcards (*) cannot be used with this tracepoint.

Events

When DTF initially connects to the router, it optionally instructs the router to filter the packets passing through the specified tracepoint to trace the events that match the events specified in the events list file located on the DTF host system.

When you run DTF on a host system, it uses a transport protocol to connect to the router you want to trace and also sends the parameters to use for the trace session. Events at the activated tracepoints are transmitted back to the host system for analysis (either live display, or recorded in a trace file for later analysis).

Events can be either traced or blocked. By default, all specified events are traced unless otherwise blocked by prefixing the event name with the exclamation point (!) character in the events list file. The special event name, denoted with an asterisk (*), is used to mean all events. If a filter list (event list) is not specified, then DTF assumes all events are traced. If a filter list is specified, then DTF uses the filter list to block the events before processing the list.

Session Trace Buffer Parameters

Session trace buffer parameters are forwarded to the VNswitch by the host system at the start of the trace session. These parameters determine how much router resources are allocated to the trace session and how much of the data passing through the tracepoints is captured. The following is a list of session parameters sent to the VNswitch:

Parameter	Description
Buffer count	Specifies the number of trace buffers used to capture trace data during the session. The larger the number of trace buffers that are used results in less trace data loss.
Buffer size	Specifies the size (in bytes) of the trace buffers. This value determines the size of the data packet that can be traced.
Capture size	Specifies the number of bytes in the data packet that is copied into the trace buffer.

Trace Data Loss

Trace data loss occurs when there are not enough trace buffers available to trace the next data packet. To minimize the effect of trace data loss, use the following guidelines, which are listed in order of effectiveness:

- Use filters to reduce the amount of data being traced.
- Increasing the buffer count allows more trace data to be buffered within the router.
- Reducing the capture size increases the packets that can be contained in each trace buffer.
- Increasing the buffer size increases buffering available to the trace system.
- Record the results instead of displaying them.

Accessing DTF

DTF version 3.0 software is included with the clearVISN Router Configurator software. It is in the install-directory\tools\supported\dtf\ subdirectory.

The latest versions of the DTF documentation and installation kits for each host platform are available over the Internet, and can be downloaded from the following World Wide Web locations:

North America:	http://www.networks.digital.com
Europe:	http://www.networks.europe.digital.com
Asia Pacific:	http://www.networks.digital.com.au

Use the search feature to find the DTF Installation Kit.

Appendix B

Command Line Interface Quick Reference

Overview

Introduction

This appendix lists all the VNswitch router commands at the command line interface (CLI) level using a command tree format. The following tables list the **Config** and **Monitor** commands:

To Find These Commands	Refer To
ARP Config	Table B-1
ARP Monitor	Table B-2
IP Config	Table B-3
IP Monitor	Table B-4
OSPF Config	Table B-5
OSPF Monitor	Table B-6
RIP Config	Table B-7

Table B-1: ARP Config Commands

To Configure ARP	Use These Commands
Auto-Refresh	**disable auto-refresh** **enable auto-refresh**
Entry	**add entry** ***interface number protocol-type ip-address mac-addr*** **change entry** ***interface number protocol-type ip-address mac-addr*** **delete entry** ***interface number ip-address*** **list entry**
List	**list all** **list config**
Refresh Timer	**set refresh-timer** ***timeout value***

Table B-2: ARP Monitor Commands

To Monitor ARP	Use These Commands
Clear	**clear** ***ifc #***
Dump	**dump** ***ifc # optional protocol #***
Hardware	**hardware**
Protocol	**protocol**
Statistics	**statistics**

Table B-3: IP Config Commands

To Configure IP	Use These Commands
Access Control	**add access-control** ***type ip-source source-mask ip-dest dest-mask first-protocol last-protocol first-port last-port***
	delete access-control ***record-number***
	list access-controls
	move access-control ***from# to#***
	set access-control on
Address	**add address** ***interface-number ip-address address-mask***
	change address ***old ip-address new ip-address new subnet mask***
	delete address ***ip-address***
	list address
Enable Enhanced Proxy ARP	**disable enhanced-proxy-arp**
	enable enhanced-proxy-arp
	set enhanced-proxy-arp off
	set enhanced-proxy-arp on
Add Enhanced Proxy ARP	**add enhanced-proxy-arp subnet**
	delete enhanced-proxy-arp subnet
BootP Forwarding	**disable bootp-forwarding**
	enable bootp-forwarding
	list bootp
BootP Server	**add bootp-server** ***server-ip-address***
	delete bootp-server server-ip-address
	list bootp
Broadcast Forwarder	**add broadcast-forwarder udp** ***udp-port interface number destination ip-address***
	delete broadcast-forwarder udp ***udp-port interface number destination ip-address***
	disable broadcast-forwarding udp ***udp-port interface number***
	enable broadcast-forwarding udp ***udp-port interface number***
	list broadcast-forwarding

To Configure IP	Use These Commands
Cache Size	**list sizes** **set cache-size *number***
Default Network Gateway	**delete default network-gateway** **set default network-gateway *gateway ip-address***
Default Subnet Gateway	**delete default subnet-gateway *subnetted-network*** **set default subnet-gateway *subnetted-network***
Directed Broadcast	**disable directed-broadcast** **enable directed-broadcast**
Filter	**add filter *ip-address ip-mask*** **delete filter *destination*** **list filter**
Internal IP Address	**set internal-ip-address *ip-address***
IP host only default network gateway	**delete ip-host-only-default network gateway** **set ip-host-only-default network gateway *IP-address***
IP host only default subnet gateway	**delete ip-host-only-default subnet gateway** **set ip-host-only-default subnet gateway *IP-address***
Reassembly Size	**list size** **set reassembly-size *number***
RFC925 Routing	**disable rfc925-routing** **enable rfc925-routing**
Route	**add route *ip-network/subnet ip-mask next-hop cost*** **change route *destination new-mask new-first-hop new-cost*** **delete route *destination*** **list route**
Routing Table Size	**list size** **set routing table-size**

