User's Guide

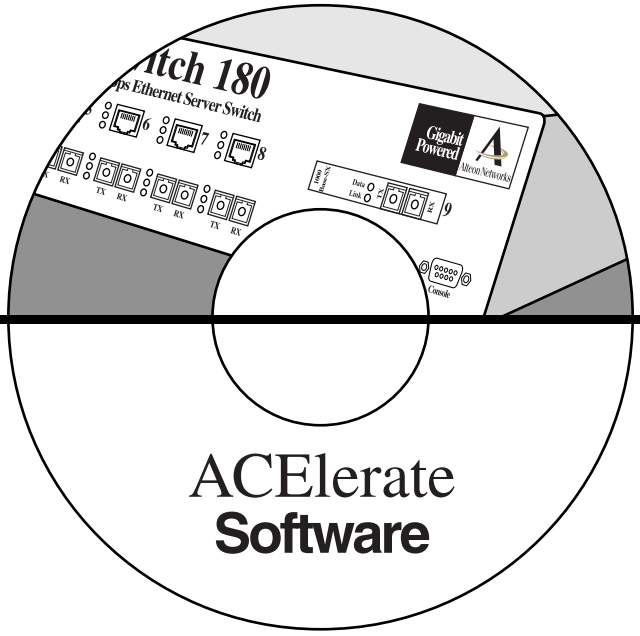

ACElerate
Software

Release 5

Part Number: 050044, Revision A

Alteon Networks
THE SERVER SWITCHING COMPANY

50 Great Oaks Boulevard
San Jose, California 95119
408-360-5500 Main
408-360-5501 Fax
www.alteon-networks.com

October 1998

Copyright 1998 Alteon Networks, Inc., 50 Great Oaks Boulevard, San Jose, California 95119, USA. All rights reserved. Part Number: 050044, Revision A.

This document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this document may be reproduced in any form by any means without prior written authorization of Alteon Networks. Documentation is provided "as is" without warranty of any kind, either express or implied, including any kind of implied or express warranty of non-infringement or the implied warranties of merchantability or fitness for a particular purpose.

U.S. Government End Users: This document is provided with a "commercial item" as defined by FAR 2.101 (Oct 1995) and contains "commercial technical data" and "commercial software documentation" as those terms are used in FAR 12.211-12.212 (Oct 1995). Government End Users are authorized to use this documentation only in accordance with those rights and restrictions set forth herein, consistent with FAR 12.211- 12.212 (Oct 1995), DFARS 227.7202 (JUN 1995) and DFARS 252.227-7015 (Nov 1995).

Alteon Networks, Inc. reserves the right to change any products described herein at any time, and without notice. Alteon Networks assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Alteon Networks, Inc. The use and purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Alteon Networks, Inc.

ACElerate™ and ACEswitch™ are trademarks of Alteon Networks, Inc. in the United States and certain other countries. Cisco® and EtherChannel® are registered trademarks of Cisco Systems, Inc. in the United States and certain other countries. Any other trademarks appearing in this manual are owned by their respective companies.

Contents

Figures

Tables

Preface

This *User's Guide* describes how to configure and use the ACElerate Release 5 software included in the Alteon Networks family of switches.

For documentation on installing the switches physically, see the hardware installation guide for your particular switch model.

Who Should Use This Book

This *User's Guide* is intended for network installers and system administrators engaged in configuring and maintaining a network. It assumes that you are familiar with Ethernet concepts, IP addressing, the IEEE 802.1d Spanning-Tree Protocol, and SNMP configuration parameters.

How This Book Is Organized

Part 1: Getting Started

These chapters introduce the major features of the switch software, and explain how to access the switch and perform basic configuration.

Chapter 1, "ACElerate Software Features," provides an overview of the major features included in Release 5 of the switch software.

Chapter 2, "The Command-Line Interface," describes how to connect to the switch and access the information and configuration menus.

Chapter 3, "First-Time Configuration," describes how to use the Setup utility for initial switch configuration and how to change the system passwords.

Part 2: The Menu System

Each chapter represents a major section within the command-line interface menu system.

Chapter 4, "Menu Basics," provides an overview of the menu system, including a menu map, global commands, and menu shortcuts.

Chapter 5, "The Information Menu," shows how to view switch configuration parameters.

Chapter 6, "The Statistics Menu," shows how to view switch performance statistics.

Chapter 7, "The Configuration Menu," shows how to configure switch system parameters, ports, VLANs, Jumbo Frames, Spanning-Tree Protocol, SNMP, Port Mirroring, IP Routing, Port Trunking, Server Load Balancing, Filtering, and more.

Chapter 8, "The Operations Menu," shows how to use commands which affect switch performance immediately, but do not alter permanent switch configurations (such as temporarily disabling ports). Also describes how to activate or deactivate optional software features.

Chapter 9, "The Boot Options Menu," describes the use of the primary and alternate switch images, how to load a new software image, and how to reset the software to factory defaults.

Chapter 10, "The Maintenance Menu," shows how to generate and access a dump of critical switch state information, how to clear it, and how to clear part or all of the forwarding database.

Part 3: Tutorials and Examples

These chapters will help you plan, implement, and administer the use of the more advanced ACElerate software features.

Chapter 11, "VLANs," describes network design and topology considerations for using VLANs.

Chapter 12, "Jumbo Frames," provides additional detail for using Jumbo Frames.

Chapter 13, "IP Routing," provides configuration background and examples for using the switch to perform routing functions.

Chapter 14, "Port Trunking," provides configuration background and examples for trunking multiple ports together.

Chapter 15, "Server Load Balancing," provides conceptual overview and configuration examples for getting the most from Server Load Balancing.

Chapter 16, "Filtering," provides conceptual overview and configuration examples for filtering and redirecting traffic.

Chapter 17, "Distributed Server Load Balancing," provides conceptual overview and configuration examples for performing server load balancing across multiple geographic sites.

Chapter 18, "Troubleshooting," describes switch configuration troubleshooting techniques.

Typographic Conventions

The following table describes the typographic styles used in this book.

Table 1 Typographic Conventions

Typeface or Symbol	Meaning	Example
`AaBbCc123`	This type is used for names of commands, files, and directories used within the text.	View the `readme.txt` file.
	It also depicts on-screen computer output and prompts.	`>> Main#`
`AaBbCc123`	This bold type appears in command examples. It shows text that must be typed in exactly as shown.	`>> Main#` **`sys`**
AaBbCc123	This italicized type appears in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command.	To establish a Telnet session, type: **`telnet`** *IP-address*
	This also shows book titles, special terms, or words to be emphasized.	Read Chapter 6 in your *User's Guide*. You *must* log in as `admin` to do this.
[ ]	Command items shown inside brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	**`ls`** [**`-a`**]

Contacting Alteon Networks

Use the following information to access Alteon Networks Online, customer support, or sales.

- Web access:

  `http://www.alteon-networks.com`

  This is the URL of Alteon Networks Online Information. This web site includes product information, software updates, release notes, and white papers. The web site also includes access to Alteon Networks Customer Support for accounts under warranty or that are covered by a maintenance contract.

- E-mail access:

  `support@alteon-networks.com`

  E-mail access to Alteon Networks Customer Support is available to accounts that are under warranty or covered by a maintenance contract.

- Telephone access to Alteon Networks Customer Support:

  1-888-Alteon0 (or 1-888-258-3660)
  1-408-360-5695

  Telephone access to Alteon Networks Customer Support is available to accounts that are under warranty or covered by a maintenance contract. Normal business hours are 8 a.m. to 6 p.m. PST.

- Telephone access to Alteon Networks Sales:

  1-888-Alteon2 (or 1-888-258-3662), and press 2 for Sales
  1-408-360-5600, and press 2 for Sales

  Telephone access is available for information regarding product sales and upgrades.

Part 1: Getting Started

CHAPTER 1

ACElerate Software Features

This chapter briefly describes the major ACElerate Release 5 software features.

Overview

The ACElerate Release 5 software offers the following features:

- Concurrent Layer 2, Layer 3 and Layer 4 switching
- Optional Application Redirection software allows the interception and redirection of client/server IP requests
- Optional Server Load Balancing software provides up to 256 real servers load balanced by up to 256 virtual servers, with each supporting multiple IP addresses and applications
- Optional Distributed Server Load Balancing software lets you balance server traffic load among up to eight remote physical sites.
- Hot Standby Support for Layer 4 Switching
- Layer 3 IP Routing software forwards frames between as many as 256 interfaces
- L3/L4 Filtering to create secure server networks
- VLAN support for up to 246 VLANs per switch
- Jumbo Frame support for frame sizes up to 9022 octets
- Cisco EtherChannel™ compatible port trunking support, allowing the creation of up to four Trunk Groups each with between two to four configured switch ports
- ACEvision web-based user interface for direct browser-to-switch interaction for configuration and monitoring
- Server Dual Homing support
- Switching Processor (SP) capability to learn up to 4095 MAC addresses
- Master Forwarding Database supports up to 8192 MAC address entries per switch
- IEEE 802.1d Spanning-Tree Protocol support
- IEEE 802.3x Flow Control support for full-duplex ports
- IEEE 802.3z Link-Negotiation support
- IEEE 802.1Q Frame Tagging when ports are enabled with VLAN tagging

- SNMP support: RFC 1213 MIB-II, RFC 1493 Bridge MIB, RFC 1398 Ethernet-like MIB, RFC 1757 RMON (Groups 1-4) MIB, and RFC 1573 Interface Extensions MIB compliant. Alteon Networks Enterprise MIB supporting the configuration and monitoring of all Alteon Networks specific features.
- Configuration and management is performed via local console port (DCE) or Telnet, and the Web-UI, with two levels of password protection
- Command-line interface Setup facility reduces the initial setup time
- TFTP download to Flash memory for software updates and upgrades

Standard Features

VLANs

Virtual Local Area Networks (*VLANs*) are commonly used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments.

The ACElerate Release 5 software supports up to 246 VLANs per switch. IEEE 802.1Q VLAN *tagging* is also supported to allow multiple VLANs per port, and to provide standards-based VLAN support for Ethernet systems.

Jumbo Frames

To reduce host frame processing overhead, the Alteon Networks switches and the ACEnic, both running operating software version 2.0 or greater, can receive and transmit frames that are larger than maximum frame size allowed on normal Ethernet.

VLANs can be configured on the same NICs and switches to separate regular traffic from Jumbo Frame traffic. End-stations with a ACEnics installed and attached to ACEswitches can communicate across both the Jumbo Frame VLANs and regular frame VLANs at the same time.

IP Routing

IP Routing allows the network administrator to seamlessly connect server IP subnets to the rest of the backbone network, using a combination of configurable IP switch interfaces and IP routing options.

The IP Routing feature enhances Alteon's Server Switching solution in the following ways:

- It provides the ability to perform Server Load Balancing (using both Layer 3 and Layer 4 switching in combination) to server subnets which are separate from backbone subnets.
- By automatically fragmenting Jumbo Frames when routing to non-Jumbo Frame subnets or VLANs, it provides another means to invisibly introduce Jumbo Frames technology into the Server Switched network.
- It provides the ability to seamlessly route IP traffic between multiple VLANs and subnets configured in the switch.

Filtering

Layer 3 (IP) and Layer 4 (Application/Protocol) filtering gives the network administrator a powerful tool to protect their server networks. Up to 224 filters can be created. Every switch port can have up to 224 of these filters applied.

Each filter can allow or deny traffic and can optionally log results, based on any combination of the following user-specified criteria:

- IP source address, by address and mask
- IP destination address, by address and mask
- Protocol type (IP, UDP, TCP, ICMP and others)
- Application source port, by name, integer or range
- Application destination port, by name, integer or range

Port Trunk Groups

Ports in a trunk group combine their bandwidth to create a single, larger virtual link. This provides the following features:

- Up to four trunk groups are supported per switch.
- Up to four ports can be trunked together to form a single virtual link with bandwidth between 2 and 4 Gigabits per second.
- Trunk groups are inherently fault tolerant: the trunk is active as long as any of its ports are available.
- Traffic on the trunk is statistically load balanced between the ports in the link.
- Trunk connections support third-party devices such as Cisco routers and switches with EtherChannel technology, and Sun's Quad Fast Ethernet Adapter.

Port Mirroring

Port mirroring provides a powerful network debugging tool. When this feature is configured, network packets being sent and/or received on a target port are duplicated and sent to a monitor port. By attaching a network analysis computer to the monitor port, you can collect detailed information about your network performance and usage.

The ACEvision Web-User Interface

With ACElerate Switching Software, the network administrator may access all switch configuration and monitoring functions through ACEvision, a web-based switch management interface. ACEvision has all of the same configuration and monitoring functions as the command-line interface, with an intuitive and easy-to-use interface structure.

Alteon Networks SNMP MIB

All configuration and monitoring data is now accessible via an enterprise Alteon Networks MIB, which can be compiled into MIB-based systems such as HP-OpenView.

RMON Lite Support

This feature provides support to RMON applications for collecting and presenting information about your network performance. Through the use of an RMON console application (available separately), you can access the following switch performance information:

- EtherStats: Real-time counters for packet and octet rates, error rates, and frame size distribution.
- History: If enabled, periodic measurements of the EtherStats are saved in switch memory. These performance snap-shots can be retrieved and displayed by your RMON application.
- Alarms and Events: Measures special user-selected conditions of which the administrator wishes to be informed (such as excessive FCS errors or high broadcast rates).

802.3x Flow Control

The ACElerate software supports 802.3x flow control on a per-port basis, on full-duplex links. 802.3x flow control provides a mechanism for Ethernet end-stations or networking devices to signal a neighbor on a full-duplex link to pause the data transmission for a short period of time. Flow control provides rudimentary capabilities for allowing a device to temporarily suspend data reception so that it can handle any data already in queues.

RFC 1573 Interface Extension MIB Compliance

Without the RFC 1573 MIB, high-speed LAN technologies such as Fast Ethernet and Gigabit Ethernet can cause frame and octet counters within the MIB-II interface to roll over in a short period of time, ruining their statistical significance.

The ACElerate software, version 2.0 and greater, supports the RFC 1573 MIB. This IF Extensions MIB allows for higher speed networking environments, providing 64-bit counters on many MIB-II statistics, plus roll-over counters for 32-bit counters.

Server Dual Homing

Server switching networks require the capability to employ resiliency and redundancy similar to FDDI network environments. The combination of Alteon Networks NICs and switches provide the Ethernet user with this capability.

For Dual Homing support, you must install two ACEnics in the same host system. These NICs are configured to provide a hot-standby failover service. The switches must be configured to support Spanning-Tree on both Gigabit Ethernet ports to support the ACEnic Dual Homing capability.

Refer to the *Installation and User's Guide* for your ACEnic Adapter for more information about this feature.

Optional Features

The following features are optional and may require additional software licences on some switches. For information on activating these features on your switch (if necessary), see "Activating Optional Software" on page 8-6.

Application Redirection Filters

Repeated client access to common web or application content across the Internet can be an inefficient use of network resources. The same filtering system that provides basic network security can also be used to intercept and redirect client traffic to cache and application servers. By redirecting client requests to a local cache or application server, you increase the speed at which clients access the information and free up valuable network bandwidth.

Server Load Balancing

With Server Load Balancing, your ACElerate powered switch is aware of the shared services provided by your server pool. The switch can then balance user session traffic among the available servers. For even greater control, traffic is distributed according to a variety of user-selectable metrics.

By helping to eliminate server over-utilization, important session traffic gets through more easily, reducing user competition for connections on overworked servers.

Intelligent health checks are performed for DNS, FTP, HTTP, NNTP, POP3 and SMTP services. If any server in a server pool fails, the remaining servers continue to provide access to vital applications and data. The failed server can be brought back up without interrupting access to services. As users are added and the server pool's capabilities are saturated, new servers can be added to the pool transparently.

Distributed Server Load Balancing

Distributed Server Load Balancing (DSLB), lets you balance server traffic load across multiple physical sites. This allows you to smoothly integrate the resources of a world-wide series of server sites, and to balance web content (or other services) intelligently among them. Alteon Network's DSLB system takes into account individual sites' health, response time, and geographic locatation for a global performance perspective.

CHAPTER 2

The Command-Line Interface

Your Alteon Networks switch is ready to perform basic switching functions right out of the box. Some of the more advanced features, however, require some administrative configuration before they can be used effectively.

The extensive ACElerate switching software included in your switch provides a variety of options for accessing and configuring the switch:

- A built-in, text-based command-line interface and menu system
- ACEvision for interactive network access through your web browser
- SNMP support for access through network management software such as HP-OpenView

The command-line interface (CLI) is the most direct method for collecting switch information and performing switch configuration. Using a basic terminal, you are presented with a hierarchy of menus that enable you to view information and statistics about the switch, and to perform any necessary configuration.

The remainder of this chapter explains how to access the CLI.

Connecting to the Switch

You can access the command-line interface in two ways:

- Using a console connection via the console port
- Using a Telnet connection over the network

Establishing a Console Connection

Requirements

To establish a console connection with the switch, you will need the following:

- An ASCII terminal or a computer running terminal emulation software set to the parameters shown in the table below:

Table 2-1 Console Configuration Parameters

Parameter	Value
Baud Rate	9600
Data Bits	8
Parity	None
Stop Bits	1

- A standard serial cable with a male DB9 connector (see your switch hardware installation guide for specifics).

Procedure

1. **Connect the terminal to the Console port using the serial cable.**
2. **Power on the terminal.**
3. **To establish the connection, press <Enter> a few times on your terminal.**

   You will next be required to enter a password for access to the switch (see "Entering Passwords" on page 2-4).

Establishing a Telnet Connection

A Telnet connection offers the convenience of accessing the switch from any workstation connected to the network. Telnet access provides the same options for user access and administrator access as those available through the console port.

To configure the switch for Telnet access, you need to have a device with Telnet software located on the same network as the switch. The switch must have an IP address. The switch can get its IP address in one of two ways:

- Dynamically, from a BOOTP server on your network
- Manually, when you configure the switch IP address (see "Setup Part 1: Basic System Configuration" on page 3-3).

Using a BOOTP Server

By default, the ACElerate software is set up to request its IP address from a BOOTP server. If you have a BOOTP server on your network, add the MAC address of the switch to the BOOTP configuration file located on the BOOTP server. The MAC address can be found on a small white label on the back panel of the switch. The MAC address can also be found in the System Information Menu (see "System Information" on page 5-2).

Running Telnet

Once the IP parameters on the switch are configured, you can access the CLI using a Telnet connection. To establish a Telnet connection with the switch, run the Telnet program on your workstation and issue the Telnet command, followed by the switch IP address:

```
telnet IP-address
```

You will next be prompted to enter a password as explained below.

If you have trouble making a Telnet connection to the switch, refer to Chapter 18, "Troubleshooting."

Entering Passwords

Once you are connected to the switch via local console or Telnet, you are prompted to enter a password. There are two levels of access to the switch: user and administrator. Each level has a different password and is granted different access privileges.

The User Account

The user has very limited power on the switch. He or she can view switch information and statistics, but can make no configuration changes. The default password for the user account is `user`.

The Administrator Account

The administrator has complete access to all menus, information, and configuration commands, including the ability to change both the user and administrator passwords. The default password for the administrator account is `admin`.

CLI vs. Setup

Once the administrator password is verified, you are given complete access to the switch. If the switch is still set to its factory default configuration, the system will ask whether you wish to run Setup (see Chapter 3, "First-Time Configuration"), a utility designed to help you through the first-time configuration process. If the switch has already been configured, the Main Menu of the CLI is displayed instead.

The following figure shows the Main Menu with administrator privileges.

```
[Main Menu]
        info  - Information Menu
        stats - Statistics Menu
        cfg   - Configuration Menu
        oper  - Operations Command Menu
        boot  - Boot Options Menu
        maint - Maintenance Menu
        diff  - Show pending config changes  [global command]
        apply - Apply pending config changes [global command]
        save  - Save updated config to FLASH [global command]
        exit  - Exit  [global command, always available]

>> Main#
```

Figure 2-1 Administrator Main Menu

Idle Timeout

By default, the switch will disconnect your console or Telnet session after five minutes of inactivity. This function is controlled by the idle timeout parameter. For information on changing this parameter, see "Configuring System Parameters" on page 7-4.

CHAPTER 3

First-Time Configuration

To help with the initial process of configuring your switch, the ACElerate software includes a Setup utility. The Setup utility prompts you step-by-step to enter all the necessary information for basic configuration of the switch.

Whenever you log in as the system administrator under the factory default configuration, you are asked whether you wish to run the Setup utility. Setup can also be activated manually from the command-line interface any time after login.

This chapter describes how to use the Setup utility and how to change system passwords.

Using the Setup Utility

Information Needed For Setup

Setup requests the following information:

- Basic system information
  - Date & time
  - Whether to use BOOTP or not
  - Whether to use Spanning-Tree Protocol or not
- Optional configuration for each port
  - Speed, duplex, flow control, and negotiation mode (as appropriate)
  - Whether to use VLAN tagging or not (as appropriate)
- Optional configuration for each VLAN
  - Name of VLAN
  - Whether the VLAN uses Jumbo Frames or not
  - Which ports are included in the VLAN

- Optional configuration of IP parameters
  - IP address, subnet mask, and broadcast address, and VLAN for each IP interface
  - IP addresses for up to four default gateways
  - Destination, subnet mask, and gateway IP address for each IP static route
  - Whether IP forwarding is enabled or not
  - Whether the RIP supply is enabled or not

Starting Setup When You Log In

The Setup prompt appears automatically whenever you login as the system administrator under the factory default settings.

1. **Connect to the switch console.**

   After connecting, the login prompt will appear as shown below.

```
Enter Password:
```

2. **Enter `admin` as the default administrator password.**

   If the factory default configuration is detected, the system prompts:

```
Connected to Alteon AceSwitch 180
15:38:00 Wed June 17, 1998

The switch is booted with factory default configuration.
  To ease the configuration of the switch, a "Set Up" facility which
  will prompt you with those configuration items that are essential
  to the operation of the switch is provided.
Would you like to run "Set Up" to configure the switch? [y/n]:
```

NOTE – If the default `admin` login is unsuccessful, or if the administrator Main Menu appears instead, the system configuration has probably been changed from the factory default settings. If you are certain that you need to return the switch to its factory default settings, see "Selecting a Configuration Block" on page 9-4.

---

3. **Enter `y` to begin the initial configuration of the switch, or `n` to bypass the Setup facility.**

Stopping and Restarting Setup Manually

Stopping Setup

To abort the Setup utility, press <Ctrl-C> during any Setup question. When you abort Setup, the system will prompt:

```
Would you like to run from top again? [y/n]
```

Enter **n** to abort Setup, or **y** to restart the Setup program at the beginning.

Restarting Setup

You can restart the Setup utility manually at any time by entering the following command at the administrator prompt:

```
# /cfg/setup
```

Setup Part 1: Basic System Configuration

When Setup is started, the system prompts:

```
"Set Up" will walk you through the configuration of
  System Date and Time, BOOTP, Spanning Tree, Port Speed/Mode,
  VLANs, and IP interfaces. [type Ctrl-C to abort "Set Up"]
------------------------------------------------------------

Will you be configuring VLANs? [y/n]
```

1. **Enter y if you will be configuring VLANs. Otherwise enter n.**

   If you decide not to configure VLANs during this session, you can configure them later using the configuration menus, or by restarting the Setup facility. For more information on VLANs issues, see Chapter 11, "VLANs."

   Next, the Setup utility prompts you to input basic system information.

2. **Enter the month of the current system date at the prompt:**

   ```
   System Date:
   Enter month [6]:
   ```

   Enter the month as a number from 1 to 12. To keep the current month, press <Enter>.

3. **Enter the day of the current date at the prompt:**

```
Enter day [17]:
```

Enter the date as a number from 1 to 31. To keep the current day, press <Enter>.

4. **Enter the year of the current date at the prompt:**

```
Enter year [98]:
```

Enter the last two digits of the year as a number from 00 to 99. "00" is considered 2000. To keep the current year, press <Enter>.

The system displays the date and time settings:

```
System clock set to 13:56:52 Wed June 17, 1998.
```

5. **Enter the hour of the current system time at the prompt:**

```
System Time:
Enter hour in 24-hour format [13]:
```

Enter the hour as a number from 00 to 23. To keep the current hour, press <Enter>.

6. **Enter the minute of the current time at the prompt:**

```
Enter minutes [56]:
```

Enter the minute as a number from 00 to 59. To keep the current minute, press <Enter>.

7. **Enter the seconds of the current time at the prompt:**

```
Enter seconds [52]:
```

Enter the seconds as a number from 00 to 59. To keep the current second, press <Enter>.

The system displays the date and time settings:

```
System clock set to 13:56:52 Wed June 17, 1998.
```

8. **Enable or disable the use of BOOTP at the prompt:**

```
BootP Option:
Current BOOTP usage:            enabled
Enter new BOOTP usage [d/e]:
```

If available on your network, a BOOTP server can supply the switch with IP parameters so that you do not have to enter them manually. Enter **d** to disable the use of BOOTP, or enter **e** to enable the use of BOOTP. To keep the current setting, press <Enter>.

9. **Turn Spanning-Tree Protocol on or off at the prompt:**

```
Spanning Tree:
Current Spanning Tree setting: ON
Turn Spanning Tree OFF? [y/n]
```

Enter **y** to turn off Spanning-Tree, or enter **n** to leave Spanning-Tree on.

Setup Part 2: Port Configuration

NOTE – The port configuration options shown in these steps are for the ACEswitch 180. When configuring port options for other switches, some of the prompts and options may be different.

1. **Select the port to configure, or skip port configuration at the prompt:**

```
Port Config:
Enter port number: (1-9)
```

If you wish to change settings for individual ports, enter the number of the port you wish to configure. To skip port configuration, press <Enter> without specifying any port and go to "Setup Part 3: VLANs" on page 3-7.

2. **If appropriate, configure Ethernet/Fast Ethernet port speed.**

If you selected a port that has an Ethernet/Fast Ethernet connector, the system prompts:

```
Fast Link Configuration:
Port Speed:
Current Port 1 speed setting:   10/100
Enter new speed ["10"/"100"/"any"]:
```

Enter the port speed from the options available, or enter **any** to have the switch auto-sense the port speed. To keep the current setting, press <Enter>.

3. **If appropriate, configure Ethernet/Fast Ethernet port duplex mode.**

   If you selected a port that has an Ethernet/Fast Ethernet connector, the system prompts:

```
Port Mode:
Current port 1 mode setting:     any
Enter new speed ["full"/"half"/"any"]
```

   Enter **full** for full-duplex, **half** for half-duplex, or **any** to have the switch auto-negotiate. To keep the current setting, press <Enter>.

4. **If appropriate, configure Ethernet/Fast Ethernet port flow control.**

   If you selected a port that has an Ethernet/Fast Ethernet connector, the system prompts:

```
Port Flow Control:
Current Port 1 flow control setting:    both
Enter new value ["rx"/"tx"/"both"/"none"]:
```

   Enter **rx** to enable receive flow control, **tx** for transmit flow control, **both** to enable both, or **none** to turn flow control off for the port. To keep the current setting, press <Enter>.

5. **If appropriate, configure Ethernet/Fast Ethernet port auto-negotiation mode.**

   If you selected a port that has an Ethernet/Fast Ethernet connector, the system prompts:

```
Port Auto Negotiation:
Current Port 1 autonegotiation:         on
Enter new value ["on"/"off"]:
```

   Enter **on** to enable auto-negotiation, or **off** to disable it. To keep the current setting, press <Enter>.

6. **If appropriate, configure Gigabit Ethernet port flow parameters.**

   If you selected a port that has a Gigabit Ethernet connector, the system prompts:

```
Gig Link Configuration:
Port Flow Control:
Current Port 1 flow control setting:    both
Enter new value ["rx"/"tx"/"both"/"none"]:
```

   Enter **rx** to enable receive flow control, **tx** for transmit flow control, **both** to enable both, or **none** to turn flow control off for the port. To keep the current setting, press <Enter>.

7. **If appropriate, configure Gigabit Ethernet port auto-negotiation mode.**

   If you selected a port that has a Gigabit Ethernet connector, the system prompts:

```
Port Auto Negotiation:
Current Port 1 autonegotiation:          on
Enter new value ["on"/"off"]:
```

   Enter **on** to enable port auto-negotiation, or **off** to disable it. To keep the current setting, press <Enter>.

8. **If configuring VLANs, turn VLAN tagging on or off for the port.**

   If you have selected to configure VLANs back in Part 1, the system prompts:

```
Port VLAN tagging config (tagged port can be a member of multiple VLANs)
Current TAG flag:                0 (untagged)
Enter new TAG flag [0/1]:
```

   Enter **1** if the port uses VLAN tagging. Enter **0** if the port does not use VLAN tagging. To keep the current setting, press <Enter>.

9. **The system prompts you to configure the next port:**

```
Enter port number: (1 to 9)
```

   When you are through configuring ports, press <Enter> without specifying any port. Otherwise, repeat the steps in this section.

Setup Part 3: VLANs

If you chose to skip VLANs configuration back in Part 1, skip to "Setup Part 4: IP Configuration" on page 3-9.

1. **Select the VLAN to configure, or skip VLAN configuration at the prompt:**

```
VLAN Config:
Enter VLAN number from 2 to 4094, NULL at end:
```

   If you wish to change settings for individual VLANs, enter the number of the VLAN you wish to configure. To skip VLAN configuration, press <Enter> without typing a VLAN number and go to "Setup Part 4: IP Configuration" on page 3-9.

2. **Enter the new VLAN name at the prompt:**

```
VLAN is newly created.
Pending new VLAN name: "VLAN 2"
Enter new VLAN name, without quotes:
```

3. **Enable or disable Jumbo Frame support for the VLAN at the prompt:**

```
VLAN Jumbo Frame Support:
Current Jumbo Frame support:            disabled
Enter new Jumbo Frame support [d/e]:
```

Enter **d** to disable Jumbo Frame support for the VLAN, or enter **e** to enable Jumbo Frame support for the VLAN. To keep the current setting, press <Enter>.

4. **Enter the VLAN port numbers.**

The system prompts you to define the first port in the VLAN:

```
Define ports in VLAN:
Current VLAN 2: empty
Enter port numbers one per line, NULL at end:
```

Type the first port number to add to the current VLAN and press <Enter>. The right angle prompt appears:

```
>
```

For each additional port in the VLAN, type the port number and press <Enter> to move to the next line. Repeat this until all ports for the VLAN being configured are entered. When you are finished adding ports to this VLAN, press <Enter> without specifying any port.

5. **The system prompts you to configure the next VLAN:**

```
VLAN Config:
Enter VLAN number from 2 to 4094, NULL at end:
```

Repeat the steps in this section until all VLANs have been configured. When all VLANs have been configured, press <Enter> without specifying any VLAN.

Setup Part 4: IP Configuration

If BOOTP was enabled back in Part 1, skip to "Setup Part 5: Final Steps" on page 3-12. Otherwise, if you disabled BOOTP, the system prompts for IP parameters.

IP Interfaces

IP interfaces are used for defining subnets to which the switch belongs.

Up to 256 IP interfaces can be configured on the switch. The IP address assigned to each IP interface provide the switch with an IP presence on your network. No two IP interfaces can be on the same IP subnet. The interfaces can be used for connecting to the switch for remote configuration, and for routing between subnets and VLANs (if used).

1. **Select the IP interface to configure, or skip interface configuration at the prompt:**

```
IP Config:

IP interfaces:
Enter interface number: (1-256)
```

If you wish to configure individual IP interfaces, enter the number of the IP interface you with to configure. To skip IP interface configuration, press <Enter> without typing an interface number and go to "Default Gateways" on page 3-10.

2. **For the specified IP interface, enter the IP address in dotted decimal notation:**

```
Current IP address:     0.0.0.0
Enter new IP address:
```

To keep the current setting, press <Enter>.

3. **At the prompt, enter the IP subnet mask in dotted decimal notation:**

```
Current subnet mask:            0.0.0.0
Enter new subnet mask:
```

To keep the current setting, press <Enter>.

4. **At the prompt, enter the broadcast IP address in dotted decimal notation:**

```
Current broadcast address:      0.0.0.0
Enter new broadcast address:
```

To keep the current setting, press <Enter>.

5. **If configuring VLANs, specify a VLAN for the interface.**

   This prompt appears if you selected to configure VLANs back in Part 1:

```
Current VLAN:     1
Enter new VLAN:
```

   Enter the number for the VLAN to which the interface belongs, or press <Enter> without specifying a VLAN number to accept the current setting.

6. **At the prompt, enter `y` to enable the IP interface, or `n` to leave it disabled**:

```
Enable IP interface? [y/n]
```

7. **The system prompts you to configure another interface:**

```
Enter interface number: (1-256)
```

   Repeat the steps in this section until all IP interfaces have been configured. When all interfaces have been configured, press <Enter> without specifying any interface number.

Default Gateways

1. **At the prompt, select a default gateway for configuration, or skip default gateway configuration:**

```
IP default gateways:
Enter default gateway number: (1-4)
```

   Enter the number for the default gateway to be configured. To skip default gateway configuration, press <Enter> without typing a gateway number and go to "IP Routing" on page 3-11.

2. **At the prompt, enter the IP address for the selected default gateway:**

```
Current IP address:     0.0.0.0
Enter new IP address:
```

   Enter the IP address in dotted decimal notation, or press <Enter> without specifying an address to accept the current setting.

3. **At the prompt, enter `y` to enable the default gateway, or `n` to leave it disabled:**

```
Enable default gateway? [y/n]
```


4. **The system prompts you to configure another default gateway:**

```
Enter default gateway number: (1-4)
```

Repeat the steps in this section until all default gateways have been configured. When all default gateways have been configured, press <Enter> without specifying any number.

IP Routing

When IP interfaces are configured for the various subnets attached to your switch, IP routing between them can be performed entirely within the switch. This eliminates the need to bounce inter-subnet communication off an external router device. Routing on more complex networks, where subnets may not have a direct presence on the switch, can be accomplished through configuring static routes or by letting the switch learn routes dynamically.

This part of the Setup program prompts you to configure the various routing parameters.

1. **At the prompt, enable or disable forwarding for IP Routing:**

```
Enable IP forwarding? [y/n]
```

Enter **y** to enable IP forwarding. To disable IP forwarding, enter **n** and skip to Step 2.

To keep the current setting, press <Enter>.

2. **At the prompt, enable or disable the RIP supply:**

```
Enable RIP supply? [y/n]
```

If your network uses Routing Interface Protocol (RIP), enter **y** to enable the RIP supply. Otherwise, enter **n** to disable it. When RIP is enabled, RIP listen is set by default.

Setup Part 5: Final Steps

1. **When prompted, decide whether to restart Setup or continue:**

```
Would you like to run from top again? [y/n]
```

Enter **y** to restart the Setup utility from the beginning, or **n** to continue.

2. **When prompted, decide whether you wish to review the configuration changes:**

```
Review the changes made? [y/n]
```

Enter **y** to review the changes made during this session of the Setup utility. Enter **n** to continue without reviewing the changes. We recommend that you review the changes.

3. **Next, decide whether to apply the changes at the prompt:**

```
Apply the changes? [y/n]
```

Enter **y** to apply the changes, or **n** to continue without applying. Changes are normally applied.

4. **At the prompt, decide whether to make the changes permanent:**

```
Save changes to flash? [y/n]
```

Enter **y** to save the changes to flash. Enter **n** to continue without saving the changes. Changes are normally saved at this point.

5. **If you do not apply or save the changes, the system prompts whether to abort them:**

```
Abort all changes? [y/n]
```

Enter **y** to discard the changes. Enter **n** to return to the "`Apply the changes?`" prompt.

NOTE – After initial configuration is complete, it is recommended that you change the default passwords as shown in the following section.

Setting Passwords

It is recommended that you change the user and administrator passwords after initial configuration and as regularly as required under your network security policies.

To change both the user password and the administrator password, you must login using the administrator password. Passwords cannot be modified from the user command mode.

NOTE – If you forget your administrator password, call your technical support representative for help using the password fix-up mode.

Changing the Default Administrator Password

The administrator has complete access to all menus, information, and configuration commands, including the ability to change both the user and administrator passwords.

The default password for the administrator account is `admin`. To change the default password, follow this procedure:

1. **Connect to the switch and log in using the `admin` password.**
2. **From the Main Menu, use the following command to access the Configuration Menu:**

```
Main# cfg
```

The Configuration Menu is displayed

```
[Configuration Menu]
        sys   - System-wide parameter menu
        port  - Port configuration menu
        ip    - IP configuration menu
        vlan  - VLAN configuration menu
        stp   - Spanning Tree menu
        snmp  - SNMP menu
        setup - Step by step configuration set up
        dump  - Dump current configuration to script file
        mirr  - Mirroring menu
        slb   - Server Load Balancing configuration menu
        trunk - Trunk Group configuration menu

>> Configuration#
```

3. **From the Configuration Menu, use the following command to select the System Menu:**

```
>> Configuration# sys
```

The System Menu is displayed.

```
[System Menu]
        date  - Set system date
        time  - Set system time
        usrpw - Set user password
        admpw - Set administrator password
        idle  - Set timeout for idle CLI sessions
        tnet  - Enable/disable Telnet access
        bootp - Enable/disable use of BOOTP
        http  - Enable/disable HTTP (Web) access
        wport - Set Web server port number
        cur   - Display current system-wide parameters

>> System#
```

4. **Select the administrator password by entering `admpw` at the `System#` prompt.**

```
System# admpw
```

5. **Enter the current administrator password at the prompt:**

```
Enter current administrator password:
```

NOTE – If you forget your administrator password, call your technical support representative for help using the password fix-up mode.

6. **Enter the new administrator password at the prompt:**

```
Enter new administrator password:
```

7. **Enter the new administrator password, again, at the prompt:**

```
Re-enter new administrator password:
```

8. **Apply and save your change by entering the following commands:**

```
System# apply
System# save
```

Changing the Default User Password

The user login has limited power on the switch. With it, you can view switch information and statistics, but you can't make configuration changes.

The default password for the user account is `user`. This password cannot be changed from the user account. Only the administrator has the abliity to change passwords, as shown in the following procedure.

1. **Connect to the switch and log in using the `admin` password.**
2. **From the Main Menu, use the following command to access the Configuration Menu:**

```
Main# cfg
```

The Configuration Menu is displayed

```
[Configuration Menu]
        sys   - System-wide parameter menu
        port  - Port configuration menu
        ip    - IP configuration menu
        vlan  - VLAN configuration menu
        stp   - Spanning Tree menu
        snmp  - SNMP menu
        setup - Step by step configuration set up
        dump  - Dump current configuration to script file
        mirr  - Mirroring menu
        slb   - Server Load Balancing configuration menu
        trunk - Trunk Group configuration menu

>> Configuration#
```

3. **From the Configuration Menu, use the following command to select the System Menu:**

```
>> Configuration# sys
```

The System Menu is displayed.

```
[System Menu]
        date  - Set system date
        time  - Set system time
        usrpw - Set user password
        admpw - Set administrator password
        idle  - Set timeout for idle CLI sessions
        tnet  - Enable/disable Telnet access
        bootp - Enable/disable use of BOOTP
        http  - Enable/disable HTTP (Web) access
        wport - Set Web server port number
        cur   - Display current system-wide parameters

>> System#
```

4. **Select the user password by entering `usrpw` at the `System#` prompt.**

```
System# usrpw
```

5. **Enter the current administrator password at the prompt.**

   Only the administrator can change the user password. Entering the administrator password confirms your authority.

```
Enter current administrator password:
```

6. **Enter the new user password at the prompt:**

```
Enter new user password:
```

7. **Enter the new user password, again, at the prompt:**

```
Re-enter new user password:
```

8. **Apply and save your changes:**

```
System# apply
System# save
```

Part 2: The Menu System

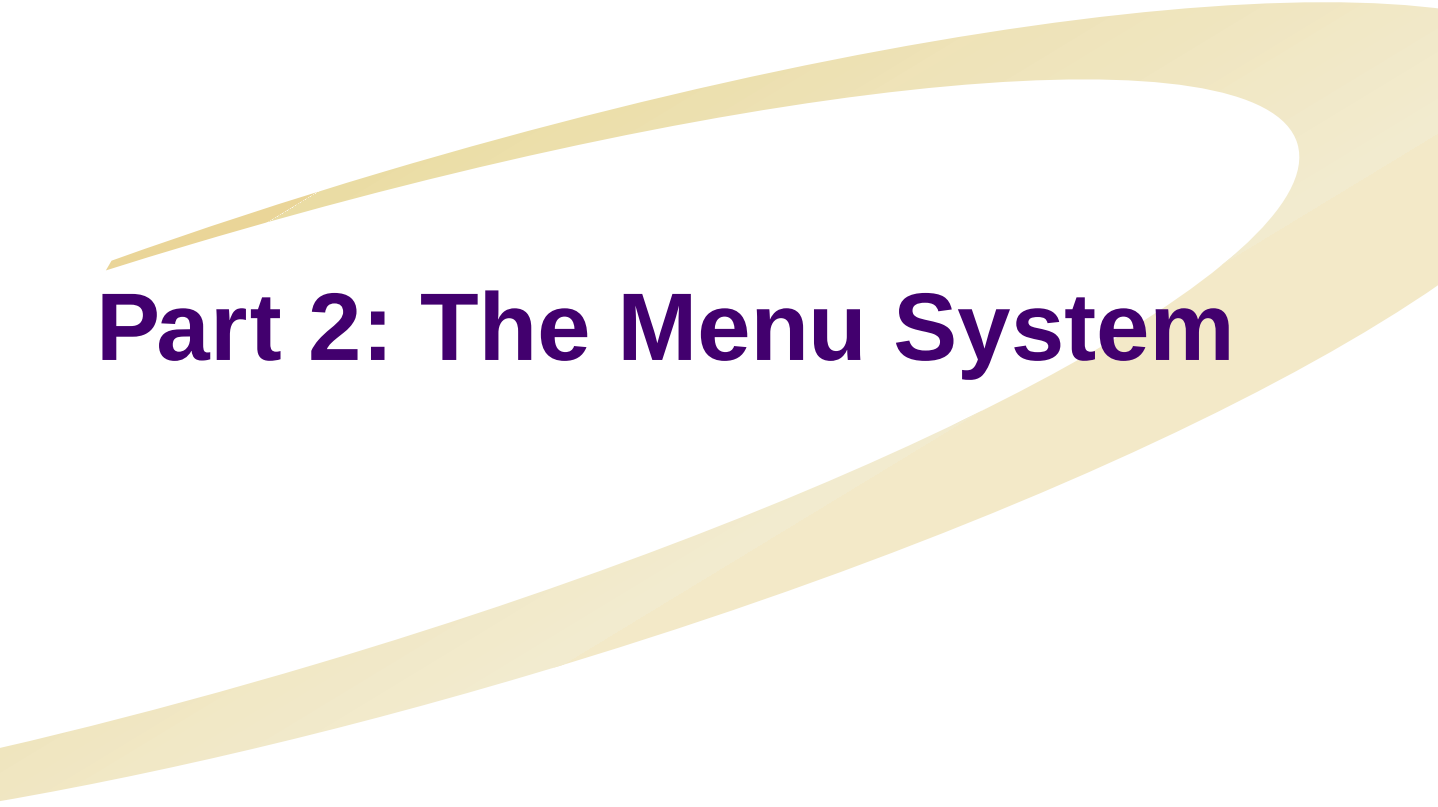

CHAPTER 4

Menu Basics

The switch's command-line interface (CLI) is used for viewing switch information and statistics. In addition, the administrator can use the CLI for performing all levels of switch configuration.

To make the CLI easy to use, the various commands have been logically grouped into a series of menus and sub-menus. Each menu displays a list of commands and/or sub-menus that are available, along with a summary of what each command will do. Below each menu is a prompt where you can enter any command appropriate to the current menu.

This chapter describes the Main Menu commands, and provides a list of commands and shortcuts that are commonly available from all the menus within the CLI.

The Main Menu

The Main Menu appears after a successful connection and login attempt. Figure 4-1 shows the Main Menu as it appears under the administrator login. Some of the features shown are not available under the user login (see "Entering Passwords" on page 2-4).

```
[Main Menu]
        info  - Information Menu
        stats - Statistics Menu
        cfg   - Configuration Menu
        oper  - Operations Command Menu
        boot  - Boot Options Menu
        maint - Maintenance Menu
        diff  - Show pending config changes  [global command]
        apply - Apply pending config changes [global command]
        save  - Save updated config to FLASH [global command]
        exit  - Exit  [global command, always available]

>> Main#
```

Figure 4-1 Administrator Main Menu

Menu Summary

- **Information Menu**

  Provides sub-menus for displaying information about how the switch is set up: from basic system settings to VLANs, Layer 4 settings, and more.

- **Statistics Menu**

  Provides sub-menus for displaying switch performance statistics. Included are port, IF, IP, ICMP, TCP, UDP, SNMP, routing, ARP, DNS, and Layer 4 statistics.

- **Configuration Menu**

  This menu is available only from an administrator login. It includes sub-menus for configuring every aspect of the switch. Changes to configuration are not active until explicitly applied. Changes can be saved to non-volatile memory.

- **Operations Command Menu**

  This menu is available only from an administrator login. Operations-level commands are used for making immediate and temporary changes to switch configuration. This menu is used for bringing ports temporarily in and out of service, performing port mirroring, and enabling or disabling Server Load Balancing functions. It is also used for activating or deactivating optional software packages.

- **Boot Options Menu**

  This menu is available only from an administrator login. This menu is used for downloading new software into the switch, selecting configuration blocks, and for resetting the switch when necessary.

- **Maintenance Menu**

  This menu is available only from an administrator login. This menu is used for debugging purposes. Chiefly, you can generate a dump of the critical state information in the switch, and clear entries in the forwarding database and the ARP and routing tables.

Menu Map

The following illustrates the administrator menu hierarchy:.

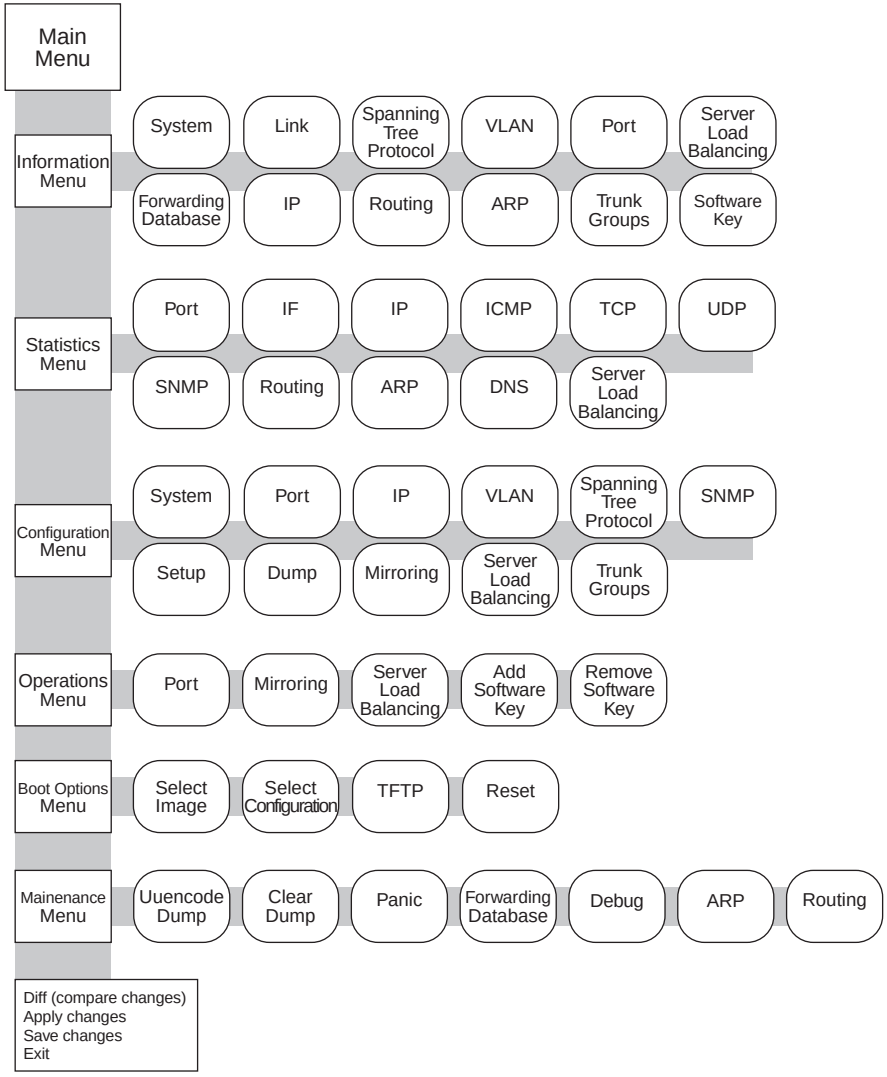

Figure 4-2 Administrator Menu Hierarchy

Global Commands

Some basic commands are recognized throughout the menu hierarchy. These commands are useful for obtaining online help, navigating through the various menus, and for applying and saving configuration changes:

Table 4-1 Global Commands

Command	Action
? *command*	Provides more information about a specific command on the current menu. When used without the *command* parameter, a summary of the global commands is displayed.
.	Display the current menu.
..	Go up one level in the menu structure.
/	If placed at the beginning of a command, go to the Main Menu. Otherwise, this is used to separate multiple commands placed on the same line.
diff	Show any pending configuration changes.
apply	Apply pending configuration changes.
save	Write configuration changes to non-volatile flash memory.
exit	Exit from the command-line interface and log out.
ping	Use this command to verify station-to-station connectivity across the network. The format is as follows: **ping** *address* [*tries* [*delay*]] Where *address* is the hostname or IP address of the device, *tries* (optional) is the number of attempts (1-32), and *delay* (optional) is the number of milliseconds between attempts. The DNS parameters must be configured if specifying hostnames (see "Domain Name System Menu" on page 7-17).
traceroute	Use this command to identify the route used for station-to-station connectivity across the network. The format is as follows: **traceroute** *address* [*max-hops* [*delay*]] Wher *address* is the hostname or IP address of the target station, *max-hops* (optional) is the maximum distance to trace (1-16 devices), and *delay* (optional) is the number of milliseconds for wait for the response. As with **ping**, the DNS parameters must be configured if speficying hostnames.
pwd	Display the command path used to reach the current menu.
lines *n*	Set the number of lines (*n*) that display on the screen at one time; the default is 24 lines. When used without a value, the current setting is displayed.
verbose *n*	Sets the level of information displayed on the screen: **0** = Quiet: Nothing appears on the screen except errors—not even prompts. **1** = Normal: Prompts and requested output are shown, but no menus. **2** = Verbose: Everything is shown. When used without a value, the current setting is displayed.

Command-Line Interface Shortcuts

As a shortcut, you can type multiple commands on a single line, separated by forward slashes (/). You can connect as many commands as required to access the menu option that you want. For example, the keyboard shortcut to access the Spanning Tree Port Configuration Menu from the Main# prompt is as follows:

```
Main# cfg/stp/port
```

Most commands can be abbreviated by entering the first characters which distinguish the command from the others in the same menu or sub-menu. For example, the command shown above could also be entered as follows:

```
Main# c/st/p
```

CHAPTER 5

The Information Menu

You can view configuration information for the switch in both the user and administrator command modes. This chapter discusses how to use the command-line interface to display switch information.

The Information Menu can be accessed from the Main Menu using the following command:

```
Main# info
```

The Information Menu is displayed:

```
[Information Menu]
        sys   - Show system information
        link  - Show link status
        stp   - Show STP information
        vlan  - Show VLAN information
        port  - Show port information
        slb   - Show Server Load Balancing information
        fdb   - Forwarding Database information menu
        ip    - Show IP information
        route - IP routing information menu
        arp   - ARP information menu
        trunk - Show Trunk Group information
        swkey - Show enabled software features

>> Information#
```

Each of these options is discussed in greater detail in the following sections.

NOTE – The sample screens shown in this chapter represent ACEswitch 180 information. Screens, menus, and parameters for other Alteon Networks switches may be slightly different.

System Information

Direct command: **/info/sys**

System information includes:

- System date and time
- Switch model name and number
- Switch name and location
- Time of last boot
- MAC address of the switch management processor
- IP address of IP interface #1
- Hardware version and part number
- Software image file and version number
- Configuration name

To view system information, at the `Information#` prompt, enter:

```
Information# sys
```

The system information is displayed:

```
System Information at 16:20:42 Wed Jan 28, 1998

Alteon ACEswitch 180
sysName:       Finance Switch
sysLocation:   Building 3A
Last boot: 15:57:56 Tue Jan 27, 1998 (reset from console)

MAC address: 00:60:cf:11:22:30    IP (If 1) address: 200.10.17.1
Hardware Revision: 2
Hardware Part No: 200009A00
Software Version 5.0.0 (FLASH image1), active configuration

>>Information#
```

Link Status

Direct command: **`/info/link`**

Link status displays configuration information about each port, including:

- Port number
- Port speed (10, 100, 10/100, or 1000)
- Duplex mode (half, full, or auto)
- Flow control for transmit and receive (no, yes, or auto)
- Link status (up or down)

To see the status of the switch ports, at the `Information#` prompt, enter:

```
Information# link
```

The current link status is displayed:

```
Port   Speed    Duplex     Flow Ctrl      Link
----   -----   --------  --TX-----RX--   ------
  1     1000     full     yes    yes       up
  2     1000     full     yes    yes       up
  3     1000     full     yes    yes       up
  4     1000     full     yes    yes      down
  5      100     full     yes    yes       up
  6       10     half      no     no       up
  7     1000     full     yes    yes      down
  8     1000*    full*    yes*    no*      up
  9     1000     full     yes    yes       up
* = value set by configuration; not autonegotiated.

>> Information#
```

Spanning-Tree Protocol Information

Direct command: **/info/stp**

The switch software uses the IEEE 802.1d Spanning-Tree Protocol (STP). In addition to seeing if STP is enabled or disabled, you can view the following STP bridge information:

- Priority
- Hello interval
- Maximum age value
- Forwarding delay
- Aging time

You can also see the following port-specific STP information:

- Port number
- Priority
- Cost
- State

To view STP information, at the `Information#` prompt, enter:

```
Information# stp
```

The current STP information is displayed:

```
Current Root:            Path-Cost Port Hello MaxAge FwdDel Aging
 8000 00:60:47:92:7e:00      100     6    2     20     15    300

Parameters:  Priority  Hello  MaxAge  FwdDel  Aging
              32768      2      20      15     300

Port  Priority   Cost      State
  1     128       10     FORWARDING
  2     128        0     FORWARDING*
  3     128        0      DISABLED *
  4     128        0      DISABLED
  5     128        5     FORWARDING
  6     128        0      DISABLED
  7     128        0      DISABLED
  8     128        0      DISABLED
  9     128        1     FORWARDING
* = STP turned off for this port.

>> Information#
```

The following table describes the STP parameters.

Table 5-1 Spanning Tree Parameter Descriptions

Parameter	Description
`Priority` (bridge)	The bridge priority parameter controls which bridge on the network will become the STP root bridge.
`Hello`	The hello time parameter specifies, in seconds, how often the root bridge transmits a configuration bridge protocol data unit (BPDU). Any bridge that is not the root bridge uses the root bridge hello value.
`MaxAge`	The maximum age parameter specifies, in seconds, the maximum time the bridge waits without receiving a configuration bridge protocol data unit before it reconfigures the STP network.
`FwdDel`	The forward delay parameter specifies, in seconds, the amount of time that a bridge port has to wait before it changes from learning state to forwarding state.
`Aging`	The aging time parameter specifies, in seconds, the amount of time the bridge waits without receiving a packet from a station before removing the station from the Forwarding Database.
`Priority` (port)	The port priority parameter helps determine which bridge port becomes the designated port. In a network topology that has multiple bridge ports connected to a single segment, the port with the lowest port priority becomes the designated port for the segment.
`Cost`	The port path cost parameter is used to help determine the designated port for a segment. Generally speaking, the faster the port, the lower the path cost. A setting of 0 indicates that the cost will be set to the appropriate default after the link speed has been autonegotiated.
`State`	The state field shows the current state of the port. The state field can be either; `BLOCKING`, `LISTENING`, `LEARNING`, `FORWARDING`, or `DISABLED`.

VLAN Information

Direct command: **/info/vlan**

VLAN information includes:

- VLAN Number
- VLAN Name
- Status
- Jumbo Frame usage
- Port membership of the VLAN

To view VLAN information for all VLANs, at the Information# prompt, enter:

```
Information# vlan
```

The current VLAN information is displayed:

```
VLAN                Name                Status Jumbo  Ports
----  --------------------------------- ------ -----  -------------
1     Default VLAN                        ena    n    6-10
1000  Engineering                         ena    n    4 5
4094  Marketing                           ena    n    1-3

>> Information#
```

To view VLAN information for a particular VLAN, at the Information# prompt, enter the VLAN number. For example:

```
Information# vlan 4094
```

The information for the selected VLAN is displayed:

```
VLAN                Name                Status Jumbo  Ports
----  --------------------------------- ------ -----  -------------
4094  Marketing                           ena    n    1-3

>> Information#
```

Port Information

Direct command: **/info/port**

Port information includes:

- Port number
- Whether the port uses VLAN Tagging or not
- RMON (ACEswitch 180 only)
- Port VLAN ID (PVID)
- Port name
- VLAN membership

To view port information, at the Information# prompt, enter:

```
Information# port
```

The port information is displayed:

```
Port  Tag  RMON  PVID       NAME                 VLAN(s)
----  ---  ----  ----  --------------  ---------------------------
  1    n    d    4094  ACEdirector     4094
  2    n    d    4094  ACME hub        4094
  3    n    d    4094  ACME router     4094
  4    n    d    1000  web-cache 1     1000
  5    n    d    1000  web-cache 2     1000
  6    n    d       1                     1
  7    n    d       1                     1
  8    n    d       1  ugroup 1           1
  9    n    d       1  ugroup 2           1

>> Information#
```

Server Load Balancing Information

Direct command: **`/info/slb`**

Server Load Balancing information includes the following:

- Failover State

  If the switch is configured as part of a pair for hot-standby redundancy, switch IP addresses, health status, and mode (active or standby) are displayed.

- Distributed Server Load Balancing State

  Remote switch number, remote switch IP address, IP subnet mask, and health status.

- Real Server State

  Real server number, real IP address, MAC address, VLAN, physical switch port, layer where health check is performed, and health check result.

- Virtual Server State

  Virtual server number, virtual IP address, virtual MAC address

- Virtual Port State

  Virtual service or port, server port mapping, real server group, group backup server.

- Redirection Filter States

  Filter number, destination port, real server port, real server group, health check layer, group backup server, URL for health checks, and real server group, IP address, backup server, and status.