To Configure IP	Use These Commands
Router ID	**list all** **set router-id** ***IP-address***

Table B-4: IP Monitor Commands

To Monitor IP	Use These Commands
Access	**access**
Counters	**counters**
Dump	**dump**
ICMP-counters	**icmp-counters**
Interface	**interface**
Route	**route** ***IP-destination-address***
Sizes	**sizes**
Static	**static**
Traceroute	**traceroute** ***IP-destination-address***

Table B-5: OSPF Config Commands

To Configure OSPF	Use These Commands
Area	**delete area** ***area#***
	list area
	set area
AS Boundry	**disable as boundry routing**
	enable as boundary
Comparison	**set comparison**
Interface	**delete interface** ***IP-address***
	disable interface ***IP-address***
	enable interface ***IP-address***
	list interface
	set interface ***IP-address***
Neighbor	**add neighbor** ***IP-address IP-address of neighbor***
	delete neighbor ***IP-address IP-address of neighbor***
	list neighbor
Nonbroadcast	**delete non-broadcast** ***IP-address***
	list non-broadcast network description
	set non-broadcast ***IP-address***
OSPF Routing Protocol	**disable ospf routing protocol**
	enable ospf routing protocol
Range	**add range** ***area# IP-address Mask***
	delete range ***area# IP-address***
Virtual Link	**delete virtual-link**
	disable virtual-link ***IP-address***
	enable virtual-link ***IP-address***
	list virtual-link
	set virtual-link

Table B-6: OSPF Monitor Commands

To Monitor OSPF	Use These Commands
Advertisement	**advertisement ls-type link-state-id advertising-router area-id**
Area	**area**
AS External	**as-external**
Database	**database** ***area-id***
Dump	**dump**
Interface	**interface** ***ip-address***
Neighbor	**neighbor** ***ip-address***
Routers	**routers**
Size	**size**
Statistics	**statistics**
Traceroute	**traceroute** ***IP-destination-address***

Table B-7: RIP Config Commands

To Configure RIP	Use These Commands
Accept RIP Route	**add accept-rip-route** ***ip-network/subnet*** **delete accept-rip-route** ***net-number*** **list rip-routes-accept**
Broadcast Address	**set broadcast-address** ***ip-interface-address***
Originate RIP Default	**set originate-rip-default**
Override Default	**disable override default** ***gateway address*** **enable override default** ***gateway address***
Override Static Routes	**disable override static-routes** ***ip-address*** **enable override static-routes** ***ip-address***
Receiving RIP	**disable receiving rip** ***ip-address*** **enable receiving rip** ***ip-address***
Receiving Dynamic Nets	**disable receiving dynamic nets** ***ip-address*** **enable enable receiving dynamic nets** ***ip-address***
Receiving Dynamic Subnets	**disable receiving dynamic subnets** ***ip-address*** **enable receiving dynamic subnets** ***ip-address***
RIP	**disable rip** **enable rip**
Sending Default Routes	**disable sending default-routes** ***ip-address*** **enable sending default-routes** ***ip-address***
Sending Net Routes	**disable sending net-routes** ***ip-address*** **enable sending net-routes** ***ip-address***
Sending Poisoned Reverse Routes	**disable sending poisoned-reverse-routes** ***ip-address*** **enable sending poisoned-reverse-routes** ***ip-address***
Sending Subnet Routes	**disable sending subnet-routes** ***ip-address*** **enable sending subnet-routes** ***ip-address***
Sending Static Routes	**disable sending static-routes** ***ip-interface-address*** **enable sending static-routes** ***ip-interface-address***

Appendix C

VNswitch Counters

Overview

Introduction

This appendix provides an overview of the VNswitch counters and the effect of packets on counters as packets flow through the router.

In This Appendix

The following topics are covered in this chapter:

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 four layers and shows the relationships between logical interfaces, bridge ports, VIs, and the IP router. The figure shows which counters are incremented for a typical path a packet can take within a VNswitch.