- Port State

  Physical port number, proxy IP address, type of Layer 4 activity, filter status and a list of applied filters. The Layer 4 activity type refers to the following:

  - `server` or `client` for Server Load Balancing
  - `redir` for Application Redirection
  - `failover` for connection to hot-standby switch
  - `none` for ports that do not use Layer 4 switching features
  - `standby` for ports waiting in hot-standby mode on the standby switch

To view Server Load Balancing information, at the `Information#` prompt, enter:

```
Information# slb
```

The Server Load Balancing information is displayed.

```
Failover state:
  primary   10.10.10.1,     up, ACTIVE (this switch)
  secondary 10.10.10.2,     up, STANDBY

Distributed SLB state:
  1: 220.3.78.3,  0.0.0.0,         FAILED

Real server state:
 20: 10.10.10.20, 08:00:20:7f:6b:35, vlan 2,port 2, health 3, FAILED
 21: 10.10.10.21, 08:00:20:0a:a7:7f, vlan 2,port 2, health 4, up

Virtual server state:
  1: 10.10.10.3,  00:60:cf:40:07:8e, dname mycompany.com
    virtual ports:
    http: rport http, group 1, backup none, content /, hname www
        real servers:
          1: 10.10.10.20,   backup none, remote, FAILED
   telnet: rport telnet, group 2, backup none
        real servers:
          3: 10.10.10.8,    backup none, up

Redirection filter state:
  1: dport http, rport http, group 1, health 4, backup none, cnt /
    real servers:
     20: 10.10.10.20,     backup none, FAILED
     21: 10.10.10.21,     backup none, up
  2: dport any, rport 0, group 1, health 3, backup none
    real servers:
     20: 10.10.10.20,     backup none, FAILED
     21: 10.10.10.21,     backup none, up

Port state:
  1: 0.0.0.0,         server, filt disabled, filters: empty
  2: 0.0.0.0,         server, filt disabled, filters: empty
  3: 10.10.10.7,      client, filt disabled, filters: empty
  4: 0.0.0.0,         client, filt disabled, filters: empty
  5: 10.10.10.6,       redir, filt  enabled, filters: 1 2
  6: 0.0.0.0,          redir, filt  enabled, filters: 1 2
  7: 0.0.0.0,           none, filt disabled, filters: empty
  8: 0.0.0.0,           none, filt disabled, filters: empty
  9: 0.0.0.0,       failover, filt disabled, filters: empty

>> Information#
```

Forwarding Database Information Menu

Direct command: **/info/fdb**

The forwarding database (FDB) contains information that maps the media access control (MAC) address of each known device to the appropriate switch port on which the device address was learned. The FDB also shows which other ports have seen frames destined for a particular MAC address.

NOTE – The master forwarding database supports up to 8192 MAC address entries per switch. Each switch port supports up to 4096 entries.

To access the Forwarding Database Menu, at the `Information#` prompt, enter:

```
Information# fdb
```

The Forwarding Database Menu is displayed:

```
[Forwarding Database Menu]
        find  - Show a single FDB entry by MAC address
        port  - Show FDB entries on a single port
        vlan  - Show FDB entries on a single VLAN
        refpt - Show FDB entries referenced by a single port
        dump  - Show all FDB entries
        stats - Show FDB statistics

>> Forwarding Database#
```

Show all FDB entries

Direct command: **/info/fdb/dump**

To show all FDB entries, at the Forwarding Database# prompt, enter:

```
Forwarding Database# dump
```

The current FDB information is displayed:

```
   MAC Address     VLAN  Port  State  Referenced from Ports...
00:a0:24:76:be:90     1    1    FWD    1 4
08:00:20:0a:a7:7f     1    2    FWD    2 3
08:00:20:73:b6:29     1    1    FWD    1 2
08:00:20:82:4d:8d     1    3    FWD    3 4
08:00:20:8a:54:2b     1         UNK    1

>> Forwarding Database#
```

An address that is in the forwarding (FWD) state, means that it has been learned by the switch. When in the trunking (TRK) state, the port field represents the trunk group number. If the state for the port is listed as unknown (UNK), the MAC address has not yet been learned by the switch, but has only been seen as a destination address. When an address is in the unknown state, no outbound port is indicated, although ports which reference the address as a destination will be listed under "Reference from ports."

Show a single FDB entry by MAC address

Direct command: **/info/fdb/find** *MAC-address*

To view information for a particular FDB entry, at the Forwarding Database# prompt, enter:

```
Forwarding Database# find
```

You are prompted to enter the MAC address of the device. Enter the MAC address using the format, xx:xx:xx:xx:xx:xx. For example, 08:00:20:12:34:56.

You can also enter the MAC address using the format, xxxxxxxxxxxx. For example, 080020123456.

Show FDB entries on a single port

Direct command: **/info/fdb/port** *port-number*

To show the FDB entries for a particular port, at the `Forwarding Database#` prompt, enter:

```
Forwarding Database# port port-number
```

Show FDB statistics

Direct command: **/info/fdb/stats**

To show Forwarding Database statistics, at the `Forwarding Database#` prompt, enter:

```
Forwarding Database# stats
```

Clearing entries from the Forwarding Database

To delete a MAC address from the FDB or to clear the entire FDB refer to "The Forwarding Database Menu" on page 10-5.

IP Information

Direct command: **/info/ip**

IP information includes:

- IP interface information: Interface number, IP address, subnet mask, broadcast address, VLAN number, and operational status.
- Default gateway information: Metric for selecting which configured gateway to use, gateway number, IP address, and health status
- IP forwarding information: Enable status, `lnet` and `lmask`
- Port status
- RIP1 information: enable status, update period, and active modes
- DNS information: primary and secondary DNS IP address, and default domain name.

To view IP information, at the `Information#` prompt, enter:

```
Information# ip
```

The IP information is displayed:

```
Interface information:
  1: 10.10.10.52,   255.255.255.0,   10.10.10.255,  vlan 1, up

Default gateway information: metric strict
  1: 10.10.10.226,  up

Current IP forwarding settings:
  OFF, lnet 0.0.0.0, lmask 0.0.0.0

Current IP port settings:
  1: ON
  2: ON
  3: ON
  4: ON
  5: ON
  6: ON
  7: ON
  8: ON
  9: ON

Current RIP settings:
  ON, update 30, LISTEN, DEFAULT, STATIC
  split horizon with poisoned reverse

Current DNS settings:
  10.10.10.200, 10.10.10.254, mycompany.com

>> Information#
```

IP Routing Information Menu

Direct command: **`/info/route`**

Routing information displays the following for each configured or learned route:

- Route destination IP address, subnet mask, and gateway address
- Type of route
- Tag indicating origin of route
- Metric for RIP tagged routes, specifying the number of hops to the destination (1-15 hops, or 16 for infinite hops)
- The IP interface that the route uses

To access the IP Routing Menu, at the `Information#` prompt, enter:

```
Information# route
```

The IP Routing Menu is displayed:

```
[IP Routing Menu]
        find  - Show a single route by destination IP address
        gw    - Show routes to a single gateway
        type  - Show routes of a single type
        tag   - Show routes of a single tag
        if    - Show routes on a single interface
        dump  - Show all routes

>> IP Routing#
```

You can display all IP routes currently held in the switch, or a portion according to one of the parameters listed on the menu.

Show All Routes

Direct command: **`/info/route/dump`**

To show all IP routes configured in the switch, at the `IP Routing#` prompt, enter:

```
IP Routing# dump
```

The IP route information is displayed:

```
  Destination         Mask          Gateway        Type      Tag     Mc If
--------------- --------------- --------------- --------- --------- -- --
0.0.0.0         0.0.0.0         205.178.13.226  indirect  static       1
0.0.0.0         255.0.0.0       0.0.0.0         martian   martian
10.0.0.0        255.0.0.0       205.178.13.15   indirect  rip        2  1
127.0.0.0       255.0.0.0       0.0.0.0         martian   martian
192.192.0.0     255.255.255.0   205.178.13.247  indirect  rip        2  1
192.192.192.0   255.255.255.0   205.178.13.2    indirect  rip        2  1
205.178.13.0    255.255.255.0   205.178.13.52   direct    fixed        1
205.178.13.52   255.255.255.255 205.178.13.52   local     addr         1
205.178.13.255  255.255.255.255 205.178.13.255  broadcast broadcast    1
205.178.14.0    255.255.255.0   205.178.13.204  indirect  rip        2  1
208.200.21.0    255.255.255.0   205.178.13.226  indirect  rip        2  1
224.0.0.0       224.0.0.0       0.0.0.0         martian   martian
255.255.255.255 255.255.255.255 255.255.255.255 broadcast broadcast

>> IP Routing#
```

The following table describes the `Type` parameters.

Table 5-2 IP Routing Type Parameters

Parameter	Description
`indirect`	The next hop to the host or subnet destination will be forwarded through a router at the `Gateway` address.
`direct`	Packets will be delivered to a destination host or subnet attached to the switch.
`local`	Indicates a route to one of the switch's IP interfaces.
`broadcast`	Indicates a broadcast route.
`martian`	The destination belongs to a host or subnet which is filtered out. Packets to this destination are discarded.
`multicast`	Indicates a multicast route.

The following table describes the `Tag` parameters.

Table 5-3 IP Routing Tag Parameters

Parameter	Description
`fixed`	The address belongs to a host or subnet attached to the switch.
`static`	The address is a static route which has been configured on the switch.
`icmp`	The address was learned via ICMP.
`snmp`	This address was configured through SNMP.
`addr`	The address belongs to one of the switch's IP interfaces.
`rip`	The address was learned by the Routing Information Protocol (RIP).
`broadcast`	Indicates a broadcast address.
`martian`	The address belongs to a filtered group.
`multicast`	Indicates a multicast address.

ARP Information Menu

Direct command: **/info/arp**

To access the Address Resolution Protocol (ARP) Menu, at the `Information#` prompt, enter:

```
Information# arp
```

The Address Resolution Protocol Menu is displayed:

```
[Address Resolution Protocol Menu]
        find  - Show a single ARP entry by IP address
        port  - Show ARP entries on a single port
        vlan  - Show ARP entries on a single VLAN
        refpt - Show ARP entries referenced by a single port
        dump  - Show all ARP entries

>> Address Resolution Protocol#
```

You can display all ARP entries currently held in the switch, or a portion according to one of the parameters listed on the menu.

Show All ARP Entries

Direct command: **/info/arp/dump**

The ARP information includes the following:

- IP address and MAC address of each entry
- Address status flag (see below)
- The VLAN and Port to which the address belongs
- The ports which have referenced the address (empty if no port has routed traffic to the IP address shown).

To show all ARP entries held in the switch, at the `Address Resolution Protocol#` prompt, enter:

```
Address Resolution Protocol# dump
```

The ARP information is displayed:

```
     IP address     Flags    MAC address    VLAN Port Referenced ports
  --------------- ----- ----------------- ---- ---- ----------------
  205.178.13.41      P 4 00:60:cf:40:07:81        1   1-9
  205.178.13.54      P 4 00:60:cf:40:07:8e        1   1-9
  205.178.13.163         00:a0:c9:89:b9:1f    2   2   empty
  205.178.13.168         00:a0:c9:4b:9e:90    2   2   empty
  205.178.13.176         00:60:08:93:e4:c0    2   2   empty
  205.178.13.184         00:60:08:c5:35:d7    2   2   empty
  205.178.13.220    P    08:00:87:0b:de:15    2       1-9
  205.178.13.223         00:60:cf:20:01:68    2   2   empty
  205.178.13.226         08:00:20:0a:a7:7f    2   2   empty
  205.178.13.235         08:00:20:7f:6b:35    2   2   empty

>> Address Resolution Protocol#
```

The `Flag` field is interpreted as follows:

Table 5-4 ARP Dump Flag Parameters

Flag	Description
P	Permanent entry created for switch IP interface.
P 4	Permanent entry created for Layer 4 proxy IP address or virtual server IP address.
R	Indirect route entry.
U	Unresolved ARP entry. The MAC address has not been learned.

Trunk Group Information

Direct command: **/info/trunk**

When trunk groups are configured, you can view the state of each port in the various trunk groups using the following command at the Information# prompt:

```
Information# trunk
```

The trunk group information is displayed:

```
Trunk group 1 port state:
  5: forwarding
  6: DOWN
  7: forwarding

Trunk group 2 port state:
  1: BLOCKING
  3: DOWN
  4: BLOCKING

Information#
```

NOTE – If Spanning-Tree Protocol on any port in the trunk group is set to `forwarding`, the remaining ports in the trunk group will also be set to `forwarding`.

Enabled Software Keys

Direct command: **/info/swkey**

You can display a list of all the optional software packages which have been activated or installed on your switch. At the `Information#` prompt, enter:

```
>> Information# swkey
```

For optional Layer 4 switching software, the information would be displayed as follows:

```
Enabled Software features:
  Layer 4: SLB + WCR
  Layer 4: DSLB

>> Information#
```

CHAPTER 6

The Statistics Menu

You can view switch performance statistics in both the user and administrator command modes. This chapter discusses how to use the command-line interface to display switch statistics.

To access the Statistics Menu, enter the following command at the `Main#` prompt:

```
Main# stats
```

The Statistics Menu is displayed:

```
[Statistics Menu]
        port  - Statistics Menu for one port
        if    - IP interface ("if") statistics
        ip    - IP statistics
        icmp  - ICMP statistics
        tcp   - TCP statistics
        udp   - UDP statistics
        snmp  - SNMP statistics
        route - Route statistics
        arp   - ARP statistics
        dns   - DNS server statistics
        slb   - Server Load Balancing statistics

>> Statistics#
```

Each of these options is discussed in greater detail in the following sections.

NOTE – The sample screens shown in this chapter represent ACEswitch 180 information. Screens, menus, and parameters for other Alteon Networks switches may be slightly different.

Port Statistics

Direct command: **`/stats/port`** *port-number*

The ACElerate software provides traffic statistics on a port-by-port basis. Traffic statistics include SNMP Management Information Base (MIB) objects from five groups:

- Bridging (dot1)
- Ethernet (dot3)
- Interface (if)
- Internet Protocol (IP)
- Link
- RMON

To view traffic statistics for a port, at the `Statistics#` prompt, enter:

```
Statistics# port port-number
```

The Port Statistics Menu is displayed:

```
[Port Statistics Menu]
        brg   - Bridging ("dot1") stats
        ether - Ethernet ("dot3") stats
        if    - Interface ("if") stats
        ip    - Internet Protocol ("IP") stats
        link  - Link stats
        rmon  - RMON stats

>>Port Statistics#
```

Select the type of statistics you want to see for the port by entering the appropriate command from the Port Statistics Menu.

IP Interface (IF) Statistics

Direct command: **`/stats/if`**

To display interface statistics for the management processors, at the `Statistics#` prompt, enter:

```
Statistics# if
```

Protocol Statistics

You can display switch management processor statistics for the following protocols:

- IP
- ICMP
- TCP
- UDP
- SNMP
- Route
- ARP
- DNS

To display statistics for a particular protocol, at the `Statistics#` prompt, enter the name of the protocol (**`IP`**, **`ICMP`**, **`TCP`**, **`UDP`**, **`SNMP`**, **`Route`**, **`ARP`** or **`DNS`**).

```
Statistics# protocol
```

Server Load Balancing Statistics

Direct command: **`/stats/slb`**

You can display the following Server Load Balancing Statistics:

- Real server statistics
- Virtual server statistics
- Filter statistics
- Switch port statistics
- Maintenance statistics

When the `Information#` prompt is displayed, enter:

```
Statistics# slb
```

The SLB Statistics Menu is displayed:

```
[Server Load Balancing Statistics Menu]
        real  - Real server stats
        group - Real server group stats
        virt  - Virtual server stats
        filt  - Filter stats
        port  - SLB switch port stats
        dist  - Distributed SLB stats
        maint - Maintenance stats

>> Server Load Balancing Statistics#
```

Real Server Statistics

Direct command: **/stats/slb/real** *real-server-number*

Real server statistics include the following:

- Number of sessions currently open on the real server
- Total sessions on the real server
- Number of times the real server has failed its health checks

To view these statistics, from the `Server Load Balancing Statistics#` prompt, enter:

```
Server Load Balancing Statistics# real real-server-number
```

The statistics for the real server you entered are displayed:

```
Real server 1 stats:
Current sessions:                   129
Total sessions:                     654
Health check failures:                0

>> Server Load Balancing Statistics#
```

Real Server Group Statistics

Direct command: **/stats/slb/group** *real-server-group-number*

Real server group statistics include the following:

- Current and total sessions for each real server in the real server group
- Current and total sessions for all real servers associated with the real server group

To view these statistics, from the `Server Load Balancing Statistics#` prompt, enter:

```
Server Load Balancing Statistics# group real-server-group-number
```

The statistics for the real server group you entered are displayed:

```
Real server group 1 stats:
  Real server IP address      Current Sessions  Total Sessions
  ----------- --------------- ----------------  --------------
        1     200.100.10.14             20              60
        2     200.100.10.15             20              77
  ----------- --------------- ----------------  --------------
     Totals                               40             137

>> Server Load Balancing Statistics#
```

Virtual Server Statistics

Direct command: **/stats/slb/virt** *virtual-server-number*

Virtual server statistics include the following:

- Current and total sessions for each real server associated with the virtual server
- Current and total sessions for all real servers associated with the virtual server

To view these statistics, from the `Server Load Balancing Statistics#` prompt, enter:

```
Server Load Balancing Statistics# virt virtual-server-number
```

The statistics for the virtual server you entered are displayed:

```
Virtual server 1 stats:
  Real server IP address      Current Sessions  Total Sessions
  ----------- --------------- ----------------  --------------
        1     200.100.10.14             20              60
        2     200.100.10.15             20              77
  ----------- --------------- ----------------  --------------
     Totals   200.100.10.1               40             137

>> Server Load Balancing Statistics#
```

NOTE – The virtual server IP address is shown in the "Totals" area below the real server IP addresses.

Filter Statistics

Direct command: **`/stats/slb/filt`** *filter-number*

You can obtain the total number of times any filter has been used. To view these statistics, from the `Server Load Balancing Statistics#` prompt, enter:

```
Server Load Balancing Statistics# filt filter-number
```

The statistics for the filter you entered are displayed:

```
Filter 1 stats:
Total firings:                      1011

>> Server Load Balancing Statistics#
```

SLB Switch Port Statistics Menu

Direct command: **/stats/slb/port** *port-number*

To view SLB port statistics, from the `Server Load Balancing Statistics#` prompt, enter:

```
Server Load Balancing Statistics# port port-number
```

The SLB Port Statistics Menu is displayed:

```
[Server Load Balancing Port Statistics Menu]
        real  - Real server stats
        group - Real server group stats
        virt  - Virtual server stats
        filt  - Filter stats
        maint - Maintenance stats

>> Server Load Balancing Port Statistics#
```

SLB Port Real Server Statistics

To view port statistics regarding the real server, from the `Server Load Balancing Port Statistics#` prompt, enter:

```
Server Load Balancing Port Statistics# real real-server-number
```

The port statistics for the real server you entered are displayed:

```
Port 1 Real server 1 stats:
Current sessions:                    9
Total sessions:                     24
Health check failures:               0

>> Server Load Balancing Port Statistics#
```

SLB Port Real Server Group Statistics

To view port statistics regarding a real server group, from the `Server Load Balancing Port Statistics#` prompt, enter:

```
Server Load Balancing Port Statistics# group real-server-group-number
```

The port statistics for the real server group you entered are displayed:

```
Port 1 Real server group 1 stats:
  Real server IP address      Current Sessions  Total Sessions
  ----------- --------------- ----------------  ---------------
       20     200.100.10.14              9              24
       21     200.100.10.15             12              23
  ----------- --------------- ----------------  ---------------
     Totals                              21              47

>> Server Load Balancing Port Statistics#
```

SLB Port Virtual Server Statistics

To view port statistics regarding the virtual server, from the `Server Load Balancing Port Statistics#` prompt, enter:

```
Server Load Balancing Port Statistics# virt virtual-server-number
```

The port statistics for the virtual server you entered are displayed:

```
Port 1 Virtual server 1 stats:
  Real server IP address        Current Sessions  Total Sessions
  ----------- ----------------- ----------------  ---------------
       20     200.100.10.14              9              24
       21     200.100.10.15             12              23
  ----------- ----------------- ----------------  ---------------
     Totals   200.100.13.1                21              47

>> Server Load Balancing Port Statistics#
```

NOTE – The virtual server IP address is shown in the "Totals" area below the real server IP addresses.

SLB Port Filter Statistics

You can obtain the total number of times any filter has been on a specific port. To view these statistics, from the `Server Load Balancing Statistics#` prompt, enter:

```
Server Load Balancing Port Statistics# filt filter-number
```

The statistics for the filter you entered are displayed:

```
Filter 1 stats:
Total firings:                        1011

>> Server Load Balancing Statistics#
```

SLB Port Maintenance Statistics

To view SLB port maintenance statistics, from the `Server Load Balancing Port Statistics#` prompt, enter:

```
Server Load Balancing Port Statistics# maint
```

The SLB Maintenance statistics are displayed:

```
Port 1 SLB Maintenance stats:
Current bindings:                     0
Binding failures:                     0
Non TCP/IP frames:                    0
TCP fragments:                        0
UDP datagrams:                        0
Incorrect VIPs:                       0
Incorrect Vports:                     0
No available real server:             0
Filtered (denied) frames:             0

Server Load Balancing Port Statistics#
```

These statistics are described in Table 6-1 on page 6-13.

Distributed SLB Statistics

Direct command: **/stats/slb/dist**

To view distributed server load balancing statistics, enter the following command at the `Server Load Balancing Statistics#` prompt:

```
Server Load Balancing Statistics# dist
```

The Distributed SLB Statistics Menu is displayed:

```
[Distributed SLB Statistics Menu]
        real  - Real server distributed stats
        group - Real server group distributed stats
        virt  - Virtual server distributed stats
        maint - Distributed maintenance stats

>> Distributed SLB Statistics#
```

Real Server Distributed Statistics

For any remote real server configured for Distributed Server Load Balancing, the following statistics can be viewed:

- Number of DNS handoffs to the remote server
- Number of HTTP redirects to the remote server

To view these statistics, from the `Distributed SLB Statistics#` prompt, enter:

```
Distributed SLB Statistics# real real-server-number
```

Where the real server number represents the real server ID on this switch, under which the remote server is configured. The statistics for the remote real server are displayed:

```
Real server 1 distributed stats:
DNS handoffs:                       3210
HTTP redirects:                       12

>> Distributed SLB Statistics#
```

Real Server Group Distributed Statistics

Real server group distributed statistics include the following:

- Number of DNS handoffs to each remote real server in the group
- Number of HTTP redirects to each remote real server in the group
- Total DNS handoffs and HTTP redirects to the remote real servers in the group

To view these statistics, from the `Distributed SLB Statistics#` prompt, enter:

```
Distributed SLB Statistics# group real-server-group-number
```

The statistics for the real server group you entered are displayed:

```
Real server group 1 distributed stats:
  Real server IP address          DNS Handoffs  HTTP Redirects
  ----------- --------------- ----------------  ---------------
        1     205.178.13.54           1240              30
        2     205.178.13.223           608              12
  ----------- --------------- ----------------  ---------------
     Totals                           1848              42

>> Distributed SLB Statistics#
```

Virtual Server Distributed Statistics

Virtual server distributed statistics include the following:

- Service: type of service running on the virtual server
- Server: type of server configuration and server ID number.
  - **v**# represents a local virtual server number
  - **r**# represents a remote site. Since each remote sites is configured on its peers as if it were a real server (with certain special properties), the number represents the real server ID on this switch, under which the remote server is configured.
- IP address of the server
- Response time: the average time (present weighted) that each service takes to respond to information exchanges with its peers. The time is specified in ticks of 65 milliseconds.
- Min sessions available: the current number of sessions available for serving client requests. This number will change as client traffic loads change, or as real servers under the virtual server or remote sites go in or out of service.

To view these statistics, from the `Distributed SLB Statistics#` prompt, enter:

```
Distributed SLB Statistics# virt virutal-server-number
```

The statistics for the virtual server you entered are displayed:

```
Virtual server 1 distributed stats:
  Service Server IP address      Response time Min sessions avail
  ------- ------ --------------- ------------- ------------------
  http     v1    205.178.13.55           16            21190
  http     r1    205.178.13.54           10            24120

  telnet   v1    205.178.13.55            4            31032

>> Distributed SLB Statistics#
```

Distributed SLB Maintenance Statistics

Distributed SLB mainteneance statistics include the following:

- The number of Distributed Site State Protocol (DSSP) updates received from remote sites.
- The number of bad DSSP updates received from remote sites. Bad updates usually indicate that there is a DSLB switch configuration problem. If bad updates occur, check your `syslog` for configuration error messages.

To view these statistics, from the `Distributed SLB Statistics#` prompt, enter:

```
Distributed SLB Statistics# maint
```

The maintenance statistics are displayed as follows:

```
Distributed SLB maintenance stats:
Updates received:              0
Bad updates received:          0

>> Distributed SLB Statistics#
```

SLB Maintenance Statistics

Direct command: **/stats/slb/maint**

SLB Maintenance statistics can be viewed from the Server Load Balancing Statistics# prompt. At the prompt, enter:

```
Server Load Balancing Statistics# maint
```

The SLB Maintenance statistics are displayed.

```
SLB Maintenance stats:
Current bindings:                    0
Binding failures:                    0
TCP fragments:                       0
UDP datagrams:                       0
Non TCP/IP frames:                   0
Incorrect VIPs:                      0
Incorrect Vports:                    0
No available real server:            0
Backup server activations:           0
Overflow server activations:         0
Filtered (denied) frames:            0

Server Load Balancing Statistics#
```

SLB Maintenance statistics are described in the following table.

Table 6-1 Server Load Balancing Maintenance Statistics

Statistic	Description
Current Bindings	Number of session bindings currently in use.
Binding Failure	Indicates instances where the switch ran out of available bindings for a port.
TCP Fragments	Indicates the number of TCP fragments encountered by the switch. Layer 4 processing might not handle TCP fragments, depending on configuration.
UDP Datagrams	Indicates that the virtual server IP address and MAC are receiving UDP frames when UDP balancing is not turned on.
Non TPC/IP Frames	Indicates the number of non-IP based frames received by the virtual server.
Incorrect VIPs	This indicates the number of times the switch has received a Layer 4 request for a virtual server which was not configured.

Table 6-1 Server Load Balancing Maintenance Statistics

Statistic	Description
Incorrect Vports	This dropped frames counter indicates that the virtual server has received frames for TCP/UDP services that have not been configured. Normally this indicates a mis-configuration on the virtual server or the client, but it may be an indication of a potential security probing application like SATAN.
No Server Available	This dropped frames counter indicates that all real servers are either out of service or at their `mcon` limit.
Backup Server Activations	This indicates the number of times a real server failure has occurred and caused a backup server to be brought online.
Overflow Server Activations	This indicates the number of times a real server has reached the `mcon` limit and caused an overflow server to be brought online.
Filtered (Denied) Frames	This indicates the number of frames that where dropped because they matched an active filter with the "`deny`" action set.

CHAPTER 7

The Configuration Menu

You must be logged in using the administrator password before you can access the configuration menus. This chapter discusses how to use command-line interface for making, viewing, and saving switch configuration changes.

To access the Configuration Menu, at the `Main#` prompt, enter:

```
Main# cfg
```

The Configuration Menu is displayed:

```
[Configuration Menu]
        sys   - System-wide parameter menu
        port  - Port configuration menu
        ip    - IP configuration menu
        vlan  - VLAN configuration menu
        stp   - Spanning Tree menu
        snmp  - SNMP menu
        setup - Step by step configuration set up
        dump  - Dump current configuration to script file
        mirr  - Mirroring menu
        slb   - Server Load Balancing configuration menu
        trunk - Trunk Group configuration menu

>> Configuration#
```

Each of these options is discussed in greater detail in the following sections.

NOTE – The sample screens shown in this chapter represent ACEswitch 180 information. Screens, menus, and parameters for other Alteon Networks switches may be slightly different.

Viewing, Applying, and Saving Changes

As you use the configuration menus to set switch parameters, the changes you make do not take effect immediately. All changes are considered "pending" until you explicitly apply them. Also, any changes are lost the next time the switch boots unless the changes are explicitly saved.

While configuration changes are in the pending state, you can do the following:

- View the pending changes
- Apply the pending changes
- Save the changes to flash memory

Viewing Pending Changes

You can view all pending configuration changes by entering **`diff`** at the menu prompt.

NOTE – The `diff` command is a global command. Therefore, you can enter **`diff`** at any prompt in the CLI.

Applying Pending Changes

To make your configuration changes active, you must apply them. To apply configuration changes, enter **`apply`** at any prompt in the CLI.

```
# apply
```

NOTE – The `apply` command is a global command. Therefore, you can enter **`apply`** at any prompt in the administrative interface.

NOTE – All configuration changes take effect immediately when applied, except for starting Spanning-Tree Protocol. To turn STP on or off, you must apply the changes, save them (see below), and then reset the switch (see "Resetting the Switch" on page 9-4).

Saving the Configuration

In addition to applying the configuration changes, you can save them to flash memory on the switch.

NOTE – If you do not save the changes, they will be lost the next time the system is rebooted.

To save the new configuration, enter the following command at any CLI prompt:

```
# save
```

When you save configuration changes, the changes are saved to the *active* configuration block. The configuration being replaced by the save is first copied to the *backup* configuration block. If you do not want the previous configuration block copied to the backup configuration block, enter the following instead:

```
# save noback
```

You can decide which configuration you want to run the next time you reset the switch. Your options include:

- The active configuration block
- The backup configuration block
- Factory default configuration

For instructions on selecting the configuration to run at the next system reset, see "Selecting a Configuration Block" on page 9-4."

Configuring System Parameters

Direct command: **/cfg/sys**

System parameters affect the operation of the switch globally. System parameters that can be modified include:

- System date and time
- User and Administrator passwords
- Idle timeout for CLI sessions
- Allow/disallow Telnet connections (from local console only)
- BOOTP usage
- Allow/disallow ACEvision web-based connections and set the web-server port number (from local console or telnet only)

To modify system parameters, at the Configuration# prompt, enter:

```
Configuration# sys
```

The System Menu is displayed:

```
[System Menu]
        date  - Set system date
        time  - Set system time
        usrpw - Set user password
        admpw - Set administrator password
        idle  - Set timeout for idle CLI sessions
        bootp - Enable/disable use of BOOTP
        http  - Enable/disable HTTP (Web) access
        wport - Set Web server port number
        cur   - Display current system-wide parameters

>> System#
```

The following table describes the System Menu options.

Table 7-1 System Options (/cfg/sys)

Option	Description
`date`	Configures the system date.
`time`	Configures the system time using a 24-hour clock format.
`usrpw`	Configures the user password; the user password can have a maximum of 15 characters.
`admpw`	Configures the administrator password; the administrator password can have a maximum of 15 characters.
`idle`	Configures the idle timeout for command-line interface sessions; the range is 1 to 60 minutes. The default is 5 minutes.
`telnet`	Enable or disable telnet access to the command-line interface sessions. This command is available only from a local console connection.
`bootp`	Enable or disable the use of BOOTP; if you enable BOOTP, the switch will query its BOOTP server for all of the switch IP parameters.
`http`	Enable or disable access to the web-based interface (ACEvision).
`wport`	Set the switch port to be used for serving switch web content (ACEvision). The default is HTTP port 80. If Distrubuted Server Load Balancing is to be used, change this parameter to use a different port (such as 8080).
`cur`	Displays current system parameters.

NOTE – By default, when `http` (above) is enabled, ACEvision serves web-interface content over port 80 on the switch IP interface. If Distributed Server Load Balancing is enabled, it also uses port 80 for peer health and performance updates. Both services cannot use the same port. If using ACEvision and DSLB together on the switch, use `wport` (above) to configure ACEvision for a different service port.

Configuring Port Parameters

Direct command: **/cfg/port** *port-number*

The Port Menu allows you to configure the settings for individual switch ports. To configure a port, at the Configuration# prompt, enter:

```
Configuration# port port-number
```

The Port Menu is displayed:

```
[Port 1 Menu]
        pref  - Set preferred phy
        back  - Set backup phy
        fast  - Fast phy menu
        gig   - Gig phy menu
        dis   - Disable port
        ena   - Enable port
        rmon  - Enable/Disable RMON for port
        tag   - 1 if port uses VLAN tagging, else 0 for untagged
        pvid  - Set default port VLAN id
        name  - Set port name
        cur   - Display current port configuration

>> Port 1#
```

NOTE – The port configuration options shown are for the ACEswitch 180. If you are configuring port options for some other models of Alteon Networks switch, some of the options might not be available or will behave differently. Any important differences are noted in the text.

The port configuration options are described in the following table.

Table 7-2 Port Configuration Options (cfg/port)

Option	Description
`pref`	If dual physical connectors are available on the port, this option defines the preferred physical connector. Choices are: • Fast Ethernet Port, RJ-45 connector • Gigabit Ethernet Port, SC fiber connector (default)
`back`	If dual physical connectors are available on the port, this option defines the physical connector to use when the preferred choice fails or is unavailable. Choices are: • Fast Ethernet Port, RJ-45 connector (default) • Gigabit Ethernet Port, SC fiber connector • None
`fast`	Configure the Fast Ethernet portion of the port (if available).
`gig`	Configure the Gigabit Ethernet portion of the port (if available).
`dis`	Disables the port. (To temporarily disable a port without changing its configuration attributes, refer to "Temporarily Disabling a Port" on page 7-8.)
`ena`	Enables the port.
`rmon`	Enable/Disable RMON support on the port.
`tag`	Set to 1 if the port uses VLAN tagging. Otherwise, set to 0.
`pvid`	Set the default VLAN number which will be used to forward frames which are not VLAN tagged.
`name`	Set a name for the port. The assigned port name appears next to the port number on some information and statistics screens.
`cur`	Displays current port parameters.

The options in Table 7-3 are used for setting options such as port speed, negotiation mode, and flow control. These options appear on the `fast` and `gig` port configuration menus for the ACEswitch 180, and on the Port Menu itself for some models of Alteon Networks switches.

Table 7-3 More Port Configuration Options (/cfg/port)

Option	Description
`speed`	Sets the link speed; the choices include: • "Any," for automatic detection (default) • 10 Mbps • 100 Mbps
`mode`	Sets the operating mode; the choices include: • "Any," for autonegotiation (default) • Full-duplex • Half-duplex
`fctl`	Sets the flow control; the choices include: • Autonegotiation (default) • Receive flow control • Transmit flow control • Both receive and transmit flow control • No flow control
`auto`	Enable or disable autonegotiation for the port.
`cur`	Displays current port parameters.

Temporarily Disabling a Port

To temporarily disable a port without changing its stored configuration attributes, enter the following command at any prompt:

```
Main# /oper/port port-number/dis
```

Because this sets a temporary state for the port, you do not need to use `apply` or `save`. See "The Operations Menu" on page 8-1 for other operations-level commands.

Configuring IP Parameters

Direct command: **/cfg/ip**

The IP Menu provides access to the switch IP parameters. IP parameters are configured to provide Telnet and SNMP management access to the switch, as well as for defining routing and forwarding information.

To configure IP parameters, at the `Configuration#` prompt, enter:

```
Configuration# ip
```

The IP Menu is displayed:

```
[IP Menu]
        if    - Interface menu
        gw    - Default gateway menu
        route - Static route menu
        frwd  - Forwarding menu
        rip1  - Routing Information Protocol menu
        port  - IP port menu
        dns   - Domain Name System menu
        log   - Set IP address of syslog host
        rearp - Set re-ARP period in minutes
        metrc - Set default gateway metric
        cur   - Display current IP configuration

>> IP#
```

These commands are described in detail in the following sections.

IP Interface Menu

Direct command: **`/cfg/ip/if`** *interface-number*

The switch can be configured with up to 256 IP interfaces. Each IP interface represents the switch on an IP subnet on your network. To configure IP interfaces, enter the following command at the IP Menu:

```
IP# if interface-number
```

The IP Interface Menu for the selected interface (1 to 256) appears:

```
[IP Interface 1 Menu]
        addr  - Set IP address
        mask  - Set subnet mask
        broad - Set broadcast address
        vlan  - Set VLAN number
        ena   - Enable IP interface
        dis   - Disable IP interface
        del   - Delete IP interface
        cur   - Display current interface configuration

>> IP Interface 1#
```

The following table describes the IP Interface Menu options.

Table 7-4 IP Interface Options (/cfg/ip/if)

Option	Description
`addr`	Configures the IP address of the switch interface using dotted decimal notation.
`mask`	Configures the IP subnet address mask for the interface using dotted decimal notation.
`broad`	Configures the IP broadcast address for the interface using dotted decimal notation.
`vlan`	Configures the VLAN number for this interface. Each interface can belong to one VLAN, though any VLAN can have multiple IP interfaces in it.
`ena`	Enable the interface.
`dis`	Disable the interface.
`del`	Delete this interface.
`cur`	Display the current interface settings.

Default Gateway Settings

Default Gateway Menu

Direct command: **`/cfg/ip/gw`** *gateway-number*

The switch can be configured with up to four default IP gateways. To configure the default IP gateways, enter the following command at the IP Menu:

```
IP# gw gateway-number
```

The Default Gateway Menu for the selected gateway (1 to 4) appears:

```
[Default gateway 1 Menu]
        addr  - Set IP address
        intr  - Set interval between ping attempts
        retry - Set number of failed attempts to declare gateway DOWN
        ena   - Enable default gateway
        dis   - Disable default gateway
        del   - Delete default gateway
        cur   - Display current default gateway configuration

>> Default gateway 1#
```

The following table describes the Default Gateway Menu options.

Table 7-5 Default Gateway Options (/cfg/ip/gw)

Option	Description
`addr`	Configures the IP address of the default IP gateway using dotted decimal notation.
`intr`	The switch pings the default gateway to verify that the gateway is up. The `intr` option lets you choose the time between health checks. The range is from 1 to 120 seconds. The default interval is 2 seconds.
`retry`	Set the number of failed health check attempts required before declaring this default gateway inoperative. The range is from 1 to 120 attempts. The default is 8 attempts.
`ena`	Enable the gateway for use.
`dis`	Disable the gateway.
`del`	Delete this gateway from the configuration.
`cur`	Display the current gateway settings.

Default Gateway Metrics

Direct command: **/cfg/ip/metrc** *metric-name*

If multiple default gateways are configured and enabled, a metric can be set to determine which how the primary gateway is selected. There are two metrics:

Table 7-6 Default Gateway Metrics (/cfg/ip/metrc)

Option	Description
strict	The gateway number determines its level of preference. Gateway #1 acts as the prefered default IP gateway until it fails or is disabled, at which point the next in line will take over as the default IP gateway.
roundrobin	This provides basic gateway load balancing. The switch sends each new gateway request to the next healthy, enabled gateway in line. All gateway requests to the same destination IP address are resolved to the same gateway.

IP Static Route Menu

Direct command: **/cfg/ip/route**

To access the IP Static Route Menu, enter the following command at the IP Menu:

```
IP# route
```

The IP Static Route Menu appears:

```
[IP Static Route Menu]
        add   - Add static route
        rem   - Remove static route
        cur   - Display current static routes

>> IP Static Route Menu#
```

The following table describes the IP Static Route Menu options.

Table 7-7 IP Static Route Options (/cfg/ip/route)

Option	Description
add	Add a static route. You will be prompted to enter a destination IP address, destination subnet mask, and gateway address. Enter all addresses using dotted decimal notation.
rem	Remove a static route. The destination address of the route to remove must be specified using dotted decimal notation.
cur	Display the current IP static routes.

IP Forwarding Menu

Direct command: **/cfg/ip/frwd**

The IP Forwarding Menu is used for setting the local network address and netmask for the route cache, and to turn IP forwarding (routing) on or off. To access the menu, enter the following at the IP Menu:

```
IP# frwd
```

The IP Forwarding Menu appears:

```
[IP Forwarding Menu]
        lnet  - Set local IP network for route cache
        lmask - Set local IP netmask for route cache
        on    - Globally turn IP Forwarding ON
        off   - Globally turn IP Forwarding OFF
        cur   - Display current static routes

>> IP Forwarding Menu#
```

The following table describes the IP Forwarding Menu options.

Table 7-8 IP Forwarding Options (/cfg/ip/frwd)

Option	Description
`lnet`	Sets the base destination IP address for a range of routes which will cached on the switch. See details below.
`lmask`	This IP address mask is used with the `lnet` to identify routes which will be included in the local route cache. See details below.
`on`	Enable IP forwarding (routing).
`off`	Disable IP forwarding (routing).
`cur`	Display the current IP forwarding settings.

Defining IP Address Ranges for the Local Route Cache

The Local Route Cache lets you more efficiently use switch resources. The `lnet` and `lmask` parameters define a range of addresses which will be cached on the switch. The `lnet` is used to define the base IP address in the range which will be cached, and the `lmask` is the mask which is applied to produce the range. To determine if a route should be added to the memory cache, the destination address is masked (bitwise AND) with the `lmask` and checked against the `lnet`.

By default, the `lnet` and `lmask` are both set to 0.0.0.0. This produces a range that includes all Internet addresses for route caching: 0.0.0.0 through 255.255.255.255.

To limit the route cache to your local hosts, you could configure the parameters as in the following examples:

Table 7-9 Local Routing Cache Address Ranges

Local Host Address Range	lnet	lmask
0.0.0.0 - 127.255.255.255	0.0.0.0	128.0.0.0
128.0.0.0 - 255.255.255.255	128.0.0.0	128.0.0.0
205.32.0.0 - 205.32.255.255	205.32.0.0	255.255.0.0

NOTE – All addresses that fall outside the defined range are forwarded to the default gateway. The default gateways must be within range.

Routing Information Protocol Menu

Direct command: **`/cfg/ip/rip1`**

The RIP1 Menu is used for configuring Routing Information Protocol version 1 parameters.

NOTE – Do not configure RIP1 parameters if your routing equipment uses RIP version 2.

To configure RIP1 parameters, enter the following from the IP Menu:

```
IP# rip1
```

The Routing Information Protocol Menu appears:

```
[Routing Information Protocol Menu]
        spply - Enable/disable supplying route updates
        lsten - Enable/disable listening to route updates
        deflt - Enable/disable listening to default routes
        statc - Enable/disable supplying static routes
        poisn - Enable/disable poisoned reverse
        updat - Set update period in seconds
        on    - Globally turn RIP ON
        off   - Globally turn RIP OFF
        cur   - Display current RIP configuration

>> Routing Information Protocol Menu#
```

The following table describes the RIP1 options.

Table 7-10 Routing Information Protocol Options (/cfg/ip/rip1)

Option	Description
`spply`	When enabled, the switch supplies routes to other routers.
`lsten`	When enabled, the switch learns routes from other routers.
`deflt`	When enabled, the switch accepts RIP default routes from other routers and gives them priority over configured default gateways. When disabled, the switch rejects RIP default routes.
`statc`	When enabled, the switch supplies RIP information about any configured ***static*** routes to other routers.
`poisn`	When enabled, the switch uses split horizon with poisoned reverse. When disabled, the switch uses only split horizon.
`updat`	Specifies the time period between routing updates. The time is specified in seconds between 1 and 120. If an entry fails to be updated on four consecutive attempts, the entry is aged out of the routing table.
`on`	Enable Routing Information Protocol (RIP).
`off`	Disable RIP.
`cur`	Display the current RIP settings.

IP Port Menu

Direct command: **`/cfg/ip/port`** *port-number*

The IP Port Menu allows you to turn IP forwarding on or off on a port by port basis. To access the menu, enter the following at the IP Menu:

```
IP# port port-number
```

The IP Forwarding Port Menu appears for the selected port:

```
[IP Forwarding Port 1 Menu]
        on    - Turn IP Forwarding ON
        off   - Turn IP Forwarding OFF
        cur   - Display current port configuration

>> IP Forwarding Port 1#
```

The following table describes the IP Forwarding Port Menu options.

Table 7-11 IP Forwarding Port Options (/cfg/ip/port)

Option	Description
`on`	Enable IP forwarding for the current port.
`off`	Disable IP forwarding for the current port.
`cur`	Display the current IP forwarding settings.

Domain Name System Menu

Direct command: **/cfg/ip/dns**

The Domain Name System (DNS) Menu is used for defining the primary and secondary DNS servers on your local network, and for setting the default domain name served by the switch services. DNS parameters must be configured prior to using hostname parameters with the `ping`, `traceroute`, and `tftp` commands.

To configure Domain Name System (DNS) parameters, enter the following at the IP Menu:

```
IP# dns
```

The Domain Name System Menu appears:

```
[Domain Name System Menu]
        prima - Set IP address of primary DNS server
        secon - Set IP address of secondary DNS server
        dname - Set default domain name
        cur   - Display current DNS configuration

>> Domain Name System#
```

The following table describes the menu options.

Table 7-12 Domain Name Service Menu Options (/cfg/ip/dns)

Option	Description
`prima`	You will be prompted to set the IP address for your primary DNS server. Use dotted decimal notation.
`secon`	You will be prompted to set the IP address for your secondary DNS server. If the primary DNS server fails, the configured secondary will be used instead. Enter the IP address using dotted decimal notation.
`dname`	Set the default domain name used by the switch. For example: `mycompany.com`
`cur`	Display the current Domain Name Service settings.

Syslog Host

Direct command: **`/cfg/ip/log`** *IP-address*

This command is used for setting the IP address of the syslog host:

```
IP# log IP-address
```

The IP address is specified using dotted decimal notation. If configured, the switch software logs the following types of messages to syslog host:

Table 7-13 Syslog Host Messages

Level	Message
NOTICE	booted
NOTICE	reset from console
NOTICE	reset from Telnet
ERR	PANIC ()
ERR	VERIFY (*text*)
ERR	ASSERT (*text*)
NOTICE	admin password changed
NOTICE	syslog host changed
NOTICE	boot config block changed
NOTICE	boot image changed
INFO	new configuration applied (general)
INFO	new configuration saved (general)
INFO	new software image downloaded
INFO	Telnet login
INFO	Telnet logout
INFO	Console login
INFO	Console logout
NOTICE	PASSWORD FIX-UP MODE IN USE

Re-ARP

Direct command: **/cfg/ip/rearp** *re-arp-interval*

The switch occasionally sends ARP (Address Resolution Protocol) requests to refresh its address database. This command is used for setting the ARP interval. From the IP Menu, enter the following:

```
IP# rearp ARP-interval
```

Where *ARP-interval* is number of minutes (from 2 to 120) between re-ARPs.

Configuring VLAN Parameters

Direct command: **/cfg/vlan** *VLAN-number*

The commands in this menu configure VLAN attributes, change the status of the VLAN, delete the VLAN, and change the Port membership of the VLAN. For a more information on configuring VLANs, see "Setup Part 3: VLANs" on page 3-7, and also Chapter 11, "VLANs."

To configure VLANs, at the `Configuration#` prompt, enter:

```
Configuration# vlan VLAN-number
```

The VLAN Menu for the VLAN you selected is displayed:

```
[VLAN 1 Menu]
        name  - Assign VLAN name
        jumbo - Enable/disable Jumbo Frame support
        del   - Delete VLAN
        ena   - Enable VLAN
        dis   - Disable VLAN
        add   - Add port to VLAN
        rem   - Remove port from VLAN
        def   - Define VLAN as list of ports
        cur   - Display current VLANs

>> VLAN 1#
```

VLAN configuration options are described in the following table.

Table 7-14 VLAN Options (/cfg/vlan)

Option	Description
`name`	Assigns a name to the VLAN or changes the existing name.
`jumbo`	Enables or disables support for Jumbo Frames on this VLAN.
`del`	Deletes the VLAN.
`ena`	Enables the VLAN.
`dis`	Disables the VLAN without removing it from the configuration.
`add` *port*	Add a port to the VLAN membership.
`rem` *port*	Remove a port from the VLAN membership.
`def` *port* [*port...*]	Define the VLAN membership as a list of specified ports. To specify multiple ports, separate each port by a space. When this command is used, the existing port list is cleared, and the specified ports are added to the VLAN. Any ports not specified are removed from the VLAN.
`cur`	Displays all currently configured VLANs.

NOTE – You cannot add a port to more than one VLAN unless the port has VLAN tagging turned on (see "Configuring Port Parameters" on page 7-6).

NOTE – All ports must belong to at least one VLAN. Any port which is removed from a VLAN and which is not a member of any other VLAN is automatically added to default VLAN #1. You cannot remove a port from VLAN #1 if the port has no membership in any other VLAN.

Configuring Spanning-Tree Parameters

Direct command: **/cfg/stp**

The ACElerate software supports the IEEE 802.1d Spanning-Tree Protocol (STP). STP is used to prevent loops in the network topology.

To configure STP parameters, at the Configuration# prompt, enter:

```
Configuration# stp
```

The Spanning-Tree Menu is displayed:

```
[Spanning Tree Menu]
        brg   - Bridge parameter menu
        port  - Port parameter menu
        on    - Globally turn Spanning Tree ON
        off   - Globally turn Spanning Tree OFF
        cur   - Display current bridge parameters

>> Spanning Tree#
```

The following table describes the Spanning-Tree Menu options.

Table 7-15 Spanning-Tree Options (/cfg/stp)

Option	Description
brg	Displays the bridge parameter menu.
port	Displays the port parameter menu.
on	Globally enables STP.
off	Globally disables STP.
cur	Displays current STP parameters.

Bridge Spanning Tree Menu

Direct command: **/cfg/stp/brg**

Spanning-Tree bridge parameters affect the global STP operation of the switch. STP bridge parameters include:

- Bridge priority
- Bridge hello time
- Bridge maximum age
- Forwarding delay
- Bridge aging time

To configure STP bridge parameters, at the `Spanning-Tree#` prompt, enter:

```
Spanning Tree# brg
```

The Bridge Spanning-Tree Menu is displayed:

```
[Bridge Spanning Tree Menu]
        prior - Set bridge Priority (0-65535)
        hello - Set bridge Hello Time (1-10 secs)
        mxage - Set bridge Max Age (6-40 secs)
        fwd   - Set bridge Forward Delay (4-30 secs)
        aging - Set bridge Aging Time (1-65535 secs, 0 to disable)
        cur   - Display current bridge parameters

>> Bridge Spanning Tree#
```

Bridge Spanning-Tree Menu options are described in the following table.

Table 7-16 Bridge Spanning-Tree Options (/cfg/stp/brg)

Option	Description
`prior`	Configures the bridge priority. The bridge priority parameter controls which bridge on the network is the STP root bridge. To make this switch the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. The range is 0 to 65535, and the default is 32768.
`hello`	Configures the bridge hello time.The hello time specifies how often the root bridge transmits a configuration bridge protocol data unit (BPDU). Any bridge that is not the root bridge uses the root bridge hello value. The range is 1 to 10 seconds, and the default is 2 seconds.
`mxage`	Configures the bridge maximum age. The maximum age parameter specifies the maximum time the bridge waits without receiving a configuration bridge protocol data unit before it reconfigures the STP network. The range is 6 to 40 seconds, and the default is 20 seconds.
`fwd`	Configures the bridge forward delay parameter. The forward delay parameter specifies the amount of time that a any bridge port has to wait before it changes from learning state to forwarding state. The range is 4 to 30 seconds, and the default is 15 seconds.
`aging`	Configures the forwarding database aging time. The aging time specifies the amount of time the bridge waits without receiving a packet from a station before removing the station from the forwarding database. The range is 1 to 65535 seconds, and the default is 300 seconds. To disable aging, set this parameter to 0.
`cur`	Displays current bridge STP parameters.

When configuring STP bridge parameters, the following formulas must be followed:

- $2*(fwd-1) \leq mxage$
- $2*(hello+1) \leq mxage$

Spanning-Tree Port Menu

Direct command: **/cfg/stp/port** *port-number*

Spanning-Tree port parameters are used to modify STP operation on an individual port basis. STP port parameters include:

- Port priority
- Port path cost

To configure STP port parameters, at the Spanning Tree# prompt, enter:

```
Spanning Tree# port port-number
```

The Spanning-Tree Port Menu is displayed:

```
[Spanning Tree Port 1 Menu]
        on     -  Turn port's Spanning Tree ON
        off    -  Turn port's Spanning Tree OFF
        prior  -  Set port Priority (0-255)
        cost   -  Set port Path Cost (10-65535)
        cur    -  Display current port Spanning Tree parameters

>> Spanning Tree Port 1#
```

The Spanning-Tree Port Menu options are described in the following table.

Table 7-17 Spanning-Tree Port Options (/cfg/stp/port)

Option	Description
on	Enables STP on the port.
off	Disables STP on the port.
prior	Configures the port priority. The port priority helps determine which bridge port becomes the designated port. In a network topology that has multiple bridge ports connected to a single segment, the port with the lowest port priority becomes the designated port for the segment. The range is 0 to 255, and the default is 128.
cost	Configures the port path cost. The port path cost is used to help determine the designated port for a segment. Generally speaking, the faster the port, the lower the path cost. The range is 1 to 65535. The default is 10 for 100Mbps ports, and 1 for gigabit ports. A value of 0 indicates that the default cost will be computed for an autonegotiated link speed.
cur	Displays current STP port parameters.

Configuring SNMP Parameters

Direct command: **`/cfg/snmp`**

The ACElerate software supports SNMP-based network management. If you are running an SNMP network management station on your network, you can manage the switch using the following standard SNMP MIBs:

- MIB II (RFC 1213)
- Ethernet MIB (RFC 1643)
- Bridge MIB (RFC 1493)
- RMON MIB (RFC 1757)

SNMP parameters that can be modified include:

- System name
- System location
- System contact
- Use of the SNMP system authentication trap function
- Read community string
- Write community string
- Trap hosts
- Trap community strings

To configure SNMP parameters, at the `Configuration#` prompt enter:

```
Configuration# snmp
```

The SNMP Menu is displayed:

```
[SNMP Menu]
        name  - Set SNMP "sysName"
        locn  - Set SNMP "sysLocation"
        cont  - Set SNMP "sysContact"
        auth  - Disable/enable SNMP "sysAuthenTrap"
        rcomm - Set SNMP read community string
        wcomm - Set SNMP write community string
        trap1 - Set IP addr of first SNMP trap host
        trap2 - Set IP addr of second SNMP trap host
        t1comm - Set community string for first trap host
        t2comm - Set community string for second trap host
        linkt - Disable/enable SNMP link up/down trap
        cur   - Display current SNMP information

>> SNMP#
```

The SNMP Menu options are described in the following table.

Table 7-18 SNMP Options (/cfg/snmp)

Option	Description
`name`	Configures the name for the system. The name can have a maximum of 64 characters.
`locn`	Configures the name of the system location. The system location can have a maximum of 64 characters.
`cont`	Configures the name of the system contact. The system contact can have a maximum of 64 characters.
`auth`	Enables or disables the use of the system authentication trap facility. The default setting is disabled.
`rcomm`	Configures the SNMP read community string. The read community string controls SNMP "get" access to the switch. It can have a maximum of 32 characters.
`wcomm`	Configures the SNMP write community string. The write community string controls SNMP "set" and "get" access to the switch. It can have a maximum of 32 characters.
`trap1`	Configures the IP address of the first SNMP trap host using dotted decimal notation. The SNMP trap host is the device that receives SNMP trap messages from the switch.
`trap2`	Configures the IP address of the second SNMP trap host using dotted decimal notation.
`t1comm`	Configures the community string for the first trap host.
`t2comm`	Configures the community string for the second trap host.
`linkt`	Enables or disables the sending of SNMP link up and link down traps.
`cur`	Displays current SNMP information.

Setup

Direct command: **`/cfg/setup`**

The setup program steps you through configuring the system date and time, `BOOTP`, IP, Spanning Tree, port, and VLAN parameters.

To start the setup program, at the `Configuration#` prompt, enter:

```
Configuration# setup
```

For a complete description of how to use Setup see Chapter 3, "First-Time Configuration."

Dump

Direct command: **/cfg/dump**

The dump program writes the current switch configuration to the terminal screen. To start the dump program, at the Configuration# prompt, enter:

```
Configuration# dump
```

The configuration is displayed in the form of a configuration script. The screen display can be captured, edited, and placed in a configuration script file.

The configuration script file can be used to configure other switches through a Telnet connection. When using Telnet to configure a new switch, paste the configuration commands from the script file at the command line prompt of the switch.

Configuring Port Mirroring

Shortcut command: **/cfg/mirr/port**

The Port Mirroring Menu is used to configure, enable, and disable the port monitor. When enabled, network packets being sent and/or received on a target port are duplicated and sent to a monitor port. By attaching a network analyzer to the monitor port, you can collect detailed information about your network performance and usage.

NOTE – Port Mirroring cannot be used simultaneously with Layer 4 services (Server Load Balancing or Application Redirection) on any switch port connected to a server either directly, or through another switch or hub.

For Server Load Balancing, this applies to any switch port configured in the "server" state. For Application Redirection, this applies to any switch port that has a cache server attached to it directly or indirectly. Use your network analyzer with a full-duplex pass-through connection or an Ethernet hub when troubleshooting a switch port for a server used for Layer 4 services.

The Port Mirroring Menu is configured from the Configuration Menu:

```
Configuration# mirr/port
```

The Port Mirroring Menu is displayed.

```
[Port Mirroring Menu]
        to    - Set "Monitoring" port
        from  - Set "Mirrored" port
        dir   - Set Direction [in, out, both]
        tmout - Set Mirroring Timeout value in seconds
        dis   - Disable Port Mirroring
        ena   - Enable Port Mirroring
        cur   - Display current Port Mirroring configuration

>> Port Mirroring#
```

The Port Mirroring Menu options are described in the following table.

Table 7-19 Port Mirroring Options (/cfg/mirr/port)

Option	Description
`to`	This defines the monitoring port. When port mirroring is enabled, packets received and/or transmitted by the mirrored port will be duplicated to the switch port specified in this command.
`from`	This defines the mirrored port. When port mirroring is enabled, packets received and/or sent by the port specified in this command will be sent to the monitor port.
`dir`	This determines which type of packets will be sent to the monitor port: `in` = packets received at the mirrored port `out` = packets sent from the mirrored port `both` = packets sent and received by the mirrored port
`tmout`	Port mirroring will be automatically disabled (regardless of port state) after the time-out period specified in this command. Valid times are from 0 (does not time-out) to 86400 seconds.
`dis`	Turns port mirroring off.
`ena`	Turns port mirroring on.
`cur`	Displays the current parameter settings.

Configuring Server Load Balancing

Direct command: **/cfg/slb**

From the Operations# prompt, enter:

```
Configuration# slb
```

The Server Load Balancing Menu is then displayed:

```
[Server Load Balancing Menu]
        real  - Real server menu
        group - Real server group menu
        virt  - Virtual server menu
        filt  - Filtering menu
        port  - Layer 4 port menu
        fail  - Layer 4 failover menu
        dist  - Distributed SLB menu
        on    - Globally turn Layer 4 processing ON
        off   - Globally turn Layer 4 processing OFF
        imask - Set virtual and real IP address mask
        mnet  - Set managment network
        mmask - Set management subnet mask
        cur   - Display current Server Load Balancing configuration

>> Server Load Balancing#
```

The following table describes the Server Load Balancing Menu options.