Figure C-1: Packet Flow

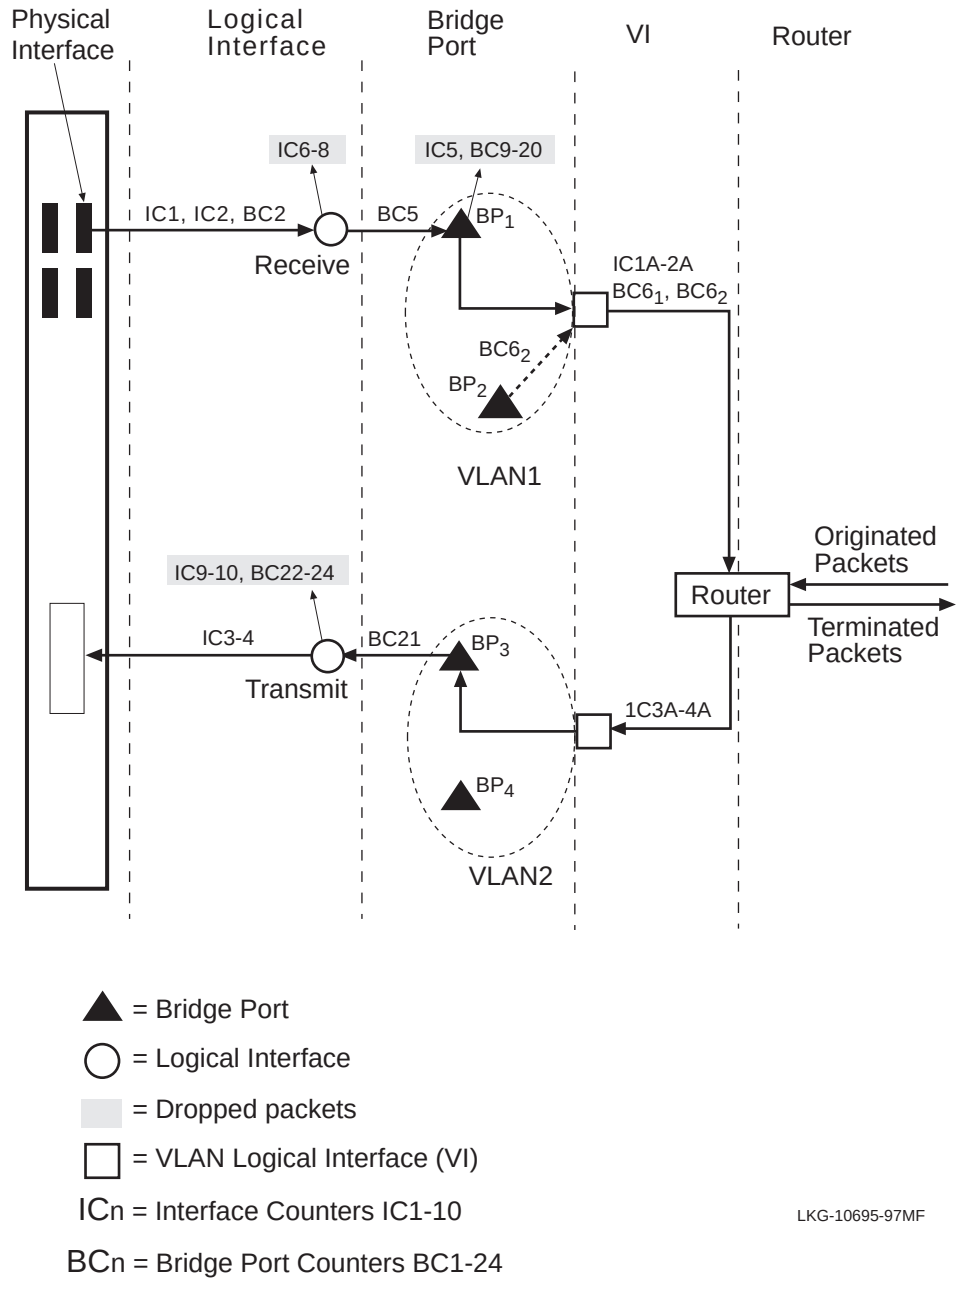

Packet Counter Overview

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 can 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 causing a single bridge error or a dropped packet counter to increment 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 packet to be dropped and the error counter to be incremented 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 are paired one-to-one with VLANs, which are groups of bridge ports. VIs submit packets for routing on behalf of any ports within their VLAN. VI receive 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, 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.

Router Packet Overview

The VNswitch offers two primary services: bridging and routing. Routing is layered on top of bridging and packets destined for routing are subject to the effects of bridging. For example, routed packets can be dropped due to a user-defined protocol filter on a bridge port and these packets are counted as being received by bridging, but not by routing.

In addition to the packets sent to the router, the router can generate packets to send out. For example, the router is capable of generating a Ping request for another node on the network. In this situation, the Ping packets are counted by IP counters, VI transmit counters, interface transmit counters, and bridge port transmit counters.

Packets destined to the router are terminated by the router. You can filter packets by enabling a filter to suppress packets destined to a network or subnet, or by disabling a protocol type. In this situation, only the interface and bridge port receive packet counters are affected.

Supported Counters

This section provides a brief description of interface counters, bridge port counters and the relationship between both counter types.

For a complete definition of interface counters and bridge port counters, refer to the *DIGITAL VNswitch 900 Series Switch Management* guide.

Interface Counters

The following table 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.

Interface Counter Number	Interface Counter Name
IC1, IC1A	Unicast packets received
IC2, IC2A	Multicast packets received
IC3, IC3A	Unicast packets transmitted
IC4, IC4A	Multicast packets transmitted
IC5	Input overflow drops
IC6	Input error drops
IC7	Input unknown protocol drops
IC8	Input congestion control drops
IC9	Output overflow drops
IC10	Output error drops

Bridge Port Counters

The following table describes each VNswitch bridge port counter (BC):

Counter Number	Counter Name
BC1	Port restarts
BC2	Total frames received by interface
BC3	IP frames fragmented
BC4	IP frames not fragmented
BC5	Frames submitted to bridging
BC6	Frames submitted to routing
BC7	Frames with unknown destination address
BC8	Frames causing learning transactions
BC9	Source address filter drops
BC10	Destination address filter drops
BC11	Protocol filter drops
BC12	Address rate limiting drops
BC13	Protocol rate limiting drops
BC14	Input buffer overflow drops
BC15	Input queue overflow drops
BC16	Source or destination port blocked drops
BC17	Terminating queue overflows
BC18	Fragmentation queue overflows
BC19	Translate flood queue overflows
BC20	Translation failures
BC21	Frames sent by bridging
BC22	Transmit queue overflows
BC23	Transmit errors
BC24	Too big to send on port drops

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

The sum of unicast and multicast packets received on an interface is the total of all received packets. It is expressed as:

IC1 + IC2 = BC2

(Unicast packets received + Multicast packets received = Total frames received by the interface)

Example 2

Some packets received on an interface may be dropped before being submitted to bridging. It is expressed as:

BC2 ≥ BC 5

(Total frames received by interface ≥ Frames submitted to bridging.)

Example 3

Some packets submitted to bridging will be submitted to routing. The rest are either dropped or bridged. It is expressed as:

BC5 ≥ BC6

(Frames submitted to bridging ≥ Frames submitted to routing)

Example 4

The packets submitted to routing by a port's VI may represent only a portion of the packets received by that VI since the VI's VLAN may contain other ports. It is expressed as:

BC6 ≤ IC1A + IC2A

(Frames submitted to routing ≤ Unicast packets received + Multicast packets received)

Transmit Relationships

Example 1

Some packets sent by bridging may be dropped at the interface layer. It is expressed as:

IC3 + IC4 ≤ BC20

(Unicast packets transmitted + Multicast packets transmitted ≤ Translation failures)

Appendix D

Configuration Examples

Overview

Introduction

This appendix provides examples of how to configure routing on a DIGITAL VNswitch 900EF module. It describes the steps a user performs to install and configure a typical VLAN routing network, including all the CLI commands necessary to configure routing for a variety of configurations.