Table 7-20 Server Load Balancing Options (/cfg/slb)

Option	Description
real	Menu for configuring real servers.
group	Menu for placing real servers into real server groups.
virt	Menu for defining virtual servers.
filt	Menu for Filtering and Application Redirection.
port	Menu for setting physical switch port states for Layer 4 activity.
fail	Menu for setting hot-standby failover parameters for Layer 4 activity.
dist	Menu for configuring Distributed Server Load Balancing

Table 7-20 Server Load Balancing Options (/cfg/slb)

Option	Description
`on`	Turn on Layer 4 software services for Server Load Balancing and Application Redirection. This option can be performed only once the optional Layer 4 software is enabled (see "Activating Optional Software" on page 8-6). Enabling Layer 4 services is not necessary for using filters only to allow, deny, or NAT traffic (see note below).
`off`	Disable Layer 4 services. All configuration information will remain in place (if `applied` or `saved`), but the software processes will no longer be active in the switch.
`imask`	Configures the real and virtual IP address mask using dotted decimal notation. See additional notes below.
`mnet`	If defined, management traffic with this source IP address will be allowed direct (non-Layer 4) access to the real servers. Specify an IP address in dotted decimal notation. A range of IP addresses is produced when used with the `mmask` below.
`mmask`	This IP address mask is used with the `mnet` to select management traffic which is allowed direct real server access.
`cur`	Displays current system parameters.

Filtering and Layer 4

Filters configured to allow, deny, or NAT traffic do not require Layer 4 software to be activated. These filters are not affected by the Server Load Balancing `on` and `off` commands in this menu.

Application Redirection filters, however, require Layer 4 software services. Layer 4 processing must be turned on before redirection filters will work.

Configuring the imask

The imask determines how many different IP addresses each real and virtual server will represent and respond to. By default, the imask setting is 255.255.255.255, which means that each real and virtual server represents a single IP address. An imask setting of 255.255.255.0 would mean that each real and virtual server represents 256 IP addresses. For example, consider the following:

- A virtual server is configured with an IP address of 172.16.10.1
- Real servers 172.16.20.1 and 172.16.30.1 are assigned to service the virtual server.
- The imask is set to 255.255.255.0

If the client request was sent to virtual IP address 172.16.10.45, the unmasked portion of the virtual IP address (0.0.0.45) gets mapped directly to whichever real IP address is selected by the Server Load Balancing algorithm. Thus, the request would be sent to either 172.16.20.45 or 172.16.30.45.

Configuring Real Server Parameters

Direct command: **/cfg/slb/real** *real-server-number*

This menu is used for configuring information about the real servers which will participate in the server pool for Server Load Balancing or Application Redirection. The required minimum of parameters to configure is as follows:

- Real server IP address
- Enabling the real server

To configure SLB real server parameters, at the `Server Load Balancing#` prompt, enter:

```
Server Load Balancing# real real-server-number
```

Where the real-server-number (1 to 256) represents a real server that you wish to configure.

The menu for the real server you entered is displayed:

```
[Real server 1 Menu]
      rip   - Set IP addr of real server
      wght  - Set server weight
      mcon  - Set maximum number of connections
      tmout - Set minutes inactive connection remains open
      bkup  - Set backup real server
      intr  - Set interval between ping attempts
      retry - Set number of failed attempts to declare server DOWN
      restr - Set number of successful attempts to declare server UP
      remot - Enable/disable remote site operation
      ena   - Enable real server
      dis   - Disable real server
      del   - Delete real server
      cur   - Display current real server configuration

>> Real server 1 #
```

Real server configuration options are described in the following table.

Table 7-21 SLB Real Server Options (/cfg/slb/real)

Option	Description
`rip`	Set the IP address of the real server in dotted decimal format. When this command is used, the address entered is `PING`ed to determine if the server is up, and the administrator will be warned if the server does not respond.
`wght`	Set the weighting value (1 to 48) that this real server will be given in the load balancing algorithms. Higher weighting values force the server to receive more connections than the other servers configured in the same real server group. By default, each real server is given a weight setting of 1. A setting of 10 would assign the server roughly 10 times the number of connections as a server with a weight of 1. Weights are not applied when using the `hash` or `minmisses` metrics (see "Server Load Balancing Metrics" on page 7-36).
`mcon`	Set the maximum number of connections that this server should simultaneously support. This option sets a threshold as an artificial barrier, such that new connections will not be issued to this server if the `mcon` limit is reached. New connections will be issued again to this server once the number of current connections has decreased below the `mcon` setting. If all servers in a real server group for a virtual server reach their `mcon` limit at the same time, client requests will be dropped by the virtual server.
`tmout`	Set the number of minutes an inactive session remains open (in even numbered increments). Every client-to-server session being load balanced is recorded in the switch's Session Binding Table. When a client makes a request, the session is recorded in the binding table, the data is transferred until the client ends the session, and the binding table entry is then removed. In certain circumstances, such as when a client application is abnormally terminated by the client's system, TCP/UDP connections will remain registered in the switch's binding table. In order to prevent table overflow, these orphaned entries must be aged out of the binding table. Using the `tmout` option, you can set the number of minutes to wait before removing orphan table entries. Settings must be specified in even numbered increments between 2 and 60 minutes. The default setting is 10. This option is also used with the Persistent option (see `/cfg/slb/virt/pbind`). When Persistent is activated, this option sets how long an idle client is allowed to remain associated with a particular server.

Table 7-21 SLB Real Server Options (/cfg/slb/real)

Option	Description
`bkup`	Set the real server used as the backup/overflow server for this real server. To prevent loss of service if a particular real server fails, use this option to assign a backup real server number. Then, if the real server becomes inoperative, the switch will activate the backup real server until the original becomes operative again. The backup server is also used in overflow situations. If the real server reaches its `mcon` (maximum connections) limit, the backup comes online to provide additional processing power until the original server becomes desaturated. The same backup/overflow server may be assigned to more than one real server at the same time.
`intr`	Set the interval between real server health verification attempts. Determining the health of each real server is a necessary function for Layer 4 switching. For TCP services, the switch verifies that real servers and their corresponding services are operational by opening a TCP connection to each service, using the defined service ports configured as part of each virtual service. For UDP services, the switch pings servers to determine their status. The `intr` option lets you choose the time between health checks. The range is from 1 to 60 seconds. The default interval is 2 seconds.
`retry`	Set the number of failed health check attempts required before declaring this real server inoperative. The range is from 1 to 63 attempts. The default is 4 attempts.
`restr`	Set the number of successful health check attempts required before declaring a UDP service operational. The range is from 1 to 63 attempts. The default is 8 attempts.
`remot`	Disable or enable remote site operation for this server. This should be enabled when the real IP address supplied above represents a remote server (real or virtual) this switch will access as part of its Distributed Server Load Balancing network. For more information see Chapter 17, "Distributed Server Load Balancing."
`ena`	You *must* perform this command to enable this real server for Layer 4 service. When enabled, the real server can process virtual server requests associated with its real server group. This option, when the `apply` and `save` commands are used, enables this real server for operation until explicitly disabled. See `/oper/slb/ena` on page 8-5 for an operations-level command.
`dis`	Disable this real server from Layer 4 service. Any disabled server will no longer process virtual server requests as part of the real server group to which it is assigned. This option, when the `apply` is are used, disables this real server until it is explicitly re-enabled. This option *does not* perform a graceful server shutdown. See `/oper/slb/dis` on page 8-5 for an operations-level command.

Table 7-21 SLB Real Server Options (/cfg/slb/real)

Option	Description
`del`	Delete this real server from the Layer 4 switching software configuration. This removes the real server from operation within its real server groups. Use this command with caution, as it will delete any configuration options that have been set for this real server. This option *does not* perform a graceful server shutdown.
`cur`	Display the current configuration information for this real server.

The Real Server Group Menu

Direct command: **`/cfg/slb/group`** *real-server-group-number*

This menu is used for combining real servers into real server groups. Each real server group should consist of all the real servers which provide a specific service for load balancing. Each group must consist of at least one real server. Each real server can belong to more than one group. Real server groups are used both for Server Load Balancing and Application Redirection.

To configure SLB real server group options, at the `Server Load Balancing#` prompt, enter:

```
Server Load Balancing# group real-server-group-number
```

Where *real-server-group-number* (1 to 256) represents the number of the real server group that you wish to configure.

The menu for the real server group you entered is displayed:

```
[Real server group 1 Menu]
        add   - Add real server
        rem   - Remove real server
        metrc - Set metric used to select next server in group
        cntnt - Set content to health check
        healt - Set health check type
        bkup  - Set backup real server
        del   - Delete real server group
        cur   - Display current group configuration

>> Real server group 1#
```

Real server group configuration options are described in the following table.

Table 7-22 SLB Real Server Group Options (/cfg/slb/group)

Option	Description
`add`	Add a real server to this real server group. You will be prompted to enter the number (1 to 256) of the real server to add to this group.
`rem`	Remove a real server from this real server group. You will be prompted for the ID number for the real server to remove from this group.
`metrc`	Set the load balancing metric used for determining which real server in the group will be the target of the next client request. See the information below.
`cntnt`	This option defines the specific content which is examined during health checks. The content depends on the type of health check specified in the `healt` option (see below).
`healt`	Set the type of health checking performed. The options are as follows: `icmp` For Layer 3 health checking, `ping` the server. `tcp` For TCP service, open and close a TCP/IP connection to the server. `http` For HTTP service, use HTTP 1.0 `GETS` to check that the URL content specified in `cntnt` is accessible on the server. `dns` For Domain Name Service, check that the domain name specified in `cntnt` can be resolved by the server. `pop3` For user mail service, check that the *user:password* account specified in `cntnt` exists on the server. `smtp` For mail-server to mail-server services, check that the user specified in `cntnt` is accessible on the server. `nntp` For newsgroup services, check that the newgroup name specified in `cntnt` is accessible on the server. `ftp` For FTP services, check that the filename specified in `cntnt` is accessible on the server through anonymous login.
`bkup`	Set the real server used as the backup/overflow server for this real server group. To prevent loss of service if the entire real server group fails, use this option to assign a backup real server number. Then, if the real server group becomes inoperative, the switch will activate the backup real server until the one of the original real servers becomes operative again. The backup server is also used in overflow situations. If all the servers in the real server group reach their `mcon` (maximum connections) limit, the backup comes online to provide additional processing power until one of the original servers becomes desaturated. The same backup/overflow server may be assigned to more than one real server group at the same time.

Table 7-22 SLB Real Server Group Options (/cfg/slb/group)

Option	Description
`del`	Delete this real server group from the Layer 4 software configuration. This removes the group from operation under all virtual servers it is assigned to. Use this command with caution: if you remove the only group assigned to a virtual server, the virtual server will become inoperative.
`cur`	Displays the current configuration parameters for this real server group.

Server Load Balancing Metrics

There are a number of metrics that can be used for selecting which real server in a group gets the next client request:

Table 7-23 Real Server Group Metrics

Option	Description
`minmisses`	Minimum misses. This metric is optimized for Application Redirection, Firewall Load Balancing and Router Load Balancing. We recommend its use for all Application Redirection situations. When `minmisses` is specified for a real server group performing Application Redirection, all requests for a specific IP destination address will be sent to the same server. This is particularly useful in caching applications, helping to maximize successful cache hits. Best statistical load balancing is achieved when the IP address destinations of load balanced frames are spread across a broad range of IP subnets. Minmisses can also be used for Server Load Balancing. When specified for a real server group performing Server Load Balancing, all requests from a specific client will be sent to the same server. This is useful for applications where client information must be retained on the server between sessions. Server load with this metric becomes most evenly balanced as the number of active clients increases.
`hash`	Like `minmisses`, the `hash` metric uses IP address information in the client request to select a server. For Application Redirection, all requests for a specific IP destination address will be sent to the same server. This is particularly useful for maximizing successful cache hits. For Server Load Balancing, all requests from a specific client will be sent to the same server. This is useful for applications where client information must be retained between sessions. The `hash` metric should be used if the statistical load balancing achieved using `minmisses` is not as optimal as desired. Although the `hash` metric can provide more even load balancing at any given instance, it is not as effective as `minmisses` when servers leave and reenter service. If the Load Balancing statistics indicate that one server is processing significantly more requests over time than other servers, consider using the `hash` metric.

Table 7-23 Real Server Group Metrics

Option	Description
`leastconns`	Least connections. With this option, the number of connections currently open on each real server is measured in real time. The server with the fewest current connections is considered to be the best choice for the next client connection request. This option is the most self-regulating, with the fastest servers typically getting the most connections over time, due to their ability to accept, process, and shut down connections faster than slower servers.
`roundrobin`	Round robin. With this option, new connections are issued to each server in turn: the first real server in this group gets the first connection, the second real server gets the next connection, followed by the third real server, and so on. When all the real servers in this group have received at least one connection, the issuing process starts over with the first real server.

NOTE – Under the `leastconns` and `roundrobin` metrics, when real servers are configured with weights (see the `wght` option on page 7-32), a higher proportion of connections are given to servers with higher weights. This can improve load balancing among servers of different performance levels. Weights are not applied when using `hash` or `minmisses`.

The Virtual Server Menu

Direct command: **`/cfg/slb/virt`** *virtual-server-number*

This menu is used for configuring the virtual servers which will be the target for client requests for Server Load Balancing. The required minimum of parameters to configure is as follows:

- Virtual server IP address
- Adding a virtual TCP/UDP port and real server group
- Enabling the virtual server

To configure SLB virtual server parameters, at the `Server Load Balancing#` prompt, enter:

```
Server Load Balancing# virt virtual-server-number
```

Where *virtual-server-number* (1 to 256) represents the number of the virtual server that you wish to configure. The menu for the virtual server you entered is then displayed:

```
[Virtual server 1 Menu]
        vip   - Set IP addr of virtual server
        dname - Set domain name of virtual server
        layr3 - Enable/disable layer 3 only balancing
        add   - Add virtual port and real server group
        rem   - Remove virtual port
        map   - Map virtual port to real port
        hname - Set hostname of virtual port
        udp   - Enable/disable UDP balancing for virtual port
        pbind - Enable/disable persistent bindings for virtual port
        ena   - Enable virtual server
        dis   - Disable virtual server
        del   - Delete virtual server
        cur   - Display current virtual configuration

>> Virtual server 1#
```

Virtual server configuration options are described in the following table.

Table 7-24 SLB Virtual Server Options (/cfg/slb/virt)

Option	Description
`vip`	Set the IP address of the virtual server using dotted decimal notation. The virtual server created within the switch will respond to `ARP`s and `PING`s from network ports as if it was a normal server. Client requests directed to the virtual server's IP address will be balanced among the real servers available to it through real server group assignments.
`dname`	Set the domain name for this virtual server. The domain name typically includes the name of the company or organization, and the Internet group code (.com, .edu, .gov, .org, etcetera). An example would be foocorp.com. It does not include the hostname portion (www, www2, ftp, etcetera). To define the hostname, see `hname` below. To clear the `dname`, specify the name as **none**.
`layr3`	Normally, the client IP address is used with the client Layer 4 port number to produce a session identifier. When the `layr3` option is used, the switch uses only the client IP address as the session identifier, associating all the connections from the same client with the same real server while any connection exits between them. This is necessary for some server applications where state information about the client system is divided across different simultaneous connections, and also in applications where TCP fragments are generated. If the real server that the client is assigned to becomes unavailable, the Layer 4 software will allow the client to connect to a different server.

Table 7-24 SLB Virtual Server Options (/cfg/slb/virt)

Option	Description
`add`	Assign a virtual port to this virtual server, and a real server group to service it. Up to eight services can be defined for each virtual server. If more services are needed for a particular virtual IP address, two or more virtual servers can be created with the same virtual IP address, each with up to eight services. For any specific virtual server IP address, each service added must be unique (for example, HTTP services cannot be added to two different virtual servers with the same virtual IP address). At least one virtual port and group is required for each virtual server. The format for this command is as follows: # **add** *virtual-port real-server-group* The virtual port is the TCP/UDP port to which the clients will be sending connection requests. The *vitual-port* number or name can be specified. You can define your own virtual port, or use one of the well-known ports as follows: Number Name Number Name 20 ftp-data 110 pop3 21 ftp 111 sunrpc 22 ssh 119 nntp 23 telnet 123 ntp 25 smtp 143 imap 37 time 144 news 42 name 161 snmp 43 whois 162 snmptrap 53 domain 179 bgp 69 tftp 194 irc 70 gopher 220 imap3 79 finger 389 ldap 80 http 443 https 109 pop2 520 rip Each real server in the real server group is expected to have a server process operational and listening to the virtual port(s) that are configured on this virtual server. See the `map` command below for information about mapping well known server TCP/UDP ports to administrator selected TCP/UDP port numbers
`rem`	Remove a virtual port from this virtual server. You must select this command to deactivate a particular virtual service from this virtual server. You will be prompted to enter the TCP/UDP port number or name for the service to be deactivated.
`map`	Map a virtual port number or name to real server port number or name. See examples, below.

Table 7-24 SLB Virtual Server Options (/cfg/slb/virt)

Option	Description
`hname`	Set the virtual port and hostname for a service added. This is used in conjunction with dname (above) to create a full host/domain name for individual services. The format for this command is as follows: # **`hname`** *virtual-port hostname* For example, to add a hostname for web services, you could specify "http" as the virtual port and "www" as the hostname. If a `dname` of "foocorp.com" was defined (above), "www.foocorp.com" would be the full host/domain name for the service. To clear the hostname for a service, use the following command: # **`hname`** *virtual-port* **`none`**
`udp`	Enable/disable UDP balancing for a virtual port. You can configure this option if the services to be load balanced include UDP instead of, or in addition to, TCP. For example, NFS in some older networking environments might use UDP instead of TCP. In those environments, you must activate UDP balancing for the particular virtual servers that clients will communicate with using UDP.
`pbind`	Enable/disable persistent bindings for a real server. This may be necessary for some server applications where state information about the client system is retained on the server over a series of sequential connections, such as with SSL (Secure Socket Layer, https), Web site search results, or multi-page web forms. This option uses the client IP address as an identifier, and associates all the connections from the same client with the same real server until the client becomes inactive and the connection is aged out of the binding table. The connection timeout value (set on the Real Server Menu) is used to control how long these inactive but persistent connections remain associated with their real servers. When the client resumes activity *after* their connection has been aged out, they will be connected to the most appropriate real server based on the load balancing algorithm. An alternative approach may be to use the real server group metrics `minmisses` or `hash` (see "Server Load Balancing Metrics" on page 7-36).
`ena`	Enable this virtual server and its services. This option activates the virtual server within the switch so that it can service client requests sent to its defined IP address.
`dis`	This option is used to disable the virtual server so that it no longer services client requests.
`del`	This command removes this virtual server from operation within the switch and deletes it from the Layer 4 switching software configuration. Use this command with caution, as it will delete the options that have been set for this virtual server.
`cur`	Displays the current parameters for the virtual server.

Direct Client Access to Real Servers

Some clients may need direct access to the real servers. This can be provided in a number of ways:

- Multiple IP addresses on the server
- Proxy IP addresses
- Port mapping

Multiple IP Addresses on the Server

One way to provide both Layer 4 access and direct access to a real server, is to assign multiple IP addresses to the real server. For example, one IP address could be established exclusively for Layer 4 Server Load Balancing, and another could be used for direct access needs.

Proxy IP Addresses

Proxy IP addresses are used primarily to eliminate Server Load Balancing topology restrictions in complex networks (see "Network Topology Considerations" on page 15-4). Proxy IP addresses can also provide direct access to real servers.

If the switch port to the client is configured with a proxy IP address (see "IP Proxy Addresses for Complex Networks" on page 15-19), the client can access each real server directly using the real server's IP address. This requires that the switch port connected to the real server is *not* set to the "`server`" state (see the `state` option under `/cfg/slb/port` on page 7-47).

Server Load Balancing is still accessed using the virtual server IP address.

Port Mapping

Port mapping provides an alternative to proxy IP addresses for direct real server access.

When Server Load Balancing is used without proxy IP addresses, the virtual server *must* process both the client-to-server requests *and* the server-to-client responses. If a client were to access the real server IP address and port directly, bypassing Layer 4 preparation, the server-to-client response could be mishandled by Layer 4 processing as it returns through the switch.

Port mapping is typically used when the switch port attached to the real server is set to the "`server`" state, indicating that Layer 4 processing is required and that proxies are not used.

First, two port processes must be executed on the real server. One real server port will handle the direct traffic, and the other will handle Layer 4 traffic. Then, the virtual server port must be mapped to the proper real server port.

In the following figure, clients can access Layer 4 services through well-known TCP port 80 at the virtual server's IP address. This is mapped to TCP port 8000 on the real server. For direct access that bypasses the virtual server and Server Load Balancing, clients can specify well-known TCP port 80 at the real server's IP address.

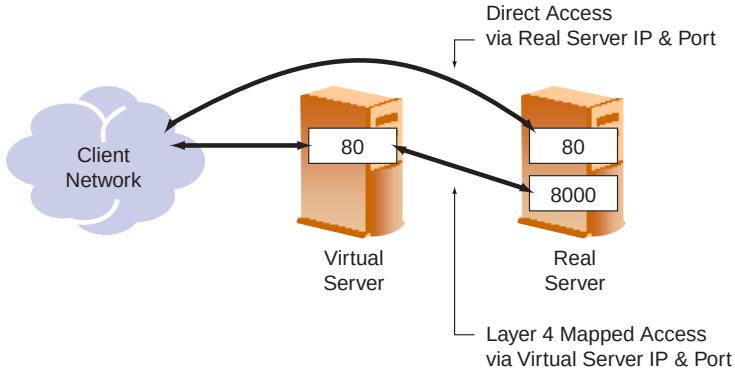

Figure 7-1 Mapped and Non-mapped Server Access

Mapping Virtual Ports to Real Ports

In addition to providing direct real server access in some situations, mapping is required when administrators choose to execute their real server processes on different TCP/UDP port than the well known TCP/UDP ports. Otherwise, virtual server ports are mapped directly to real server ports by default and require no mapping configuration.

The format for the `map` command is as follows:

```
Virtual server 1# map virtual-server-port real-server-port
```

The Filter Menu

Direct command: **`/cfg/slb/filt`** *filter-number*

The switch supports up to 224 traffic filters. Each filter can be configured to allow, deny, redirect or NAT traffic according to a variety of address and protocol specifications, and each physical switch port can be configured to use any combination of filters.

The required minimum of parameters to configure is as follows:

- Set the address, masks, and/or protocol which will be affected by the filter
- Set the action which the filter takes
- Enable the filter
- Add the filter to a switch port
- Enable filtering on the switch port

To configure Filtering parameters, at the `Server Load Balancing#` prompt, enter:

```
Server Load Balancing# filt filter-number
```

The menu is displayed for the selected filter:

```
[Filter 1 Menu]
        sip   - Set source IP address
        smask - Set source IP mask
        dip   - Set destination IP address
        dmask - Set destination IP mask
        proto - Set IP protocol
        sport - Set source TCP/UDP port or range
        dport - Set destination TCP/UDP port or range
        actio - Set action
        group - Set real server group for redirection
        rport - Set real server port for redirection
        proxy - Enable/disable client proxy
        inver - Enable/disable filter inversion
        nat   - Set which addresses are network address translated
        cache - Enable/disable caching sessions that match filter
        log   - Enable/disable logging
        ena   - Enable filter
        dis   - Disable filter
        del   - Delete filter
        cur   - Display current filter configuration

>> Filter 1#
```

Filter configuration options are described in the following table.

Table 7-25 Filter Options (/cfg/slb/filt)

Option	Description
`sip`	If defined, traffic with this source IP address will be affected by this filter. Specify an IP address in dotted decimal notation, or "**any**". A range of IP addresses is produced when used with the `smask` below.
`smask`	This IP address mask is used with the `sip` to select traffic which this filter will affect. See details below for more information on producing address ranges.
`dip`	If defined, traffic with this destination IP address will be affected by this filter. Specify an IP address in dotted decimal notation, or "**any**". A range of IP addresses is produced when used with the `dmask` below.
`dmask`	This IP address mask is used with the `dip` to select traffic which this filter will affect. See details below for more information on producing address ranges.
`proto`	If defined, traffic from the specified protocol is affected by this filter. The protocol number, name, or "**any**" can be specified: Number Name 1 icmp 2 igmp 6 tcp 17 udp 89 ospf
`sport`	If defined, traffic with the specified TCP or UDP source port will be affected by this filter. The port number, range, name, or "**any**" can be specified. The well-known ports are as follows: Number Name — Number Name 20 ftp-data — 110 pop3 21 ftp — 111 sunrpc 22 ssh — 119 nntp 23 telnet — 123 ntp 25 smtp — 143 imap 37 time — 144 news 42 name — 161 snmp 43 whois — 162 snmptrap 53 domain — 179 bgp 69 tftp — 194 irc 70 gopher — 220 imap3 79 finger — 389 ldap 80 http — 443 https 109 pop2 — 520 rip A number range can be specified by placing a dash between the low and high port number. For example: `31000-33000`
`dport`	If defined, traffic with the specified real server TCP or UDP destination port will be affected by this filter. The port number, range, name, or "**any**" can be specified, just as with `sport` above.

Table 7-25 Filter Options (/cfg/slb/filt)

Option	Description
actio	Specify the action this filter takes: `allow` Allow the frame to pass. `deny` Discard frames that fit this filter's profile. This can be used for building basic security profiles. `redir` Redirect frames that fit this filter's profile, such as for web-cache redirection. In addition, Layer 4 processing must be activated (see the `/cfg/slb/on` command on page 7-30). `nat` Perform generic Network Address Translation (NAT). This can be used to map the source or destination IP address and port information of a private network scheme to/from the advertised network IP address and ports. This is used in conjunction with the `nat` option below and can also be combined with proxies (see "Network Address Translation Examples" on page 16-20).
group	This option applies only when `redir` is specified at the filter action. Define a real server group (1 to 256) to which redirected traffic will be sent.
rport	This option applies only when `redir` is specified at the filter action. This defines the real server TCP or UDP port to which redirected traffic will be sent. For valid Layer 4 health checks, this must be configured whenever TCP protocol traffic is redirected. Also, if transparent proxies are used for Network Address Translation (NAT) on the switch (see the `pip` option on page 7-48), `rport` must be configured for all Application Redirection filters.
proxy	This option applies only when `redir` or `nat` is specified at the filter action. Enable or disable proxy IP address translation for traffic matching the filter criteria. By default, this option is enabled. If disabled, any proxy defined for the switch port using the `pip` command (see page 7-48) is not performed for traffic that meets the filter criteria. This is useful when some types of traffic must retain origial IP address information, or when other forms of translation (such as Application Redirection or NAT) are preferred.
inver	Invert the filter logic. If the conditions of the filter are met, *don't* act. If the conditions fo the filter are *not met,* perform the assigned action.
nat	When `nat` is set as the filter action (above), this command specifies whether the source or the destination information is remapped. If **`source`** is specified, the frame's source IP address (`sip`) and port number (`sport`) are replaced with the `dip` and `dport` values. If **`dest`** is specified, the frame's destination IP address (`dip`) and port number (`dport`) are replaced with the `sip` and `sport` values.
cache	Enable or disable caching of session information for this filter. Enabling this option can increase session performance, but takes some of the session binding resources. If you experience an unacceptable number of binding failures as shown in the Server Load Balancing Maintenance Statistics (`/stats/slb/maint`), you may wish to disable the session cache on filters which have lower performance priority.

Table 7-25 Filter Options (/cfg/slb/filt)

Option	Description
`log`	When enabled, a message is sent to the syslog whenever the filter encounters traffic that meets the profile. This is used primarily with filters that deny traffic for security purposes. To prevent a high-volume of `syslog` messages, do not use this option with filters that are triggered frequently.
`ena`	Turn this filter on.
`dis`	Disable this filter.
`del`	Remove this filter from the switch configuration.
`cur`	Displays current filter parameters.

Defining IP Address Ranges for Filters

You can specify a range of IP address for filtering both the source and/or destination IP address for traffic. When a range of IP addresses is needed, the `sip` (source) or `dip` (destination) defines the base IP address in the desired range, and the `smask` (source) or `dmask` (destination) is the mask which is applied to produce the range.

For example, to determine if a client request's destination IP address should be redirected to the cache servers attached to a particular switch, the destination IP address is masked (bitwise AND) with the `dmask` and then compared to the `dip`.

As another example, you could configure the switch with two filters so that each would handle traffic filtering for one half of the Internet. To do this, you could define the following parameters:

Table 7-26 Filtering IP Address Ranges

Filter	Internet Address Range	dip	dmask
#1	0.0.0.0 - 127.255.255.255	0.0.0.0	128.0.0.0
#2	128.0.0.0 - 255.255.255.255	128.0.0.0	128.0.0.0

The SLB Port Menu

Direct command: **/cfg/slb/port** *port-number*

To configure switch port parameters, at the Server Load Balancing# prompt, enter:

```
Server Load Balancing# port port-number
```

The menu for the port you entered is displayed.

```
[SLB port 1 Menu]
        state - Set Layer 4 state for port
        pip   - Set Proxy IP address for port
        filt  - Enable/disable filtering for port
        add   - Add filter to port
        rem   - Remove filter from port
        cur   - Display current port configuration

>> SLB port 1#
```

Configuration options are described in the following table.

Table 7-27 SLB Port Options (/cfg/slb/port)

Option	Description	
state	Set the port state for Layer 4 activity:	
	server	For Server Load Balancing, the port connects to servers. Traffic not associated with virtual servers is switched normally. In this state, the switch port remaps real server IP addresses and port values to virtual server IP addresses and ports.
	client	For Server Load Balancing, the port connects to clients. Traffic not associated with virtual servers is switched normally. In this state, the switch port binds servers to clients and remaps virtual server IP addresses and port values to real server IP addresses and port.
	redir	For Application Redirection, the port connects to clients. Frames matching redirection filters on this port are forwarded to real server groups.
	none	The port performs only normal Layer 2/Layer 3 switching. Virtual server traffic is not supported.
	failover	The port connects to a secondary switch for hot-standby. Inter-switch information, configuraiton, and topology changes are sent across this link.

Table 7-27 SLB Port Options (/cfg/slb/port)

Option	Description
`pip`	Set the proxy IP address for this port using dotted decimal notation. When defined, client address information in Layer 4 requests is replaced with this proxy address. In Server Load Balancing applications, this forces response traffic to return through switch as required, rather than around it as possible in complex routing environments. For examples, see "IP Proxy Addresses for Complex Networks" on page 15-19. Proxies are also useful for Application Redirection and Network Address Tranlastion (NAT). For examples, see "IP Proxy Addresses for Transparent Proxies or Complex Networks" on page 16-18, and also "Network Address Translation Examples" on page 16-20.
`filt`	Enable or disable filtering on this port.
`add`	Add a filter for use on this port.
`rem`	Remove a filter from use on this port.
`cur`	Displays current system parameters.

NOTE – Client ports, redirect ports, server ports, and fail-over link ports are mutually exclusive. They cannot be run simultaneously on the same port.

NOTE – When changing the filters on a given port, it may take some time before the port session information is updated so that the filter changes take effect. To make port filter changes take effect immediately, clear the session binding table for the port (see the `clear` command under Table 8-3 on page 8-5).

The SLB Failover Menu

Direct command: **`/cfg/slb/fail`**

To configure SLB failover parameters, at the `Server Load Balancing#` prompt, enter:

```
Server Load Balancing# fail
```

The SLB Failover Menu is displayed:

```
[SLB failover Menu]
        prima - Set IP addr of primary switch
        secon - Set IP addr of secondary switch
        resp  - Set seconds before silent peer is assumed down
        spply - Enable/disable RIP1 Supply when standby
        on    - Globally turn SLB failover ON
        off   - Globally turn SLB failover OFF
        cur   - Display current failover configuration

>> SLB failover menu#
```

Failover configuration options are described in the following table.

Table 7-28 SLB Failover Options (/cfg/slb/fail)

Option	Description
`prima`	Set IP address of primary switch.
`secon`	Set IP address of secondary switch.
`resp`	Set seconds before silent peer is assumed down.
`spply`	Enable or disable IP route supply for the secondary switch. When disabled, the switch does not supply IP routes to other routers while in standby mode.
`on`	Globally turn SLB failover ON.
`off`	Globally turn SLB failover OFF.
`cur`	Displays current system parameters.

The Distributed SLB Menu

Direct command: **`/cfg/slb/dist`**

To configure Distributed Server Load Balancing parameters, at the `Server Load Balancing#` prompt, enter:

```
Server Load Balancing# dist
```

The Distributed SLB Menu is displayed:

```
[Distributed SLB menu Menu]
        site  - Remote Site menu
        dns   - Enable/disable DNS handoffs
        ttl   - Set Time To Live of DNS resource records
        http  - Enable/disable HTTP redirects
        intr  - Set interval between remote site updates
        on    - Globally turn Distributed SLB ON
        off   - Globally turn Distributed SLB OFF
        cur   - Display current distributed SLB configuration

>> Distributed SLB menu#
```

Distributed SLB configuration options are described in the following table.

Table 7-29 Distributed SLB Options (/cfg/slb/dist)

Option	Description
`site`	Display the Remote Site Menu for one of up to eight remote sites.
`dns`	Enable or disable DNS handoffs to peer sites by this switch. This should be enabled for proper DSLB operation. If disabled, whenever the switch receives a DNS request for a configured service, it will respond only with its own virtual IP address, regardless of performance or load considerations.
`ttl`	Specify the duration (from 0 to 120 minutes) that the DNS response from the switch will remain in the cache of DNS servers. A lower value may increase the ability of the DSLB system to adjust to sudden changes in traffic load, but will generate more DNS traffic. Higher numbers may reduce the amount of DNS traffic, but may slow DSLB's response to sudden traffic changes.
`http`	Enable or disable HTTP Redirects to peer sites by this switch. When enabled, this switch will redirect client requests to peer sites if its own real servers fail or have reached their maximum connection limits. If disabled, the switch will not perform HTTP Redirects, but will instead drop requests for new connections and cause the client's browser to eventually issue a new DNS request.
`intr`	Set the time between Distributed Site State Protocol (DSSP) updates between this switch and its peers. The range is between 1 and 120 minutes.
`on`	Activate Distributed Server Load Balancing (DSLB) for this switch. This option can be performed only once the optional DSLB software is activated (see "Activating Optional Software" on page 8-6).
`off`	Turn DSLB off for this switch. Any active remote sites will still perform DSLB services with each other, but will not handoff requests to this switch.
`cur`	Displays current DSLB parameters.

The Remote Site Menu

Up to eight remote sites can be configured. To configure remote site parameters, from the `Distributed SLB Menu#` prompt, enter:

```
Distributed SLB menu# site site-number
```

The menu for the selected site will be displayed:

```
[Remote site 1 Menu]
        prima - Set primary switch IP address of remote site
        secon - Set secondary switch IP address of remote site
        updat - Enable/disable remote site updates
        ena   - Enable remote site
        dis   - Disable remote site
        del   - Delete remote site
        cur   - Display current remote site configuration

>> Remote site 1#
```

Remote site configuration options are described in the following table.

Table 7-30 Remote Site Options (/cfg/slb/dist/site)

Option	Description
`prima`	Define the IP interface IP address of the primary switch at the remote site used for Distributed Server Load Balancing. Use dotted decimal notation.
`secon`	If the remote site is configured with a redundant switch, enter the IP address of the remote secondary switch here. If the remote site primary switch fails, the local switch will address the remote site secondary switch instead.
`updat`	Enable or disable remote site updates. If enabled, this switch will send regular Distributed Site State Protocol (DSSP) updates to its remote peers using HTTP port 80. If disabled, the switch will not send state updates. If your local firewall does not permit this traffic, disable the updates.
`ena`	Enable this remote site for use with Distributed Server Load Balancing.
`dis`	Disable this remote site. The switch will no longer use this remote site for Distributed Server Load Balancing.
`del`	Remove this remote site from operation, and delete its configuration.
`cur`	Displays current system parameters.

NOTE – When `updat` (above) is enabled, Distributed Server Load Balancing uses service port 80 on the IP interface for DSSP updates. By default, ACEvision also uses port 80. Both services cannot use the same port. If both are enabled, configure ACEvision use a different service port (see the `/cfg/sys` options under Table 7-1 on page 7-5).

Configuring Port Trunking

Trunk groups can provide super-bandwidth connections between Alteon Networks switches or other trunk capable devices. A "trunk" is a group of ports that act together, combining their bandwidth to creating a single, larger port. Up to four trunk groups can be configured on the switch. The following restrictions apply:

- Any physical switch port can belong to no more than one trunk group.
- Up to four ports can belong to the same trunk group.
- Best performance is achieved when all ports in any given trunk group are configured for the same speed.
- Trunking from non-Alteon Networks devices must comply with Cisco® EtherChannel® technology.

To configure trunking parameters, enter the following at the Configuration Menu:

```
Configuration# trunk trunk-group-number
```

The Trunk Group Menu for the selected group is displayed:

```
[Trunk group 1 Menu]
        add   - Add port to trunk group
        rem   - Remove port from trunk group
        ena   - Enable trunk group
        dis   - Disable trunk group
        del   - Delete trunk group
        cur   - Display current Trunk Group configuration

>> Trunk group 1#
```

Trunk group configuration options are described in the following table.

Table 7-31 Trunk Group Options (/cfg/trunk)

Option	Description
`add`	Add a physical port to the current trunk group.
`rem`	Remove a physical port from the current trunk group.
`ena`	Enable the current trunk group.
`dis`	Turn the current trunk group off.
`del`	Remove the current trunk group configuration.
`cur`	Displays current trunk group parameters.

CHAPTER 8

The Operations Menu

The Operations Menu is generally used for commands which affect switch performance immediately, but do not alter permanent switch configurations. For example, you can use the Operations Menu to immediately disable a port (without the need to apply or save the change), while making sure that when the switch is reset, the port returns to its normally configured operation.

The Operations Menu is available from the Main Menu. At the Main Menu prompt, enter:

```
# oper
```

The Operations Menu is displayed:

```
[Operations Menu]
        port  - Operational Port items menu
        mirr  - Operational Mirroring menu
        slb   - Operational Server Load Balancing menu
        swkey - Enter key to enable software feature
        rmkey - Enter software feature to be removed

>> Operations#
```

Each of these options is discussed in greater detail in the following sections.

Operations-Level Port Options

Direct command: **`/oper/port`** *port-number*

Operations-level port options are used for temporarily disabling or enabling a port, and for changing RMON status on a port. The Operations Port Menu is available from the `Operations#` prompt:

```
>> Operations # port port-number
```

The Operations Port Menu appears:

```
[Operations Port 1 Menu]
        dis   - Disable port
        ena   - Enable port
        rmon  - Enable/Disable RMON for port
        cur   - Current port state

>> Operations Port 1#
```

The options are described in the following table.

Table 8-1 Operations Port Menu Options (/oper/port)

Option	Description
`dis`	Temporarily disable the port. The port will be returned to its configured operation mode when the switch is reset.
`ena`	Temporarily enable the port. The port will be returned to its configured operation mode when the switch is reset.
`rmon`	Temporarily toggle RMON support on or off for the port. The port will be returned to its configured operation mode when the switch is reset.
`cur`	Display the current settings for the port.

Operations-Level Port Mirroring Options

Direct command: **/oper/mirr**

The Port Mirroring Menu is used to configure, enable, and disable the port monitor. When enabled, Layer 2 network packets being sent and/or received on a target port are duplicated and sent to a monitor port. By attaching a network analyzer to the monitor port, you can collect detailed information about your network performance and usage.

NOTE – Layer 3 and Layer 4 traffic is not mirrored through this facility.

NOTE – Port Mirroring cannot be used simultaneously with Layer 4 services (Server Load Balancing or Application Redirection) on any switch port connected to a server either directly, or through another switch or hub. For Server Load Balancing, this applies to any switch port configured in the "`server`" state. For Application Redirection, this applies to any switch port that has a cache server attached to it directly or indirectly. Use your network analyzer with a full-duplex pass-through connection or an Ethernet hub when troubleshooting a switch port for a server used for Layer 4 services.

Port Mirroring parameters are configured from the Operations Menu:

```
>> Operations # mirr
```

The Port Mirroring Menu is displayed.

```
[Port Mirroring Menu]
        to    - Set "Monitoring" port
        from  - Set "Mirrored" port
        dir   - Set Direction [in, out, both]
        tmout - Set Mirroring Timeout value
        dis   - Disable Port Mirroring
        ena   - Enable Port Mirroring
        cur   - Display current Port Mirroring configuration

>> Port Mirroring#
```

The Port Mirroring Menu options are described in the following table.

Table 8-2 Port Mirroring Menu Options (/oper/mirr)

Option	Description
`to`	This defines the monitoring port. When port mirroring is enabled, packets received and/or transmitted by the mirrored port will be duplicated to the switch port specified in this command.
`from`	This defines the mirrored port. When port mirroring is enabled, packets received and/or sent by the port specified in this command will be sent to the monitor port.
`dir`	This determines which type of packets will be sent to the monitor port: `in` = packets received at the mirrored port `out` = packets sent from the mirrored port `both` = packets sent and received by the mirrored port
`tmout`	Port mirroring will be automatically disabled (regardless of port state) after the time-out period specified in this command. Valid times are from 0 (does not time-out) to 86400 seconds.
`dis`	Turns port mirroring off.
`ena`	Turns port mirroring on.
`cur`	Displays the current parameter settings.

Operations-Level Server Load Balancing Options

Direct command: **`/oper/slb`**

When the optional Layer 4 software is enabled, the operations-level Server Load Balancing options are used for temporarily disabling or enabling real servers, selecting the switch for active or hot-standby mode, and synchronizing the configuration on the active and hot-standby switches. The menu is available from the `Operations#` prompt:

```
>> Operations # slb
```

The Server Load Balancing Operations Menu appears:

```
[Server Load Balancing Operations Menu]
        ena   - Enable real server
        dis   - Disable real server
        actve - Set switch to active
        stnby - Set switch to standby
        synch - Synchronize configuration
        clear - Clear session table on port
        cur   - Current SLB operational state

>> Server Load Balancing Operations#
```

The options are described in the following table.

Table 8-3 Server Load Balancing Operations Menu Options (/oper/slb)

Option	Description
`ena`	Temporarily enable a real server. The real server will be returned to its configured operation mode when the switch is reset.
`dis`	Temporarily disable a real server, removing it from operation within its real server group and virtual server. The real server will be returned to its configured operation mode when the switch is reset.
`actve`	In a network with redundant switches, the command selects this switch as the current primary switch.
`stnby`	In a network with redundant switches, the command selects this switch as the backup switch.
`synch`	In a network with redundant switches, the active switch configuration is copied to the backup switch automatically according to an internal schedule. If desired, the `synch` command can be issued manually to initiate an immediate synchronization of switch configuration states.
`clear`	Clear the session table for a specific port, and allow port filter changes to take effect immediately. Note: This disrupts current Server Load Balancing and Application Redirection sessions.
`cur`	Display the current settings for the port.

Activating Optional Software

Direct command: **/oper/swkey**

The swkey option is used for activating any optional software you have purchased for your switch.

Before you can activate optional software, you must obtain a software license from your Alteon Networks representative or authorized reseller. One software license is needed for each switch where the optional software is to be used. You will receive a Licence Certificate for each software license purchased.

To obtain a software key, you must register each License Certificate with Alteon Networks, and provide the MAC address of the ACElerate switch that will run the optional software. Alteon Networks will then provide a License Password.

NOTE – Each License Password will work only on the specific switch which has the MAC address you provided when registering your Licence Certificate.

Once you have your License Password, perform the following actions:

1. **Connect to the switch's command-line interface and log in as the administrator (see Chapter 2, "The Command-Line Interface").**

2. **At the Main# prompt, enter:**

```
Main# oper
```

3. **At the Operations# prompt, enter:**

```
Operations# swkey
```

4. **When prompted, enter your 16-digit software key code. For example:**

```
Enter Software Key: 123456789ABCDEF
```

If the correct code is entered, you will see the following message:

```
Valid software key entered.
Software feature enabled.
```

Removing Optional Software

Direct command: **/oper/rmkey**

The `rmkey` option is used for deactivating any optional software. Deactivated software is still present in switch memory and can be reactivated at any later time.

To deactivate optional software, enter the following at the Operations Menu:

```
Operations# rmkey
```

When prompted, enter the code for software to be removed. For example:

```
Enter Software Feature to be removed: [SLB]|DSLB|WCR: SLB
```

CHAPTER 9

The Boot Options Menu

To use the Boot Options Menu, you must be logged in to the switch as the administrator. The Boot Options Menu provides options for:

- Selecting a switch software image to be used when the switch is next reset
- Selecting a configuration block to be used when the switch is next reset
- Downloading a new software image to the switch via TFTP

To access the Boot Options Menu, at the Main Menu prompt, enter:

```
Main# boot
```

The Boot Options Menu is displayed:

```
[Boot Options Menu]
        image - Select software image to use on next boot
        conf  - Select config block to use on next boot
        tftp  - Download new software image via TFTP
        reset - Reset switch [WARNING: Restarts Spanning Tree]
        cur   - Display current boot options

>> Boot Options#
```

Each of these options is discussed in greater detail in the following sections.

Updating the Switch Software Image

The switch software image is the executable code running on the switch. A version of the image ships with the switch, and comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch.

Upgrading the software image on your switch requires the following:

- Loading the new image onto a TFTP server on your network
- Downloading the new image from the TFTP server to your switch
- Selecting the new software image to be loaded into switch memory the next time the switch is reset

Downloading New Software to Your Switch

The switch can store up to two different software images, called `image1` and `image2`, as well as boot software, called `boot`. When you download new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

To download a new software to your switch, you will need the following:

- The image or boot software loaded on a TFTP server on your network
- The hostname or IP address of the TFTP server
- The name of the new software image or boot file

NOTE – The DNS parameters must be configured if specifying hostnames. See "Domain Name System Menu" on page 7-17).

When the above requirements are met, use the following procedure to download the new software to your switch.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# tftp
```

2. **Enter the name of the switch software to be replaced:**

```
Enter name of switch software image to be replaced
 ["image1"/"image2"/"boot"]:
```

3. **Enter the hostname or IP address of the TFTP server.**

```
Enter hostname or IP address of TFTP server:
```

4. **Enter the name of the new software file on the server.**

```
Enter name of file on TFTP server:
```

The exact form of the name will vary by TFTP server. However, the file location is normally relative to the TFTP directory (usually `/tftpboot`).

5. **The system prompts you to confirm your request.**

You should next select a software image to run, as described below.

Selecting a Software Image to Run

You can select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# image
```

2. **Enter the name of the image you want the switch to use upon the next boot.**

The system informs you of which image is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use switch software "image1" on next reset.
Specify new image to use on next reset ["image1"/"image2"]:
```

Selecting a Configuration Block

When you make configuration changes to the switch, you must save the changes so that they are retained beyond the next time the switch is reset. When you perform the `save` command, your new configuration changes are placed in the *active* configuration block. The previous configuration is copied into the *backup* configuration block.

There is also a *factory* configuration block. This holds the default configuration set by the factory when your switch was constructed. Under certain circumstances, it may be desirable to reset the switch configuration to the default. This can be useful when a custom-configured switch is moved to a network environment where it will be reconfigured for a different purpose.

Use the following procedure to set which configuration block you want the switch to load the next time it is reset:

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# conf
```

2. **Enter the name of the configuration block you want the switch to use:**

The system informs you of which configuration block is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use active configuration block on next reset.
Specify new block to use ["active"/"backup"/"factory"]:
```

Resetting the Switch

You can reset the switch to make your software image file and configuration block changes occur.

NOTE – Resetting the switch causes the Spanning-Tree Protocol to restart. This process can be lengthy, depending on the topology of your network.

To reset the switch, at the `Boot Options#` prompt, enter:

```
>> Boot Options# reset
```

You are prompted to confirm your request.

CHAPTER 10

The Maintenance Menu

The Maintenance Menu is used to manage dump information and forwarding database information. It also includes a debugging menu to help with troubleshooting.

Dump information contains internal switch state data that is written to flash memory on the switch after any one of the following occurs:

- The switch administrator forces a switch *panic*. The `panic` option, found in the Maintenance Menu, causes the switch to dump state information to flash memory, and then causes the switch to reboot.
- The switch administrator enters the switch reset key combination on a device attached to the console port. The switch reset key combination is <Shift-Ctrl-6>.
- The watchdog timer forces a switch reset. The purpose of the watchdog timer is to reboot the switch if the switch software freezes.
- The switch detects a hardware or software problem that requires a reboot.

To use the Maintenance Menu, you must be logged in to the switch as the administrator. To access the Maintenance Menu, at the `Main#` prompt, enter:

```
Main# maint
```

The Maintenance Menu is displayed:

```
[Maintenance Menu]
        uudmp - Uuencode FLASH dump
        cldmp - Clear FLASH dump
        panic - Dump state information to FLASH and reboot
        fdb   - Forwarding Database Manipulation Menu
        debug - Debugging Menu
        arp   - ARP Cache Manipulation Menu
        route - IP Route Manipulation Menu

>> Maintenance#
```

Uuencode Flash Dump

Direct command: **/maint/uudmp**

Using this command, dump information is presented in uuencoded format. This format makes it easy to capture the dump information as a file or a string of characters. You can then contact Alteon Networks Customer Support for help analyzing the information.

If you want to capture dump information to a file, set your communication software on your workstation to capture session data prior to issuing the `uudmp` command. This will ensure that you do not lose any information. Once entered, the `uudmp` command will cause approximately 1460 lines of data to be displayed on your screen and copied into the file.

Using the `uudmp` command, dump information can be read multiple times. The command does not cause the information to be updated or cleared from flash memory.

NOTE – Dump information is not cleared automatically. In order for any subsequent dump information to be written to flash memory, you must manually clear the dump region. For more information on clearing the dump region, see "Clearing Dump Information" on page 10-3.

To access dump information, at the `Maintenance#` prompt, enter:

```
Maintenance# uudmp
```

The dump information is displayed on your screen and, if you have configured your communication software to do so, captured to a file. If the dump region is empty, the following appears:

```
No FLASH dump available.
```

Clearing Dump Information

Direct command: **/maint/cldmp**

To clear dump information from flash memory, at the Maintenance# prompt, enter:

```
Maintenance# cldmp
```

The switch clears the dump region of flash memory and displays the following message:

```
FLASH dump region cleared.
```

If the flash dump region is already clear, the switch displays the following message:

```
FLASH dump region is already clear.
```

Using the Panic Command

Direct command: **/maint/panic**

The panic command causes the switch to immediately dump state information to flash memory and automatically reboot.

To select panic, at the Maintenance# prompt, enter:

```
Maintenance# panic
```

Enter **y** to confirm the command:

```
Confirm dump and reboot [y/n]: y
```

The following messages are displayed:

```
Starting system dump...done.

Reboot at 11:54:08 Thursday June 26, 1997...

Boot version 1.0.1

Alteon ACEswitch 180

Rebooted because of console PANIC command.

Booting complete 11:55:01 Thursday June 26, 1997:
```

Unscheduled System Dumps

If there is an unscheduled system dump to flash memory, the following message is displayed when you log on to the switch:

```
Note: A system dump exists in FLASH. The dump was saved
      at 13:43:22 Fri Jun 27, 1997. Use /maint/uudmp to
      extract the dump for analysis and /maint/cldmp to
      clear the FLASH region. The region must be cleared
      before another dump can be taken.
```

The Forwarding Database Menu

Direct command: **`/maint/fdb`**

The Forwarding Database Menu can be used to view information, and to delete a MAC address from the forwarding database or clear the entire forwarding database. This is helpful in identifying problems associated with MAC address learning and packet forwarding decisions.

To access the FDB Manipulation Menu, at the `Maintenance#` prompt, enter:

```
Maintenance# fdb
```

The FDB Manipulation Menu is displayed:

```
[Forwarding Database Menu]
        find  - Show a single FDB entry by MAC address
        port  - Show FDB entries for a single port
        vlan  - Show FDB entries for a single VLAN
        refpt - Show FDB entries referenced by a single port
        dump  - Show all FDB entries
        stats - Show FDB statistics
        del   - Delete an FDB entry
        clear - Clear entire FDB

>> Forwarding Database#
```

Delete an FDB entry

To delete a MAC address from the FDB, at the `Forwarding Database#` prompt, enter:

```
Forwarding Database# del MAC-address
```

Clear entire FDB

To clear the entire FDB, at the `Forwarding Database#` prompt, enter:

```
Forwarding Database# clear
```

The FDB is cleared of all the entries.

The other information viewing choices on the Forwarding Database Menu are covered under "Forwarding Database Information Menu" on page 5-10.

Using the Miscellaneous Debug Menu

The Miscellaneous Debug Menu displays trace buffer information about certain events that can be helpful in understanding the switch operation. You can view the following information using the debug menu:

- Events traced by the Management Processor (MP)
- Events traced by the Switching Processor (SP)
- Events traced to a buffer area when a reset occurs

If the switch resets for any reason, the MP trace buffer and SP trace buffers are saved into the snap trace buffer area.

The output from these commands can be interpreted by the Alteon Networks Customer Support organization.

Snap Trace Information

A snap trace is taken when the switch resets and a message is sent to the console. Possible causes for a snap trace to be taken are:

- Watchdog timer: The processor is reset if the Management Processor fails to refresh the on-board timer. A snap trace is initiated which resets the switch.
- Software reset: Upon encountering certain error conditions or anomalies, the software triggers a panic. A snap trace is generated which dumps information to a file, and resets the switch.

Actions that can be taken if a snap trace is generated are:

- Record console messages and send them to Alteon Networks Customer Support.
- Retrieve the dump file by using the Maintenance Menu and choosing `uudmp`. Refer to "Uuencode Flash Dump" on page 10-2 for more information.

Accessing the Miscellaneous Debug Menu

Direct command: **`/maint/debug`**

To access the Miscellaneous Debug Menu, at the `Maintenance#` prompt, enter:

```
Maintenance# debug
```

The Miscellaneous Debug Menu is displayed:

```
[Miscellaneous Debug Menu]
        tbuf  - Display MP trace buffer
        snap  - Display MP snap (or post-mortem) trace buffer
        sptb  - Display SP trace buffer

>> Debug#
```

Display Management Processor Trace Buffer

Direct command: **/maint/debug/tbuf**

To view events traced by the MP, at the Debug# prompt, enter:

```
Debug# tbuf
```

Header information similar to the following is displayed:

```
MP trace buffer at 18:27:37 Mon Dec 29, 1997; mask: 0x2ffff748
```

The buffer information is displayed after the header.

Display Switching Processor Trace Buffer

Direct command: **/maint/debug/sptb** *port-number*

To view events traced by the SP, at the Debug# prompt, enter:

```
Debug# sptb port-number
```

Header information similar to the following is displayed:

```
Port 1 trace buffer at 18:27:41 Mon Dec 29, 1997; mask:0x018007eb
```

The buffer information is displayed after the header.

Display MP Snap Trace Buffer

Direct command: **/maint/debug/snap**

To view buffer information traced at the time that a reset occurred, at the Debug# prompt, enter:

```
Debug# snap
```

Using the ARP Cache Manipulation Menu

Direct command: **/maint/arp**

To access the ARP Cache Manipulation Menu, at the Maintenance# prompt, enter:

```
Maintenance# arp
```

The Address Resolution Protocol Menu is displayed:

```
[Address Resolution Protocol Menu]
        find  - Show a single ARP entry by IP address
        port  - Show ARP entries on a single port
        vlan  - Show ARP entries on a single VLAN
        refpt - Show ARP entries referenced by a single port
        dump  - Show all ARP entries
        del   - Delete an ARP entry
        clear - Clear ARP cache

>> Address Resolution Protocol#
```

Show ARP Entries

See "ARP Information Menu" on page 5-16.

Delete an ARP Entry

Direct command: **/maint/arp/del** *IP-address*

To remove a single ARP entry from switch memory, enter the following at the ARP Menu:

```
>> Address Resolution Protocol# del I-address
```

Clear All ARP Entries

Direct command: **/maint/arp/clear**

To clear the entire ARP list from switch memory, enter the following at the ARP Menu:

```
>> Address Resolution Protocol# clear
```

Using the IP Route Manipulation Menu

Direct command: **/maint/route**

To access the IP Route Manipulation Menu, at the Maintenance# prompt, enter:

```
Maintenance# route
```

The IP Routing Menu is displayed:

```
[IP Routing Menu]
        find  - Show a single route by destination IP address
        gw    - Show routes to a single gateway
        type  - Show routes of a single type
        tag   - Show routes of a single tag
        if    - Show routes on a single interface
        dump  - Show all routes
        clear - Clear route table

>> IP Routing#
```

Show Routes

See "IP Routing Information Menu" on page 5-13.

Clear the Routing Table

Direct command: **/maint/route/clear**

To clear all dynamic routes from switch memory, enter the following at the IP Routing Menu:

```
>> IP Routing# clear
```

Part 3: Tutorials and Examples

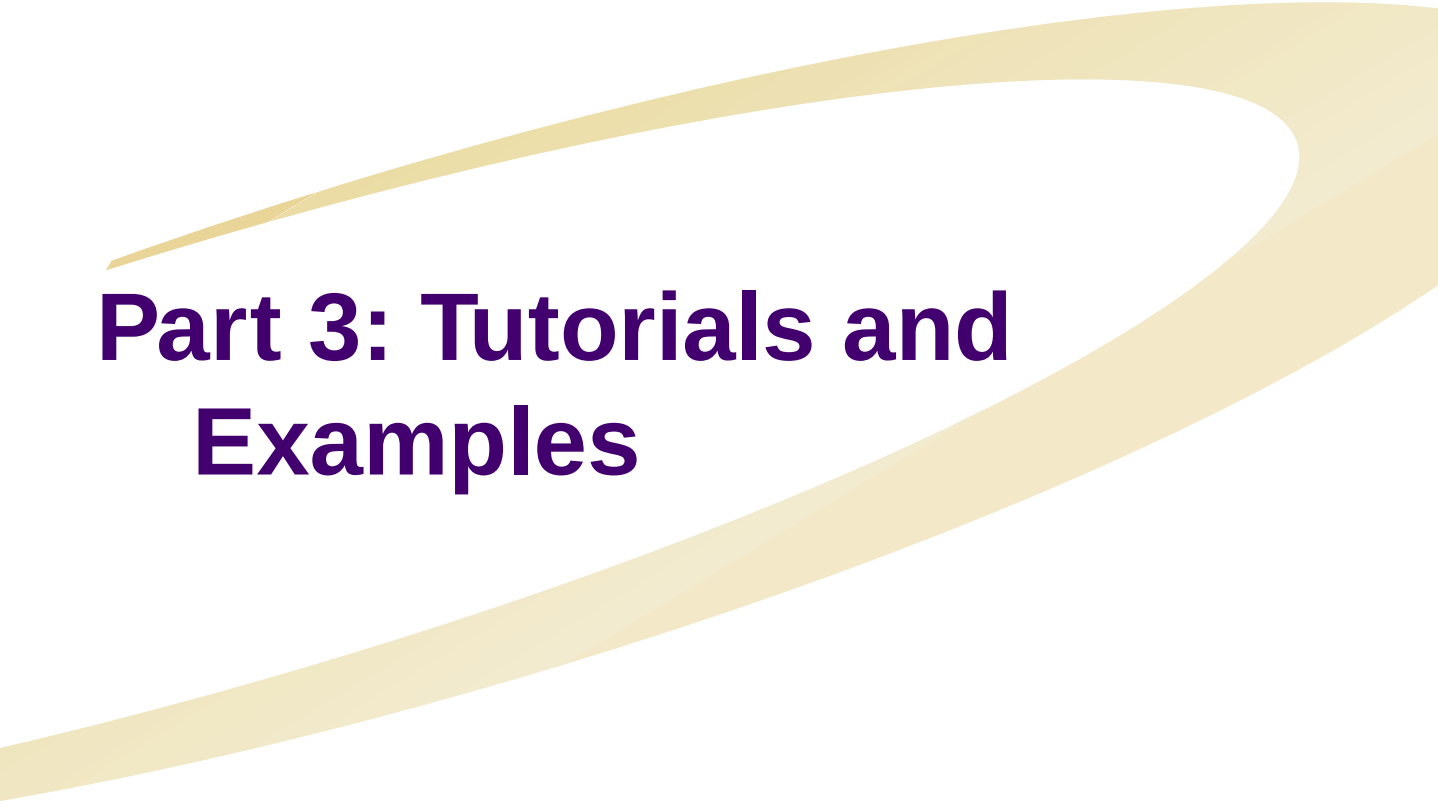

CHAPTER 11

VLANs

Virtual Local Area Networks (*VLANs*) are commonly used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments.

Basic VLANs can be configured during initial switch configuration (see "Using the Setup Utility" on page 3-1). More comprehensive VLAN configuration can be done from the command-line interface (see "Configuring VLAN Parameters" on page 7-19 as well as "Configuring Port Parameters" on page 7-6).

VLAN ID Numbers

The ACElerate software Release 5.0 (and greater) supports up to 246 VLANs per switch. Even though the maximum number of VLANs supported at any given time is 246, each can be identified with any number between 1 and 4094.

VLANs are defined on a per-port basis. Each port on the switch can belong to one or more VLANs, and each VLAN can have any number of switch ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN *tagging* enabled (see below).

Each port in the switch has a configurable default VLAN number, known as its *PVID*. The factory default value of all PVIDs is 1. This places all ports on the same VLAN initially, although each port's PVID is configurable to any VLAN number between 1 and 4094.

Any non-tagged frames (those with no VLAN specified) are classified with the sending port's PVID.

VLAN Tagging

The ACElerate software supports 802.1Q VLAN *tagging*, providing standards-based VLAN support for Ethernet systems.

Tagging places the VLAN identifier in the frame header, allowing multiple VLANs per port. When you configure multiple VLANs on a port, you must also enable tagging on that port.

Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags.

VLANs and Spanning-Tree

When *Spanning-Tree* is enabled on the switch, it detects and eliminates logical loops in a bridged or switched network. When multiple paths exist, Spanning-Tree configures the network so that a switch uses only the most efficient path. If that path fails, Spanning-Tree automatically sets up another active path on the network to sustain network operations.

If you configure the switch with Spanning-Tree, there will be a single instance of Spanning-Tree per switch regardless of number of configured VLANs in an enabled state.

VLANs and the IP Interfaces

Careful consideration must be made when creating VLANs within the switch, such that communication with the switch Management Processor (MP) remains possible where it is required.

Access to the switch for remote configuration, trap messages, and other management functions can only be accomplished from stations that are on VLANs which include an IP interface to the switch (see "IP Interface Menu" on page 7-10). Likewise, access to management functions can be cut off to any VLAN by excluding IP interfaces from its membership.

For example, if all IP interfaces are left on VLAN #1 (the default), and all other ports are configured for VLANs other than VLAN #1, then switch management features are effectively cut off. If an IP interface is added to one of the other VLANs, the stations in that VLAN all have access to switch management features.

VLAN Topologies and Design Issues

By default, the ACElerate software has a single VLAN configured on every port. This groups all ports into the same broadcast domain. This VLAN has an 802.1Q VLAN PVID of 1. Since in this default only a single VLAN is configured per port, VLAN tagging is turned off.

Since VLANs are most commonly used to create individual broadcast domains and/or separate IP subnets, it is useful for host systems to be able to have presence on more than one VLAN simultaneously. Alteon Networks switches and ACEnics have the unique capability of being able to support multiple VLANS on a per port or per interface basis, allowing very flexible configurations.

You can configure multiple VLANs on a single ACEnic, with each VLAN being configured through a logical interface and logical IP address on the host system. Each VLAN configured on the NIC must also be configured on the switch port to which it is connected. If multiple VLANs are configured on the port, tagging must be turned on.

Using this flexible multi-VLAN system, you can logically connect users and segments to a host with a single ACEnic that supports many logical segments or subnets.

Example #1: Multiple VLANS with Tagging NICs

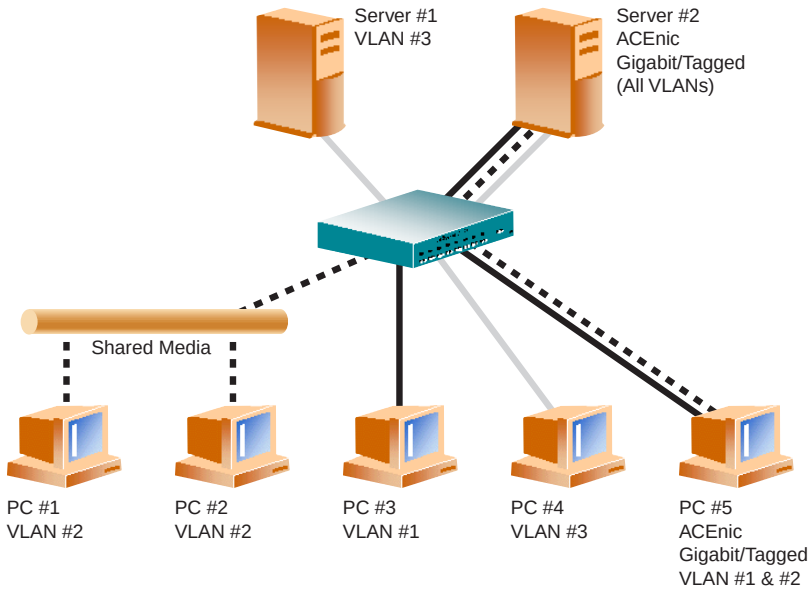

Figure 11-1 Example #1: Multiple VLANs with Tagging NICs

The following items describe the features of this VLAN:

- The ACEswitch is configured for three VLANs that also represent three different IP subnets.
- Two servers and five clients are attached to the switch.
- Server #1 is part of the VLAN #3 and only has presence in one IP subnet. The port that it is attached to is configured only for VLAN #3, so VLAN tagging is off.
- Server #2 is a high-use server that needs to be accessed from all VLANs and IP subnets. The server has an Alteon Networks ACEnic installed with VLAN tagging turned on. The NIC is attached to one of the ACEswitch's Gigabit Ethernet ports, which is configured for VLANs #1, #2, and #3, and also has tagging turned on. Because of the VLAN tagging capabilities of both the NIC and the switch, the server is able to communicate on all three IP subnets in this network, but continues to maintain broadcast separation between all three VLANs and subnets.
- PCs #1 and #2 are attached to a shared media hub that is then connected to the switch. They belong to VLAN #2, and are logically in the same IP subnet as Server #2 and PC #5. Tagging is not enabled on their switch port.
- PC #3 is a member of VLAN #1, and can only communicate with Server #2 and PC #5.
- PC #4 is a member of VLAN #3, and can only communicate with Server #1 and Server #2.
- PC #5 is a member of both VLAN #1 and VLAN #2, and has an Alteon Networks ACEnic Gigabit Ethernet Adapter installed. It is able to communicate with Server #2 via VLAN #1, and to PC #1 and PC #2 via VLAN #2. The switch port to which it is connected is configured for both VLAN #1 and VLAN #2, and has tagging turned on.
- VLAN tagging is only required on ports that are connected to other ACEswitches, or on ports that connect to tag-capable end-stations, such as servers with Alteon Networks ACEnic Gigabit Ethernet Adapters.

Example #2: Parallel Links with VLANs

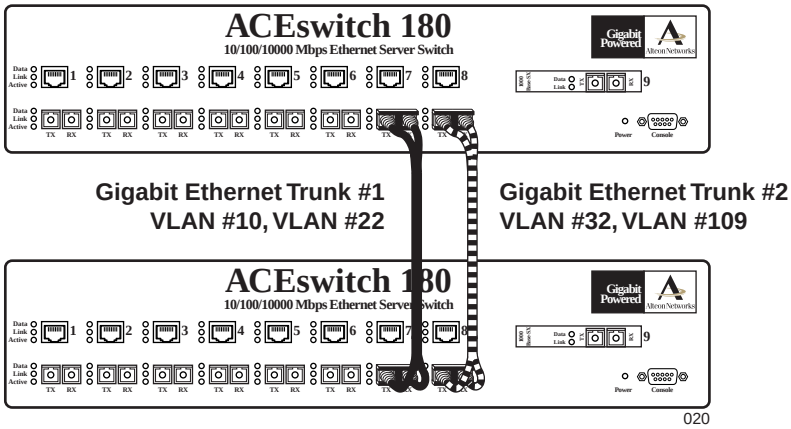

Figure 11-2 Example #2: Parallel Links with VLANs

The following items describe the features of this example:

- Example #2 shows how, through the use of VLANs, it is possible to create configurations where there are multiple links between two switches, without creating broadcast loops.
- Two ACEswitches are connected with two different Gigabit Ethernet links. Without VLANs, this configuration would create a broadcast loop, but the Spanning-Tree Protocol (STP) Topology Resolution process resolves parallel loop-creating links.
- With VLANs, neither switch-to-switch link shares the same VLAN and thus, are separated into their own broadcast domains.
- Ports #1 and #2 on both switches are on VLAN #10; Ports #3 and #4 on both switches are on VLAN #22. Ports #5 and #6 on both switches are on VLAN #32; and port #9 on both switches are on VLAN #109.
- It is necessary to turn off Spanning-Tree on at least one of the switch-to-switch links, or alternately turned off in both switches. Spanning-Tree executes on a per-network level, not a per-VLAN level. STP Bridge PDUs will be transmitted out both connected Gigabit Ethernet ports and be interpreted by the connected switch that there is a loop to resolve.
- Spanning-Tree is not VLAN-aware. Therefore, any VLAN configuration that might involve a parallel link from an STP perspective must be taken into account during network design. Alteon Networks recommends that you avoid topologies such as these, if at all possible.

CHAPTER 12

Jumbo Frames

To reduce host frame processing overhead, the Alteon Networks ACEnics and ACElerate powered switches, both running operating software version 2.0 or greater, can receive and transmit frames that are far larger than the maximum normal Ethernet frame. By sending one Jumbo Frame instead of myriad smaller frames, the same task is accomplished with less processing.

The switches and the ACEnic support Jumbo Frame sizes up to 9022 octets. These can be transmitted and received between ACEnic-enabled hosts through the switch across any VLAN.

Isolating Jumbo Frame Traffic using VLANs

Jumbo Frame traffic must not be used on a VLAN where there is any device that cannot process frame sizes larger than Ethernet maximum frame size.

Additional VLANs can be configured on the NICs and switches to support non-Jumbo Frame VLANs for servers and workstations that do not support extended frame sizes. End-stations with an ACEnics installed and attached to switches can communicate across both the Jumbo Frame VLANs and regular frame VLANs at the same time.

In the example illustrated in Figure 12-1, the two servers can handle Jumbo Frames but the two clients cannot; therefore Jumbo Frames should only be enabled and used on the VLAN represented by the solid lines, but not for the VLAN with the dashed lines. Jumbo Frames are not supported on ports configured for half-duplex mode.

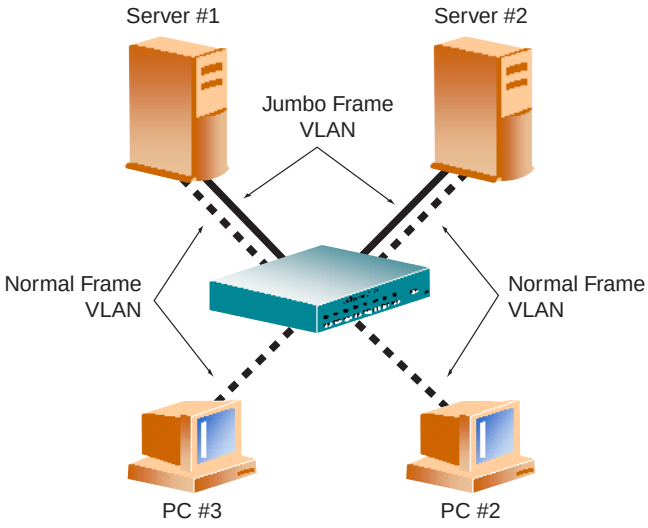

Figure 12-1 Jumbo Frame VLANs

Routing Jumbo Frames to Non-Jumbo Frame VLANs

When IP Routing is used to route traffic between VLANs, the switch will fragment jumbo UDP datagrams when routing from a Jumbo Frame VLAN to a non-Jumbo Frame VLAN. The resulting Jumbo Frame to regular frame conversion makes implementation even easier.

CHAPTER 13

IP Routing

IP Routing Benefits

IP Routing allows the network administrator to seamlessly connect server IP subnets to the rest of the backbone network, using a combination of configurable IP switch interfaces and IP routing options.

The IP Routing feature enhances Alteon's Server Switching solution in the following ways:

- It provides the ability to perform Server Load Balancing (using both Layer 3 and Layer 4 switching in combination) to server subnets which are separate from backbone subnets.
- By automatically fragmenting UDP Jumbo Frames when routing to non-Jumbo Frame VLANs or subnets, it provides another means to invisibly introduce Jumbo Frames technology into the Server Switched network.
- It provides the ability to seamlessly route IP traffic between multiple VLANs configured in the switch.

Example of Routing Between IP Subnets

The physical layout of most corporate networks has evolved over time. Classic hub/router topologies have given way to faster switched topologies, particularly now that switches are increasingly intelligent. ACElerate powered switches, in fact, are now smart enough and fast enough to perform routing functions on par with wire speed Layer 2 switching.

The combination of faster routing and switching in a single device provides another service: it allows you to build versatile topologies that account for legacy configurations.

For example, consider the following topology migration:

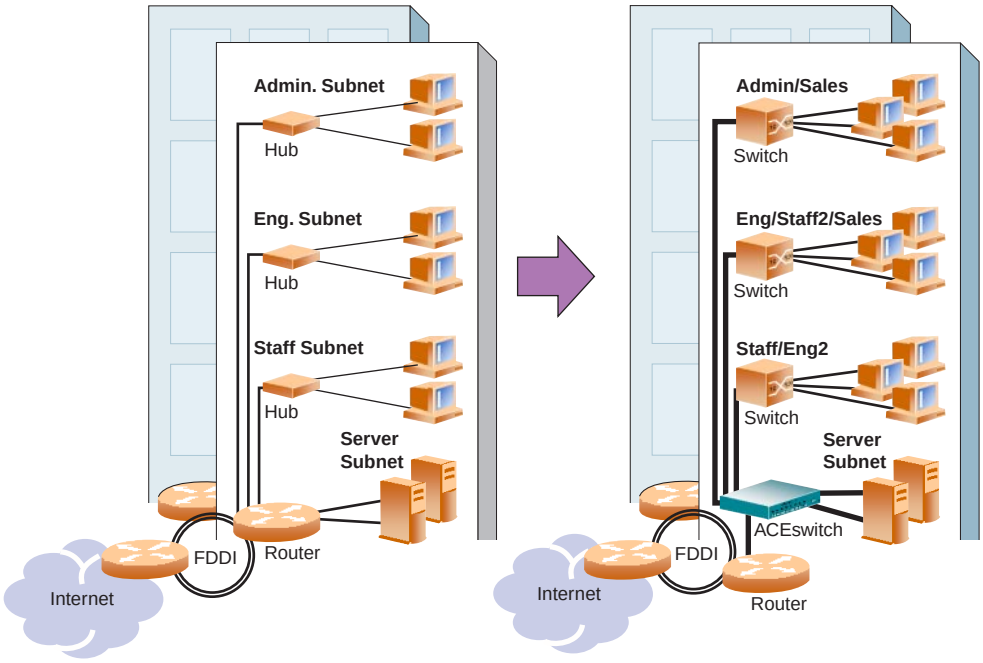

Figure 13-1 The Router Legacy Network

In this example, a corporate campus has migrated from a router-centric topology to a faster, more powerful switch-based topology. As is often the case, the legacy of network growth and redesign has left the system with a hodge-podge of illogically distributed subnets. This is a situation that switching alone cannot cure. Instead, the router is flooded with cross-subnet communication. This compromises efficiency in two ways:

- Routers can be slower than switches. The cross-subnet side trip from the switch to the router and back again adds two hops for the data, slowing throughput considerably.
- Traffic to the router increases, worsening any congestion.

Even if every end-station on the network could be moved to better logical subnets (a daunting task), competition for access to common server pools on different subnets still burdens the routers.

This problem is solved by using Alteon Networks switches with built-in IP Routing capabilities. Cross-subnet LAN traffic can now be routed within the ACElerate powered switches with wire speed Layer 2 switching performance. This not only eases the load on the router, but saves the network administrators from reconfiguring each and every end-station with new IP addresses.

Take a closer look at the ACEswitch 180 in the example configuration:

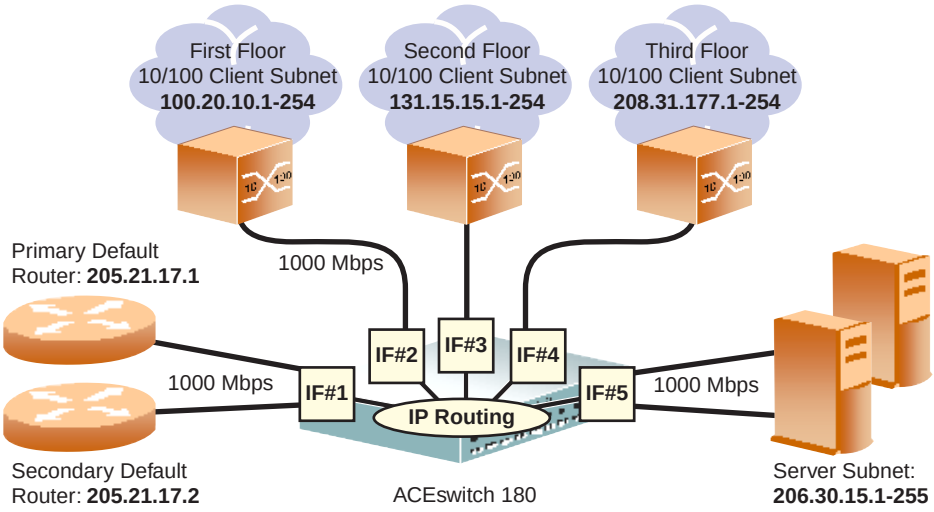

Figure 13-2 Switch-Based Routing Topology

The ACEswitch 180 connects the Gigabit Ethernet and Fast Ethernet trunks from various switched subnets throughout one building. Common servers are placed on another subnet attached to the switch. A primary and backup router are attached to the switch on yet another subnet.

Without Layer 3 IP Routing on the switch, cross-subnet communication is relayed to the default gateway (in this case, the router) for the next level of routing intelligence. The router fills in the necessary address information and sends the data back to the switch, which relays the packet to the proper destination subnet using Layer 2 switching.

With Layer 3 IP Routing in place on the ACEswitch, routing between different IP subnets can be accomplished entirely within the switch. This leaves the routers free to handle inbound and outbound traffic for this group of subnets.

As an added benefit, UDP Jumbo Frame traffic is automatically fragmented to regular Ethernet frame sizes when routing to non-Jumbo Frame subnets. For instance, this allows servers to communicate with each other using Jumbo Frames, and to non-Jumbo Frame devices using regular frames, all transparently to the user.

Example ACEswitch 180 Configuration for Subnet Routing

Prior to configuration, you must be connected to the switch command-line interface as the administrator (see Chapter 2, "The Command-Line Interface").

NOTE – For details about any of the menu commands described in this example, see "Configuring IP Parameters" on page 7-9.

---

1. **Assign an IP address (or document the existing one) for each real server, router, and client workstation.**

   In our example topology in Figure 13-2 on page 13-3, the following IP addresses are used:

Table 13-1 Subnet Routing Example: IP Address Assignments

Subnet	Devices	IP Addresses
#1	Primary and Secondary Default Routers	205.21.17.1 and 205.21.17.2
#2	First Floor Client Workstations	100.20.10.1-254
#3	Second Floor Client Workstations	131.15.15.1-254
#4	Third Floor Client Workstations	208.31.177.1-254
#5	Common Servers	206.30.15.1-254

2. **On the switch, assign an IP interface for each subnet attached to the switch.**

   Since there are five IP subnets connected to the switch, five IP interfaces are needed:

Table 13-2 Subnet Routing Example: IP Interface Assignments

Interface	Devices	IP Interface Address
IF #1	Primary and Secondary Default Routers	205.21.17.3
IF #2	First Floor Client Workstations	100.20.10.16
IF #3	Second Floor Client Workstations	131.15.15.1
IF #4	Third Floor Client Workstations	208.31.177.2
IF #5	Common Servers	206.30.15.200

These are configured using the following commands at the CLI:

```
>> Main# /cfg/ip/if 1                    (Select IP interface 1)
>> IP Interface 1# addr 205.21.17.3      (Assign IP address for the interface)
>> IP Interface 1# ena                   (Enable IP interface 1)
>> IP Interface 1# ../if 2               (Select IP interface 2)
>> IP Interface 2# addr 100.20.10.16     (Assign IP address for the interface)
>> IP Interface 2# ena                   (Enable IP interface 2)
>> IP Interface 2# ../if 3               (Select IP interface 3)
>> IP Interface 3# addr 131.15.15.1      (Assign IP address for the interface)
>> IP Interface 3# ena                   (Enable IP interface 3)
>> IP Interface 3# ../if 4               (Select IP interface 4)
>> IP Interface 4# addr 208.31.177.2     (Assign IP address for the interface)
>> IP Interface 4# ena                   (Enable IP interface 4)
>> IP Interface 4# ../if 5               (Select IP interface 5)
>> IP Interface 5# addr 206.30.15.200    (Assign IP address for the interface)
>> IP Interface 5# ena                   (Enable IP interface 5)
```

3. **Set each server and workstation's default gateway to point to the appropriate switch IP interface (the one in the same subnet as the server or workstation).**

4. **On the switch, configure the default gateways to point to the routers.**

   This allows the switch to send outbound traffic to the routers:

```
>> IP Interface 5# /cfg/ip/gw 1          (Select primary default gateway)
>> Default gateway 1# addr 205.21.17.1   (Point to primary router)
>> Default gateway 1# ena                (Enable primary default gateway)
>> Default gateway 1# ../gw 2            (Select secondary default gateway)
>> Default gateway 2# addr 205.21.17.2   (Point to secondary router)
>> Default gateway 2# ena                (Enable secondary default gateway)
```

5. **On the switch, enable, apply, and verify the configuration.**

```
>> Default gateway 2# ../fwrd            (Select the IP Forwarding Menu)
>> IP Forwarding# on                     (Turn IP forwarding on)
>> IP Forwarding# apply                  (Make your changes active)
>> IP Forwarding# ../cur                 (View current IP settings)
```

   Examine the resulting information. If any settings are incorrect, make any appropriate changes.

6. **On the switch, save your new configuration changes.**

```
>> IP# save                              (Save for restore after reboot)
```

Another Option: Adding VLANs to the Routing Example

The routers, servers, and clients in the example above are all in the same broadcast domain. If limiting broadcasts is desired in your network, you could use VLANs to create distinct broadcast domains. For example, you could create one VLAN for the routers, one for the servers, and one for the client trunks.

In this exercise, we are adding to the previous configuration.

1. **Determine which switch ports and IP interfaces belong to which VLANs.**

   The following table adds ports and VLANs information:

Table 13-3 Subnet Routing Example: Optional VLAN Ports

VLAN	Devices	IP Interface	Switch Port
#1	First Floor Client Workstations	3	1
	Second Floor Client Workstations	4	2
	Third Floor Client Workstations	5	3
#2	Primary Default Router	1	4
	Secondary Default Router	2	5
#3	Common Servers #1	6	6
	Common Servers #2	7	7

2. **On the switch, set the default VLAN for each port:**

```
>> # /cfg/port 1                    (Select port for First Floor)
>> Port 1# pvid 1                   (Set default to VLAN 1)
>> Port 1# ../port 2                (Select port for Second Floor)
>> Port 2# pvid 1                   (Set default to VLAN 1)
>> Port 2# ../port 3                (Select port for Third Floor)
>> Port 3# pvid 1                   (Set default to VLAN 1)
>> Port 3# ../port 4                (Select port for default router 1)
>> Port 4# pvid 2                   (Set default to VLAN 2)
>> Port 4# ../port 5                (Select port for default router 2)
>> Port 5# pvid 2                   (Set default to VLAN 2)
>> Port 5# ../port 6                (Select port for common server 1)
>> Port 6# pvid 3                   (Set default to VLAN 3)
>> Port 6# ../port 7                (Select port for common server 2)
>> Port 7# pvid 3                   (Set default to VLAN 3)
```

3. **On the switch, enable the VLANs.**

```
>> Port 7# /cfg/vlan 1                    (Select VLAN 1, the client VLAN)
>> VLAN 1# ena                            (enable VLAN 1)
>> VLAN 1# ../vlan 2                      (Select VLAN 2, the def. router VLAN)
>> VLAN 2# ena                            (enable VLAN 2)
>> VLAN 2# ../vlan 3                      (Select VLAN 3, the server VLAN)
>> VLAN 3# ena                            (enable VLAN 3)
```

4. **On the switch, add each IP interface to the appropriate VLAN.**

Now that the ports are separated into three VLANs, the IP interface for each subnet must be placed in the appropriate VLAN. From Table 13-3 on page 13-6, the settings are made as follows:

```
>> VLAN 3# /cfg/ip/if 1                   (Select IP interface 1 for def. routers)
>> IP Interface 1# vlan 2                 (Set to VLAN 2)
>> IP Interface 1# ../if 2                (Select IP interface 2 for first floor)
>> IP Interface 2# vlan 1                 (Set to VLAN 1)
>> IP Interface 2# ../if 3                (Select IP interface 3 for second floor)
>> IP Interface 3# vlan 1                 (Set to VLAN 1)
>> IP Interface 3# ../if 4                (Select IP interface 4 for third floor)
>> IP Interface 4# vlan 1                 (Set to VLAN 1)
>> IP Interface 4# ../if 5                (Select IP interface 5 for servers)
>> IP Interface 5# vlan 3                 (Set to VLAN 3)
```

5. **On the switch, apply and verify the configuration.**

```
>> IP Interface 5# apply                  (Make your changes active)
>> IP Interface 5# /info/vlan             (View current VLAN information)
>> Information# port                      (View current port information)
```

Examine the resulting information. If any settings are incorrect, make the appropriate changes.

6. **On the switch, save your new configuration changes.**

```
>> Information# save                      (Save for restore after reboot)
```

CHAPTER 14

Port Trunking

Port Trunking Overview

Basics

Trunk groups can provide super-bandwidth, multi-link connections between Alteon Networks switches or other trunk-capable devices. A “trunk group” is a group of ports that act together, combining their bandwidth to create a single, larger virtual link.

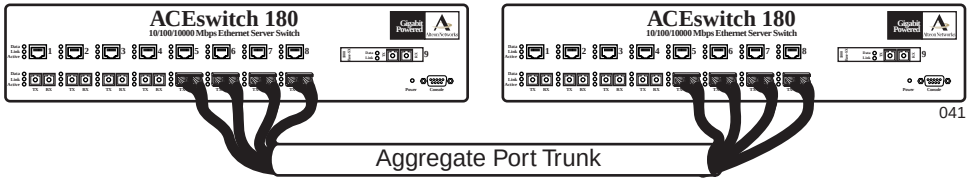

Figure 14-1 Port Trunk Group

When using port trunk groups between two A180 switches, for example, the network administrator can create a virtual link between the switches operating up to 4 Gigabits per second, depending on how many physical ports are combined. The switch supports up to 4 trunk groups per switch, each with 2 to 4 links.

Trunk groups are also useful for connecting an Alteon Networks switch to third-party devices that support link aggregation, such as Cisco routers and switches with EtherChannel technology (*not* ISL Trunking technology), and Sun's Quad Fast Ethernet Adapter. Alteon Network’s trunk group technology is compatible with these devices when they are configured manually.

Statistical Load Distribution

Network traffic is statistically load balanced between the ports in a trunk group. The ACElerate powered switch uses both the Layer 2 MAC address and Layer 3 IP address information present in each transmitted frame for determining load distribution.

The addition of Layer 3 IP address examination is an important advance for traffic distribution in trunk groups. In some port trunking systems, only Layer 2 MAC addresses are considered in the distribution algorithm. Each packet's particular combination of source and destination MAC addresses results in selecting one line in the trunk group for data transmission. If there are enough Layer 2 devices feeding the trunk lines, then traffic distribution becomes relatively even. In some topologies, however, only a limited number of Layer 2 devices (such as a handful of routers and servers) feed the trunk lines. When this occurs, the limited number of MAC address combinations encounterd results in a lopsided traffic distribution that can reduce the effective combined bandwidth of the trunked ports.

By adding Layer 3 IP address information to the distribution algorithm, a far wider variety of address combinations is seen. Even with just a few routers feeding the trunk, the normal source/destination IP address combinations (even within a single LAN) can be widely varied. This results in a wider statistical load distribution and maximizes the use of the combined bandwidth available to trunked ports.

Built-In Fault Tolerance

Since each trunk group is comprised of multiple physical links, the trunk group is inherently fault tolerant. As long as one connection between the switches is available, the trunk remains active.

Statistical load balancing is maintained whenever a port in a trunk group is lost or returned to service.

Port Trunking Example

In this example, three ports will be trunked between two ACEswitch 180s.

Prior to configuring each switch in this example, you must connect to the appropriate switch's command-line interface as the administrator (see Chapter 2, "The Command-Line Interface").

NOTE – For details about any of the menu commands described in this example, see "Configuring Port Trunking" on page 7-52.

1. **Connect the switch ports which will be involved in the trunk group:**

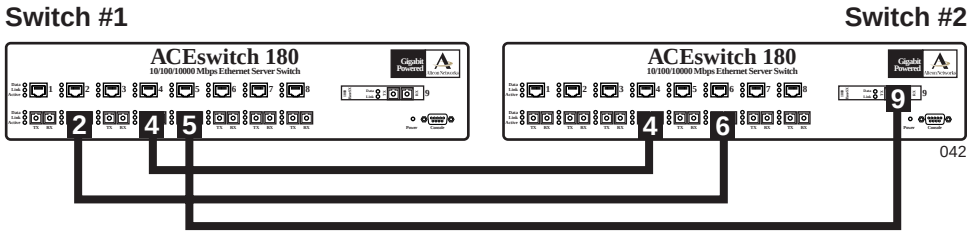

Figure 14-2 Example Port Trunk Group Configuration

2. **On Switch #1, define a Trunk Group.**