In This Appendix

The following topics are covered in this appendix:

Common Example Elements

The examples in this appendix use the same hardware and software components, the same initial settings, the same connection method, and the same network topology.

Hardware Components

The discussion in this appendix uses the same hardware components for all examples. The hardware consists of a DIGITAL VNswitch 900EF module configured with the factory default settings and connected to a DIGITAL MultiSwitch 900.

Bridge Settings

The VNswitch is factory configured for plug-n-play bridging as the factory default setting. This means that all the bridge ports are configured in a single default VLAN, with routing disabled.

Connecting to the Configuration Console

The first task is to connect to the module's configuration console. This can be done a variety of ways, including connecting through the MultiSwitch 900 and via Telnet. (Refer to Starting and Terminating Console Sessions in Chapter 1 for details.)

Network Topology

All the examples use the same network topology. The network consists of three VLANs. The first VLAN is dedicated to users with a VLAN name of *DEFAULT*. The second VLAN is dedicated to the engineering department with a VLAN name of *engineering*. The third VLAN is dedicated to an FDDI backbone with a VLAN name of *backbone*. The existing FDDI backbone is already running RIP and the example describes how you configure a new VNswitch 900EF to support the new Default and Engineering groups. Figure D-1 illustrates the network topology used in all the examples.

Figure D-1: Example Network Topology

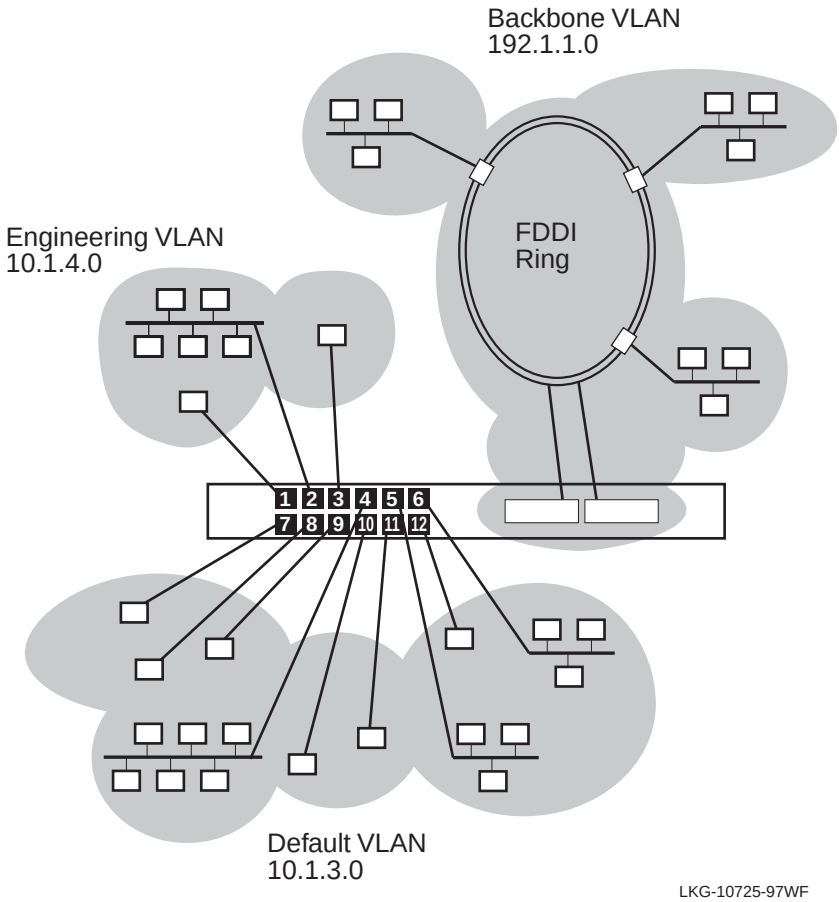

Configuring IP and RIP on a VLAN

The following examples assume that the connection to the VNswitch console is achieved using the Redirect option from the MultiSwitch 900 menu.

Enabling Routing Globally

To operate a VNswitch router on a VLAN, you must first enable routing globally. The router is initially configured with routing globally disabled. Issuing the **enable routing** command and answering **yes** to the prompt, the router automatically invokes a restart. To enable routing globally, enter:

config

```
Config>enable routing
```

Press Return. The following is displayed and requires action:

```
Enable RIP listening after restart [No]?
Default Gateway [0.0.0.0]?

When the box reboots the MAC address assigned to the
interface associated with the HST address may be different
from the one currently being used.  Therefore you may need
to flush the ARP cache on your host before you can reconnect
via Telnet.

***WARNING*** This will invoke an automatic RESTART
Are you sure you want to do this (Yes or No): Yes

System Restart ...
```

After the system is restarted, the VNswitch Installation Menu appears.

Routing on your VNswitch is now enabled.

Creating a VSD

After the system restarts, routing is enabled on the DEFAULT VSD, which is attached to the first VLAN interface (VI). This VI has the original HST IP address assigned to it, unless you did not have an IP address originally, or you choose not to transfer it after enabling routing.

The next task is to configure three VLANS and assign them to VIs. The backbone VLAN contains the FDDI for connection to the FDDI backbone. The engineering VLAN contains the first three Ethernet ports (Eth/1, Eth/2 and Eth/3), which will be used by the engineering group. The DEFAULT VLAN is for users, and contains all the other ports.

Since you are not configuring over a Telnet connection, you can move the ports between VLANs without losing your console connection. Also, since the DEFAULT VLAN already exists, you do not need to create it.

Once you have enabled routing and entered the VSD Config process, you are now ready to create two VSDs. To create two VSDs named backbone and engineering, enter:

```
Main>config

Config>vlans

VSD Config>create vsd
VSD Name: [ ] backbone
Bridge Ports (range 1-13): [ ]? 13
VNbus tag (range 66-128): [ ]
Routing over VI (none, any, or one of 15-45): [any]? *
VSD 2 created.

VSD Config>create vsd
VSD Name: [ ] engineering
Bridge Ports (range 1-13): [ ]? 1-3
VNbus tag (range 66-128): [ ]
Routing over VI (none, any, or one of 16-45): [any]? *
VSD 3 created.

VSD Config>list all
VSD Name        Ports       VNbus tag     ifc
 1  DEFAULT     4-12            65        14
 2  backbone    13                        15
 3  engineering 1-3                       16
```

***** The default "Any" assigns the lowest available VI number.

Configuring IP and RIP

With the ports assigned to the correct VLANs, the IP addresses are assigned to the appropriate VIs and RIP is enabled and configured on each (See Figure D-2). The configuration of IP is dynamic (therefore, the commands take effect immediately after they are entered). In this example, you use a standard 24-bit subnet mask for all subnets. By default, RIP is set to advertise and receive routes on an interface, so you do not need to change any of the RIP interface settings in the following task. To configure IP and RIP, enter:

```
Config>ip

Internet protocol user configuration
IP config>

IP config>add address 14 10.1.3.1 255.255.255.0
IP config>add address 15 192.1.1.1 255.255.255.0
IP config>add address 16 10.1.4.1 255.255.255.0

IP Config>enable rip

IP Config>list address
IP addresses for each interface:
Ifc  0                                IP disabled on this interface
...
Ifc 14   10.1.3.1      255.255.255.0    Network broadcast,   fill 1
Ifc 15   192.1.1.1     255.255.255.0    Network broadcast,   fill 1
Ifc 16   10.1.4.1      255.255.255.0    Network broadcast,   fill 1
Ifc 17                                IP disabled on this interface
...
Ifc 45 IP disabled on this interface
Router-ID: Not set
Internal IP address: Not set
```

The module is now transmitting and receiving RIP packets on each VLAN. To check that the IP routing tables in the module contain all the routes, at the `Monitor>` prompt, enter **ip dump**.

Figure D-2: Example Configuring IP

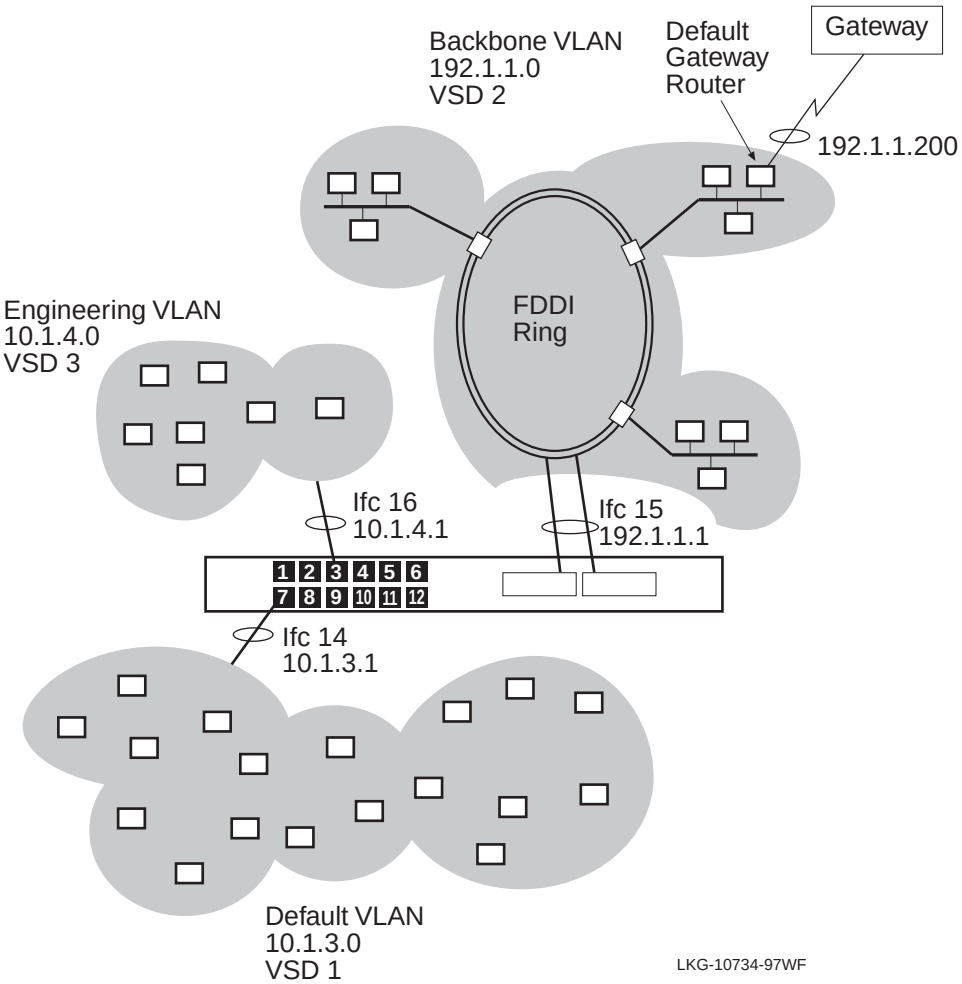

Modifying IP and RIP for Send-Only Operation

The engineering group wants to run the gated program on their UNIX workstations; however, you do not want to allow the possibility of any unofficial routes being announced by these workstations from being learned by RIP on the VNswitch and propagated to the rest of the network. In addition, you want to ensure that a default route is always propagated to the engineering VLAN, but not the backbone VLAN, so that the UNIX workstations, which are listening to the default route announcements, can find the local router.

The following task disables the reception of RIP packets on the engineering VLAN, enables the announcement of default routes to the engineering network (by default RIP does not announce default routes), and enables the fabrication of a default route if there is not one already in the routing table. To modify IP and RIP for Send-Only operation, enter:

```
IP Config>disable receiving rip 10.1.4.1

IP Config>enable sending default-routes 10.1.4.1

IP Config>set originate-rip-default
Always originate default route? [No]: yes
Originate default of cost [1]?
```

Modifying IP and RIP to Define a Static Default Route

There is a single router (192.1.1.200) on the backbone VLAN that is the gateway to the rest of the network and the Internet. This gateway router does not announce a default route in its RIP announcements. You want to set up a static default route from the VNswitch to the gateway router and have the default route announced in the engineering and DEFAULT VLANs.

The following task defines the static default route to the gateway router (at a cost of 1) and enables announcement of the default route into the engineering and DEFAULT VLANs respectively (by default announcement of default routes in RIP is disabled). To modify IP and RIP to define a static default route, enter:

```
IP Config>add route
IP destination [0.0.0.0]?
Address mask [0.0.0.0]?
Via gateway at [0.0.0.0]? 192.1.1.200
Cost [1]?

IP Config>enable sending default-routes 10.1.4.1

IP Config>enable sending default-routes 10.1.3.1
```

Modifying IP and RIP to Receive a Default Route

There is a single router (192.1.1.200) on the backbone VLAN that is the gateway to the rest of the network and the Internet. This gateway router announces a default route in its RIP announcements. You want to receive this default and readvertise it on the DEFAULT and engineering VLANs.

The following steps enable reception of the default route on the backbone interface (by default the default route is ignored in received RIP packets) and enable announcement of the default route into the engineering and DEFAULT VLANs respectively (by default announcement of default routes in RIP is disabled). To modify IP and RIP to receive a default route, enter:

```
IP Config>enable override default 192.1.1.1

IP Config>enable sending default-routes 10.1.3.1

IP Config>enable sending default-routes 10.1.4.1
```

Using Access Controls

This example demonstrates the use of access controls using the same network configuration as in the RIP Configuration section. Assume that you want to prevent Telnet access to the engineering VLAN from the DEFAULT VLAN. To do this you use IP access controls to disable use of the Telnet protocol (which uses TCP port 23) from the DEFAULT VLAN subnets.

Disabling Telnet Access from the Default VLAN

The commands to set up and enable the appropriate access controls are shown in the following example. The order is important in access-control processing (controls at the top of the list are checked first and processing halts when the first matching control is found). Also, the default action (if no access-control matches the packet) is to discard the packet.

The following task blocks any Telnet protocol packets on the default subnet 10.1.3.0 (DEFAULT VLAN) and the destination address is to the engineering VLAN. This task also enables access for all other protocols to all other destinations. To disable Telnet access from the DEFAULT VLAN, enter:

```
IP Config> add access-control
```

```
Enter type [E]?
Internet source [0.0.0.0]? 10.1.3.0
Source mask [255.255.255.255]? 255.255.255.0
Internet destination [0.0.0.0]? 10.1.4.0
Destination mask [255.255.255.255]? 255.255.255.0
Enter starting protocol number ([CR] for all) [-1]? 6
Enter ending protocol number [6]?
Enter starting port number ([CR] for all) [-1]? 23
Enter ending port number [23]?
IP Config>add access-control I 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 -1 -1
IP Config>list access-control
Access Control is: enabled
List of access control records:
```

	Ty	Source	Mask	Destination	Mask	Beg Pro	End Pro	Beg Prt	End Prt
1	E	10.1.3.0	FFFFFF00	10.1.4.0	FFFFFF00	6	6	23	23
2	I	0.0.0.0	00000000	0.0.0.0	00000000	0	255	0	65535

NOTES

This task blocks Telnet access from the engineering VLAN to the DEFAULT VLAN since responses from the DEFAULT VLAN are blocked.

The second **add access control** command in the example above is normally required.

Enabling Access Controls

For access controls to take effect, you must set access control on and then restart the module. To enable access controls, enter:

```
IP Config>set access-control on

IP Config>exit
Config>exit

Main>

Main>restart
Are you sure you want to restart the system? (Yes or [No]): yes
```

Modifying Access Controls to Enable Telnet from a Single Host

The network manager has a PC in the DEFAULT VLAN that requires Telnet access to the engineering VLAN. In this example, an extra access control is added to allow that particular PC (10.1.3.15) Telnet access to any destination.

This task allows Telnet access from source address 10.1.3.15 to any destination. The task allows Telnet packets from any source address to reach the 10.1.3.15 PC. The task also moves the two new access controls to the top of the list so that they are executed before the one that blocks all Telnet access from the DEFAULT VLAN. To modify access controls to enable Telnet from a single host ,enter:

```
IP Config> add access-control
```

```
Enter type [E]? I
Internet source [0.0.0.0]? 10.1.3.15
Source mask [255.255.255.255]?
Internet destination [0.0.0.0]?
Destination mask [255.255.255.255]? 0.0.0.0
Enter starting protocol number ([CR] for all) [-1]? 6
Enter ending protocol number [6]?
Enter starting port number ([CR] for all) [-1]? 23
Enter ending port number [23]?
IP Config>add access-control I 0.0.0.0 0.0.0.0 10.1.3.15 255.255.255.255 6 6 23 23
IP Config>list access-control
Access Control is: enabled
List of access control records:
```

	Ty	Source	Mask	Destination	Mask	Beg Pro	End Pro	Beg Prt	End Prt
1	E	10.1.3.0	FFFFFF00	10.1.4.0	FFFFFF00	6	6	23	23
2	I	0.0.0.0	00000000	0.0.0.0	00000000	0	255	0	65535
3	I	10.1.3.15	FFFFFF00	0.0.0.0	00000000	6	6	23	23
4	I	0.0.0.0	00000000	10.1.3.15	FFFFFFFF	6	6	23	23

```
IP Config>move access-control
Enter index of control to move [1]? 3
Move record AFTER record number [0]?
About to move:
```

	Ty	Source	Mask	Destination	Mask	Beg Pro	End Pro	Beg Prt	End Prt
3	I	10.1.3.15	FFFFFFFF	0.0.0.0	00000000	6	6	23	23

```
to be the first element in the list
Are you sure this is what you want to do(Yes or [No]): yes
```

```
IP Config>move access-control
Enter index of control to move [1]? 4
Move record AFTER record number [0]?1
About to move:

                                                          Beg     End     Beg       End
      Ty    Source       Mask        Destination    Mask    Pro     Pro     Prt       Prt
 4    I     0.0.0.0      00000000    10.1.3.15      FFFFFFFF  6       6     23        23
 to be after:
 1    I     10.1.3.15    FFFFFFFF    0.0.0.0        00000000  6       6     23        23
Are you sure this is what you want to do(Yes or [No]): yes
IP Config>list access-control
Access Control is: enabled
List of access control records:

                                                          Beg     End     Beg       End
      Ty    Source       Mask        Destination    Mask    Pro     Pro     Prt       Prt
 1    I     10.1.3.15    FFFFFFFF    0.0.0.0        00000000  6       6     23        23
 2    I     0.0.0.0      00000000    10.1.3.15      FFFFFFFF  6       6     23        23
 3    E     10.1.3.0     FFFFFF00    10.1.4.0       FFFFFF00  6       6     23        23
 4    I     0.0.0.0      00000000    0.0.0.0        00000000  0     255       0     65535
```

Configuring OSPF

This example uses the same network configuration as described in Modifying Access Controls to Enable Telnet from a Single Host. In this example, however, OSPF replaces RIP as the routing protocol. The existing backbone VLAN is part of the OSPF backbone area. The DEFAULT and engineering VLANs are placed in a new OSPF area, 1.1.1.1. This example assumes that routing is enabled, the engineering and backbone VLANs are configured, and the IP addresses are configured (refer to Configuring IP and RIP on a VLAN).

Configuring OSPF Areas

In this example, the VNswitch is connected to two areas. The first area is the OSPF backbone area that has the special area ID of 0.0.0.0. The second area is a new OSPF area that contains the DEFAULT and engineering VLANs. This new area can have any unique area ID (except the reserved one of 0.0.0.0). In this example, the area 10.1.0.0 is used to show that the area contains subnets of 10.1.0.0. (See Figure D-3.)

This task sets areas 0.0.0.0 and 10.1.0.0. Both areas are none-stub areas, and the backbone area is using simple password authentication. To configure OSPF areas, enter:

```
Config>ospf
Open SPF-Based Routing Protocol configuration console

OSPF Config>set area
Area number [0.0.0.0]?
Authentication Type [0]? 1

OSPF Config>set area
Area number [0.0.0.0]? 10.1.0.0
Authentication Type [0]?
Is this a stub area? [No]:

OSPF Config>list area
                --Area configuration--
Area ID    AuType        Stub? Default-cost Import-summaries?
 0.0.0.0 1=Simple-pass No   N/A        N/A
10.1.0.0 0=None        No   N/A        N/A
```

Figure D-3: Example Configuring OSPF Areas

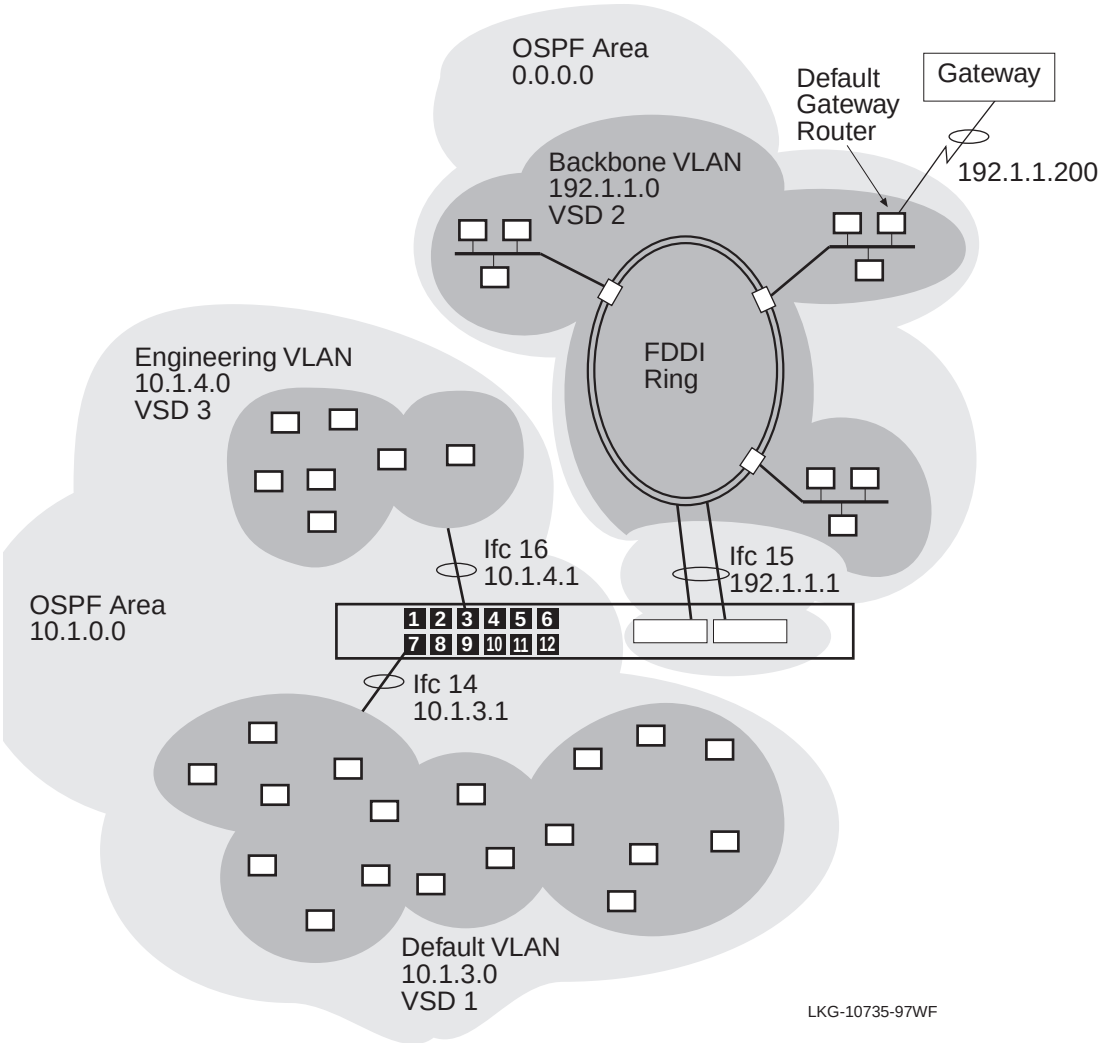

Configuring OSPF Interfaces

In this example, the OSPF interfaces are configured and each interface is associated with a single area. The backbone VLAN is part of the OSPF backbone area and the engineering and DEFAULT VLANs are part of the 10.1.0.0 area. Since the backbone area is configured to run with password authentication, the **set interface** command prompts for a password (a string of up to 8 characters), which should be the same as the password being used by other routers in the OSPF backbone area.

This task configures the backbone interface (supplying a password of **mypasswd**). The task also configures the DEFAULT and engineering interfaces respectively. The three **enable interface** commands enable each of the interfaces (by default the interfaces are disabled). To configure OSPF interfaces, enter:

```
OSPF Config>set interface
Interface IP address [0.0.0.0]? 192.1.1.11
Attaches to area [0.0.0.0]?
Retransmission Interval (in seconds) [5]?
Transmission Delay (in seconds) [1]?
Router Priority [1]?
Hello Interval (in seconds) [10]?
Dead Router Interval (in seconds) [40]?
Type Of Service 0 cost [1]?
Authentication Key []? mypasswd
Retype Auth. Key []? mypasswd

OSPF Config>set interface
Interface IP address [0.0.0.0]? 10.1.4.11
Attaches to area [0.0.0.0]? 10.1.0.0
Retransmission Interval (in seconds) [5]?
Transmission Delay (in seconds) [1]?
Router Priority [1]?
Hello Interval (in seconds) [10]?
Dead Router Interval (in seconds) [40]?
Type Of Service 0 cost [1]?
Authentication Key []?
Retype Auth. Key []?
```

```
OSPF Config>set interface
Interface IP address [0.0.0.0]? 10.1.3.11
Attaches to area [0.0.0.0]? 10.1.0.0
Retransmission Interval (in seconds) [5]?
Transmission Delay (in seconds) [1]?
Router Priority [1]?
Hello Interval (in seconds) [10]?
Dead Router Interval (in seconds) [40]?
Type Of Service 0 cost [1]?
Authentication Key []?
Retype Auth. Key []?

OSPF Config>enable interface 192.1.1.11

OSPF Config>enable interface 10.1.4.11

OSPF Config>enable interface 10.1.3.11

OSPF Config>list interface
        --Interface configuration--
IP address  Sta  Area   Cost Rtrns  TrnsDly  Pri  Hello  Dead
192.1.1.11 Ena  0.0.0.0   1     5      1     1    10    40
10.1.4.11  Ena  10.1.0.0  1     5      1     1    10    40
10.1.3.11  Ena  10.1.0.0  1     5      1     1    10    40

             Authentication Keys
IP address AuType     Key (Hex/Ascii)
192.1.1.11   1=Simple-pass 0x6D79706173737764   "mypasswd"
10.1.4.11  0=None     0x0000000000000000 ""
10.1.3.11  0=None     0x0000000000000000 ""
OSPF Config>
```

Enabling OSPF

This example enables OSPF and restarts the VNswitch so that the OSPF configuration takes effect. The **enable ospf** command prompts for the number of external routes and the number of routers in the OSPF domain so that it can allocate sufficient memory for its databases. In this example, an arbitrary number of 500 external routes and 50 routers is used. To enable OSPF, enter:

```
OSPF Config>enable ospf
Estimated # external routes [0]? 500
Estimated # OSPF routers [0]? 50

OSPF Config>exit
Config>exit

Main>restart
Are you sure you want to restart the system? (Yes or [No]): yes
```

Modifying OSPF to Propagate RIP Routes

Suppose that there are other routers in engineering that are only running RIP and you want to ensure that the networks they are advertising are reachable from the backbone and DEFAULT VLANs (which are only running OSPF). To do this, you configure RIP to run on the engineering VLAN (as described in the previous example) and configure OSPF to advertise the RIP routes (shown below). In addition, you also want to advertise any static routes you may have configured.

This task configures OSPF as an autonomous system (AS) boundary router and instructs it to import RIP and static routes and advertise them in OSPF. The task also sets route comparison to type 1, which causes OSPF to advertise the imported routes as OSPF external type 1 routes rather than as external type 2 routes. To modify OSPF to propagate RIP routes, enter:

```
OSPF Config>enable as
Import RIP routes? [No]: yes
Import static routes? [No]: yes
Import direct routes? [No]:
Import subnet routes? [No]: yes
Always originate default route? [No]:

OSPF Config>set comparison
Compare to type 1 or 2 externals [2]? 1
OSPF Config>
```

At the completion of this task, perform the procedures for a **restart** command.

Index

A

B

C

D

E

F

I

L

T

U

V