```
>> Main # /cfg/trunk 1                    (Select trunk group #1)
>> Trunk group 1# add 2                   (Add port 2 to trunk group #1)
>> Trunk group 1# add 4                   (Add port 4 to trunk group #1)
>> Trunk group 1# add 5                   (Add port 5 to trunk group #1)
>> Trunk group 1# ena                     (Enable trunk group #1)
```

3. **On Switch #1, apply and verify the configuration.**

```
>> Trunk group 1# apply                   (Make your changes active)
>> Trunk group 1# cur                     (View current trunking configuration)
```

Examine the resulting information. If any settings are incorrect, make the appropriate changes.

4. **On Switch #1, save your new configuration changes.**

```
>> Trunk group 1# save                    (Save for restore after reboot)
```

5. **On Switch #2, repeat the process.**

```
>> Main # /cfg/trunk 3                    (Select trunk group #3)
>> Trunk group 3# add 4                   (Add port 4 to trunk group #3)
>> Trunk group 3# add 6                   (Add port 6 to trunk group #3)
>> Trunk group 3# add 9                   (Add port 9 to trunk group #3)
>> Trunk group 3# ena                     (Enable trunk group #3)
>> Trunk group 3# apply                   (Make your changes active)
>> Trunk group 3# cur                     (View current trunking configuration)
>> Trunk group 3# save                    (Save for restore after reboot)
```

Switch #1 trunk group #1 is now connected to Switch #2 trunk group #3.

NOTE – In this example, both switches are Alteon Networks switches. If a third-party device supporting link aggregation is used (such as Cisco routers and switches with EtherChannel technology, or Sun's Quad Fast Ethernet Adapter), then trunk groups on the third-party device should be configured manually. Connection problems could arise when using automatic trunk group negotiation on the third-party device.

---

6. **Examine the trunking information on each switch.**

```
>> /info/trunk                            (View trunking information)
```

Information about each port in each configured trunk group will be displayed. Make sure that trunk groups consist of the expected ports, and that each port is in the expected state.

CHAPTER 15

Server Load Balancing

This chapter describes how to configure and use the optional Layer 4 software for Server Load Balancing. For information on activating this optional software if required, see "Activating Optional Software" on page 8-6.

Server Load Balancing Overview

Benefits

Server Load Balancing benefits your network in a number of ways:

- Increased efficiency for server utilization and network bandwidth

  With Server Load Balancing, your ACElerate powered switch is aware of the shared services provided by your server pool. The switch can then balance user session traffic among the available servers. For even greater control, traffic is distributed according to a variety of user-selectable rules.

  By helping to eliminate server over-utilization, important session traffic gets through more easily, reducing user competition for connections on overworked servers.

- Increased reliability of services to users

  If any server in a server pool fails, the remaining servers continue to provide access to vital applications and data. The failed server can be brought back up without interrupting access to services.

- Increased scalability of services

  As users are added and the server pool's capabilities are saturated, new servers can be added to the pool transparently.

Identifying Your Needs

Server Load Balancing may be the right option for addressing these vital network concerns:

- A single server no longer meets the demand for its particular application.
- The connection from your LAN to your server overloads the server's capacity.
- Your NT and UNIX servers hold critical application data and must remain available even in the event of a server failure.
- Your web site is vital, being used as a way to do business and for taking orders from customers. It must not become overloaded or unavailable.
- You want to use multiple servers or hot-standby servers for maximum network uptime.
- You must be able to scale your applications to meet client and LAN request capacity.
- You can't afford to continue using an inferior load balancing technique such as DNS Round Robin, or a software-only system.

How Server Load Balancing Works

In an average network that employs multiple servers without server load balancing, each server usually specializes in providing one or two unique services. If one of these servers provides access to applications or data which is in high demand, it can become overutilized. Placing this kind of strain on a server can decrease the performance of the entire network as user requests are rejected by the server and then resubmitted by the user stations. Ironically, over-utilization of key servers often happens in networks where other servers are actually under-utilized.

The solution to getting the most from your servers is the Layer 4 switching feature of Server Load Balancing. With this software feature, your switch is aware of the services provided by each server, and can direct user session traffic to the appropriate server based on a variety of balancing algorithms.

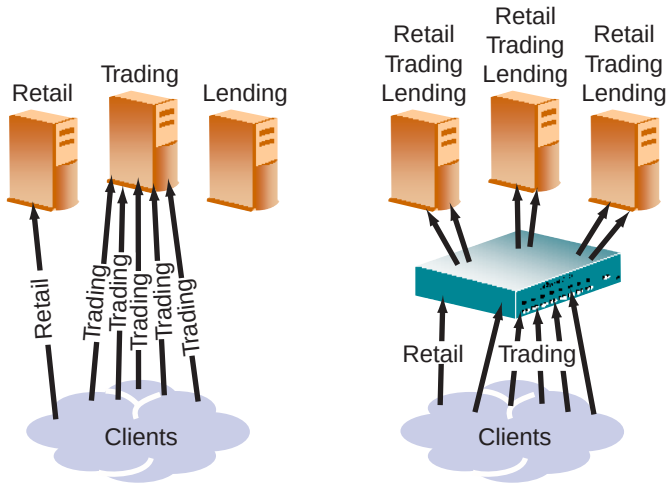

Figure 15-1 Traditional vs. Server Load Balanced network configurations

To provide Server Load Balancing for any particular type of service, each server in the pool must have access to identical content, either directly (duplicated on each server) or through a back-end network (mounting the same file system or database server).

The switch with Layer 4 software acts as a front-end to the servers, interpreting user session requests and distributing them among the available servers. To accomplish this, the switch is configured to act as a virtual server and is given a virtual IP address (or range of addresses) for each collection of services it will distribute. There can be as many as 256 virtual servers on the switch, each distributing up to eight different services.

Each virtual server is assigned a list of the real IP addresses (or range of addresses) of the real servers in the pool where its services reside. When the user stations request connections to a service, they will communicate with a virtual server on the switch. When the switch receives the request, it binds the session request to the real IP address of the best available real server, and remaps the fields in each frame from virtual addresses to real addresses.

Network Topology Considerations

When deploying Layer 4 switching features, there are a number of key aspects to consider:

- All client requests to a virtual IP address and all responses from the real servers *must* pass through the switch. If alternate paths exist between the client and the real servers (as shown below), the Layer 4 switch can be configured with proxies in order to guarantee that Layer 4 traffic uses the correct path (see "IP Proxy Addresses for Complex Networks" on page 15-19).

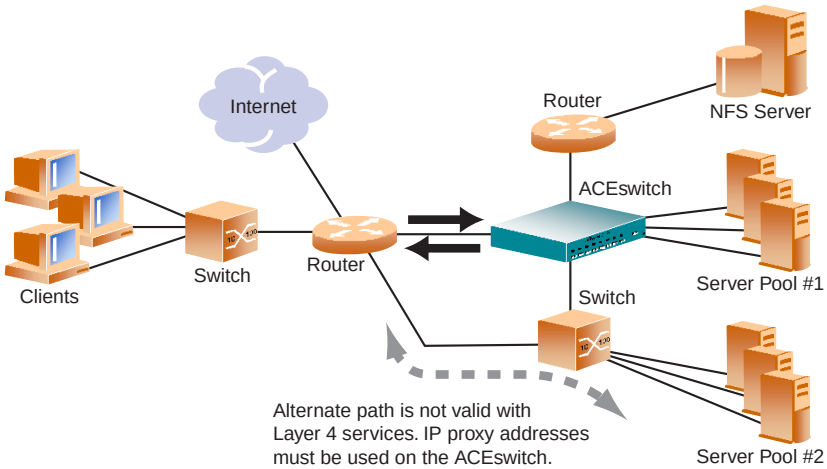

Figure 15-2 Client/Server traffic must pass through the ACEswitch

- Identical content must be available to each server in the same pool. Either of these methods can be used:
  - Static applications and data are duplicated on each real server in the pool.
  - Each real server in the pool has access to the same data through use of a shared file system or back-end database server.
- Some services require that a series of client requests go to the same real server so that session-specific state data can be retained between connections. Services of this nature include web search results, multi-page forms that the user fills in, or custom web-based applications typically created using `cgi-bin` scripts. Connections for these types of services must be configured as "persistent" (see the `pbind` option in Table 7-24 on page 7-40), or must use the `minmisses` or `hash` metrics (see "Server Load Balancing Metrics" on page 7-36).

- Clients and servers cannot be connected through the same switch port. Each port in use on the switch must be configured for one, and only one, of the following network topologies:
  - Layer 4 server port. This represents ports where real servers are connected to the Layer 4 switch, directly or through a hub, router, or another switch. Real server responses to client requests are processed on these ports. These ports must not lead to client devices. These ports can simultaneously provide Layer 2 switching and IP Routing functions.
  - Layer 4 client port. This represents ports where client request traffic is processed. Address translation from the virtual IP to the real server IP address occurs here. Maximizing the number of these ports on the Layer 4 switch will improve the switch's potential for effective Server Load Balancing. These ports can also simultaneously provide Layer 2 switching and IP Routing functions.
  - Non-Layer 4 port. The port does not process either client requests or real server responses. An example of such port usage is for an NFS server which the real servers access for data, and also real servers which are accessed through proxy IP addresses.

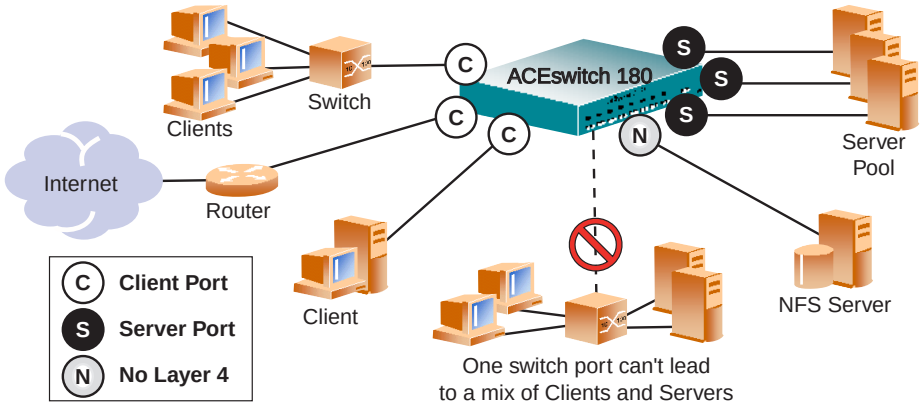

Figure 15-3 Port designations using Layer 4 Server Load Balancing

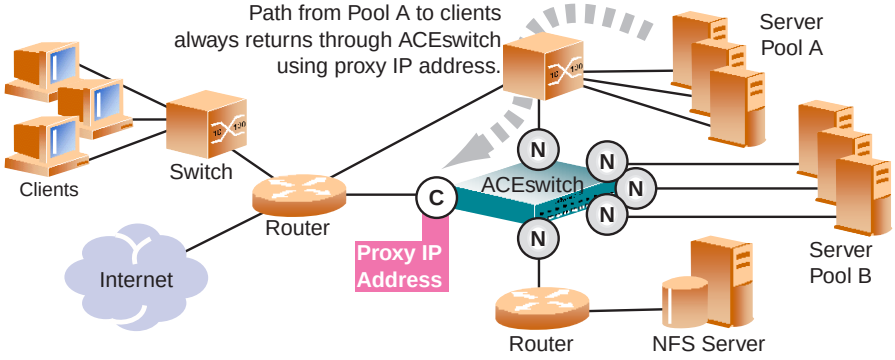

Figure 15-4 Port designations using proxy IP addresses

Server Load Balancing Examples

Web Hosting Configuration

Consider a situation where customer web sites are being hosted by a popular web hosting company and/or Internet Service Provider (ISP). The web content is relatively static and is kept on a single NFS server for easy administration. As the customer base increases, so does the number of simultaneous web connection requests.

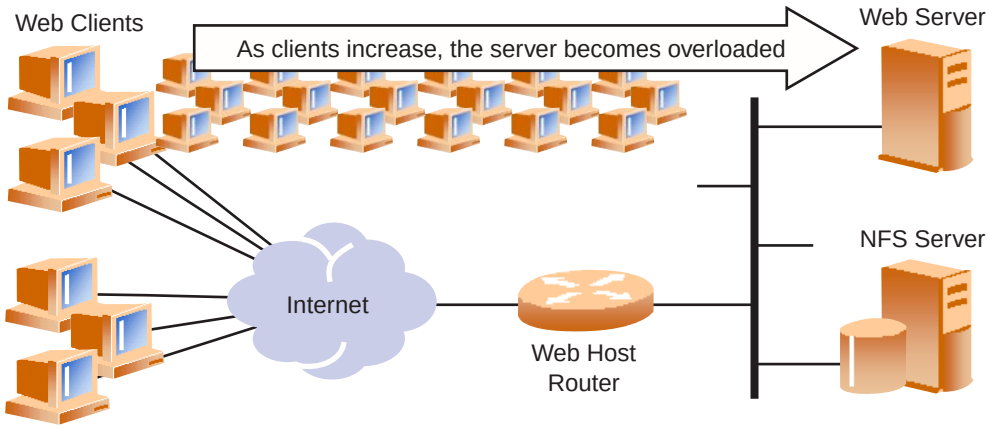

Figure 15-5 Web hosting configuration without Layer 4 switching

Such a company has three primary needs:

- Increased server availability
- Server performance scalable to match new customer demands
- Easy administration of network and servers

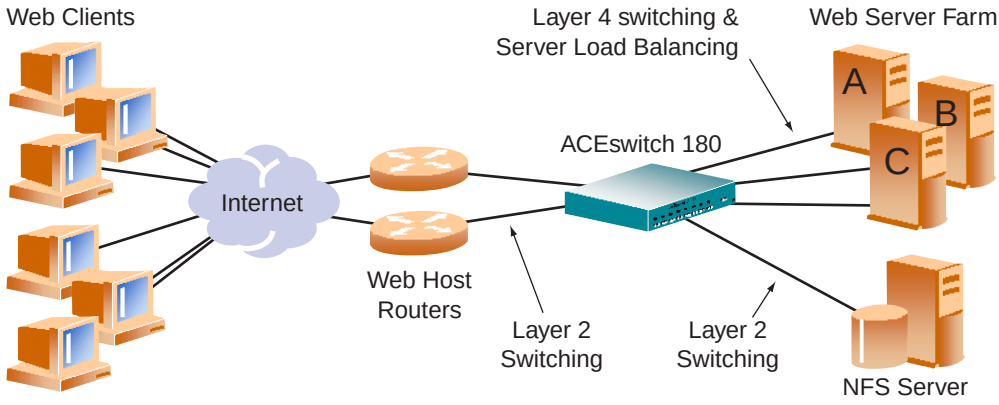

Figure 15-6 Web hosting with Layer 4 solutions

Each concern about this company's site can be addressed by adding an ACEswitch 180 with optional Layer 4 software.

- Reliability is increased by providing multiple paths from the clients to the Layer 4 switch, and by access to a pool of servers that have identical content. If one server fails, the others can take up the additional load.
- Performance is improved by balancing the web request load across multiple servers. More servers can be added at any time to increase processing power.
- For ease of maintenance, servers can be added or removed dynamically without interrupting shared services.

Example ACEswitch 180 Configuration for the Web Hosting Solution

In the following examples, many of the Server Load Balancing options are left to their default values. See "Additional Server Load Balancing Options" on page 15-17 for more options.

The following is required prior to configuration:

- You must be connected to the switch command-line interface as the administrator (see Chapter 2, "The Command-Line Interface").
- The optional Server Load Balancing software must be enabled (see "Activating Optional Software" on page 8-6).

NOTE – For details about any of the menu commands described in this example, see "Configuring Server Load Balancing" on page 7-29.

---

1. **Assign an IP address to each of the real servers in the server pool.**

   The real servers in any given real server group must have an IP route to the switch that will perform the Server Load Balancing functions. This is most easily accomplished by placing the switches and servers on the same IP subnet, although advanced routing techniques can be used as long as they do not violate the topology rules outlined in "Network Topology Considerations" on page 15-4.

   For this example, the three web-host real servers have the following IP addresses on the same IP subnet:

   Table 15-1 Web Host Example: Real Server IP addresses

Real Server	IP address
Server A	200.200.200.2
Server B	200.200.200.3
Server C	200.200.200.4

2. **Define an IP interface on the switch.**

   The switch must have an IP route to all of the real servers which receive Layer 4 switching services. For Server Load Balancing, the switch uses this path to determine the level of TCP/IP reachability of the real servers.

   To configure an IP interface for this example, enter this command from the CLI:

```
>> Main# /cfg/ip/if 1                          (Select IP interface #1)
>> IP Interface 1# addr 200.200.200.100       (Assign IP address for the interface)
>> IP Interface 1# ena                        (Enable IP interface #1)
```

NOTE – The IP interface and the real servers must belong to the same VLAN. This example assumes that all ports and IP interfaces use default VLAN #1, requiring no special VLAN configuration for the ports or IP interface.

3. **On the switch, define each Real Server.**

   For each real server, you must assign a real server number, specify its actual IP address, and enable the real server. For example:

```
>> IP Interface 1# /cfg/slb/real 1          (Server A is real server 1)
>> Real server 1 # rip 200.200.200.2        (Assign Server A IP address)
>> Real server 1 # ena                      (Enable real server 1)
>> Real server 1 # ../real 2                (Server B is real server 2)
>> Real server 2 # rip 200.200.200.3        (Assign Server B IP address)
>> Real server 2 # ena                      (Enable real server 2)
>> Real server 2 # ../real 3                (Server C is real server 3)
>> Real server 3 # rip 200.200.200.4        (Assign Server C IP address)
>> Real server 3 # ena                      (Enable real server 3)
```

4. **On the switch, define a Real Server Group.**

   This combines the three real servers into one service group:

```
>> Real server 3 # /cfg/slb/group 1         (Select real server group 1)
>> Real server group 1# add 1               (Add real server 1 to group 1)
>> Real server group 1# add 2               (Add real server 2 to group 1)
>> Real server group 1# add 3               (Add real server 3 to group 1)
```

5. **On the switch, define a Virtual Server.**

   All client requests will be addressed to a virtual IP on a virtual server defined on the switch. Clients acquire the virtual IP through normal DNS resolution. HTTP uses well-known TCP port 80. In this example, HTTP is configured as the only service running on this virtual IP, and is associated with our real server group. For example:

```
>> Real server group 1 # /cfg/slb/virt 1    (Select virtual server 1)
>> Virtual server 1# vip 200.200.200.1      (Assign a virtual server IP address)
>> Virtual server 1# add http 1             (Associate virtual port to real group)
>> Virtual server 1# ena                    (Enable the virtual server)
```

NOTE – This configuration is not limited to HTTP web service. Other TCP/IP services can be configured in a similar fashion. For a list of other well-known services and ports, see the command option information on page 7-39.

6. **On the switch, define the Port States.**

In this example, the following ports are being used on the ACEswitch 180:

Table 15-2 Web Host Example: ACEswitch 180 Port Usage

Port	Host	Port Setting
1	Server A	`server`
2	Server B	`server`
3	Server C	`server`
4	Back-end NFS server. All three real servers get their web content from this machine. This port does not require Layer 4 switching.	`none`
5	Client router A. This connects the switch to the Internet where all client requests originate.	`client`
6	Client router B. This also connects the switch to the Internet where all client requests originate.	`client`

The ports are configured as follows:

```
>> Virtual server 1# /cfg/slb/port 1          (Select physical switch port 1)
>> SLB port 1# state server                   (Assign port 1 to server traffic)
>> SLB port 1# ../port 2                      (Select physical switch port 2)
>> SLB port 2# state server                   (Assign port 2 to server traffic)
>> SLB port 2# ../port 3                      (Select physical switch port 3)
>> SLB port 3# state server                   (Assign port 3 to server traffic)
>> SLB port 3# ../port 4                      (Select physical switch port 4)
>> SLB port 4# state none                     (Assign port 4 to no Layer 4 for NFS)
>> SLB port 4# ../port 5                      (Select physical switch port 5)
>> SLB port 5# state client                   (Set port 5 for client traffic)
>> SLB port 5# ../port 6                      (Select physical switch port 6)
>> SLB port 6# state client                   (Set port 6 for client traffic)
```

7. **On the switch, enable, apply, and verify the configuration.**

```
>> SLB port 6# ..                             (Select the SLB Menu)
>> Server Load Balancing# on                  (Turn Server Load Balancing on)
>> Server Load Balancing# apply               (Make your changes active)
>> Server Load Balancing# cur                 (View current settings)
```

Examine the resulting information. If any settings are incorrect, make any appropriate changes.

8. **On the switch, save your new configuration changes.**

```
>> Server Load Balancing# save          (Save for restore after reboot)
```

9. **On the switch, check the Server Load Balancing information.**

```
>> Server Load Balancing# /info/slb     (View SLB information)
```

Check that all Server Load Balancing parameters are working according to expectation. If necessary, make any appropriate configuration changes and then check the information again.

High-Availability Web Application Configuration

Consider a situation where a corporation depends on intranet access to an important database application. They can't tolerate any downtime, and the system must remain easy to maintain.

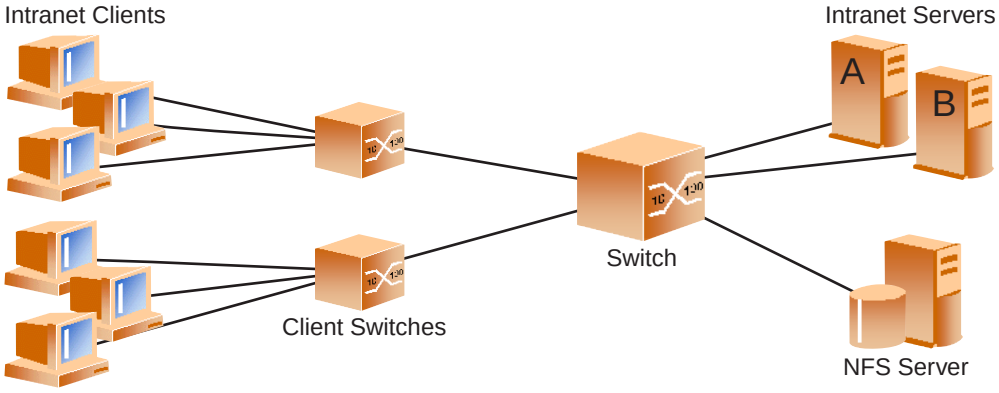

Figure 15-7 Intranet configuration without redundancy

Such a company has three primary needs:

- Server Load Balancing to increase performance
- Hot Standby on each ACEswitch, and Dual Homing NICs in each server to provide network redundancy
- Hot Standby (or Backup) servers to ensure reliability down to a single server

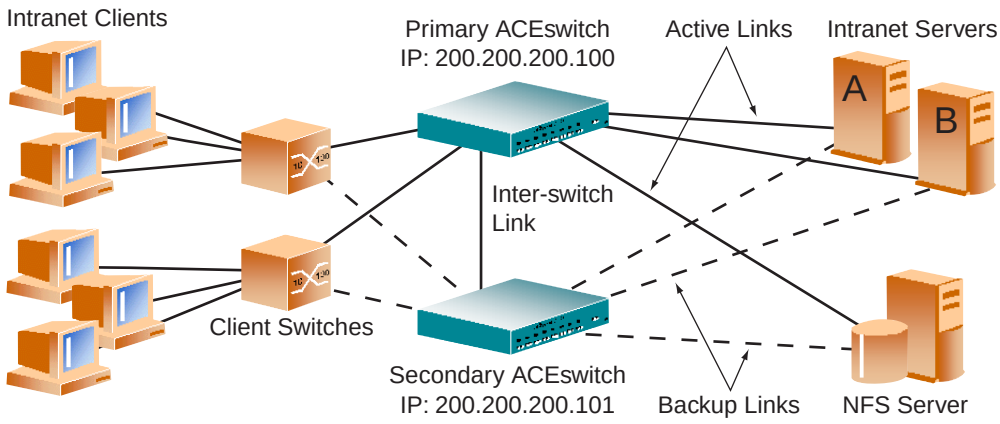

Figure 15-8 Intranet configuration with ACEswitch high-availability solution

In this example, high-availability is accomplished as follows:

- Each real server is attached to the same port number on both ACEswitches using Alteon Networks' ACEnics configured with Dual Homing to maintain server IP addresses. Alternately, dual third-party NICs with different IP addresses can be used in the servers.
- Both switches are directly connected to each other to ensure a fast data path at failover.
- The switch configured as the primary switch downloads its configuration to the secondary.
- Single link failover is resolved within one second as the primary switch instructs the secondary to turn on a standby link.
- Primary switch or multiple link failover is resolved within two seconds.

Example ACEswitch 180 Configuration for the Intranet Solution

In the following examples, many of the Server Load Balancing options are left to their default values. See "Additional Server Load Balancing Options" on page 15-17 for more options.

The following is required prior to configuration:

- When called to configure either switch, you must be connected to the appropriate switch's command-line interface as the administrator (see Chapter 2, "The Command-Line Interface").
- Optional Layer 4 software must be enabled (see "Activating Optional Software" on page 8-6).

NOTE – For details about any of the menu commands described in this example, see "Configuring Server Load Balancing" on page 7-29.

---

1. **Assign an IP address to each of the real servers in the server pool.**

   The real servers in any given real server group must be in the same VLAN and must have an IP route to the switches that will perform the Server Load Balancing functions. This is most easily accomplished by placing the switches and servers on the same IP subnet, although advanced routing techniques can be used as long as they do not violate the topology rules outlined in "Network Topology Considerations" on page 15-4.

   For this example, the real servers have the following IP addresses on the same IP subnet:

   Table 15-3 High-Availability Example: Real Server IP addresses

Real Server	IP address
Server A	200.200.200.2
Server B	200.200.200.3

NOTE – When third-party NICs without Dual Homing are used, the real servers will have two IP addresses: one for each installed NIC. To account for this, configure each NIC IP address as a different real server. Configure the primary as outlined below, and configure the secondary as a backup server (see "Backup/Overflow Servers" on page 15-18).

2. **On the both switches, define an IP interface on the switch.**

   Each switch must have an IP route to all of the real servers which receive Layer 4 services.

   Table 15-4 High-Availability Example: ACEswitch 180 IP addresses

ACEswitch 180	IP interface address
Primary Switch	200.200.200.100
Secondary Switch	200.200.200.101

   For example, from the Main Menu on the primary switch:

```
>> Main# /cfg/ip/if 1                          (Select IP interface #1)
>> IP Interface 1# addr 200.200.200.100       (Assign IP address for the interface)
>> IP Interface 1# ena                        (Enable IP interface #1)
```

   And from the Main Menu on the secondary switch:

```
>> Main# /cfg/ip/if 1                          (Select IP interface #1)
>> IP Interface 1# addr 200.200.200.101       (Assign IP address for the interface)
>> IP Interface 1# ena                        (Enable IP interface #1)
```

3. **On the primary switch, define each Real Server.**

   For each real server, you must assign a server number and its IP address, and you must enable the server. For example:

```
>> IP Interface 1# /cfg/slb/real 1          (Server A is real server 1)
>> Real server 1 # rip 200.200.200.2        (Assign Server A IP address)
>> Real server 1 # ena                      (Enable real server 1)
>> Real server 1 # ../real 2                (Server B is real server 2)
>> Real server 2 # rip 200.200.200.3        (Assign Server B IP address)
>> Real server 2 # ena                      (Enable real server 2)
```

4. **On the primary switch, define a Real Server Group.**

   This combines both real servers into one service group:

```
>> Real server 2 # /cfg/slb/group 1         (Select real server group 1)
>> Real server group 1# add 1               (Add real server 1 to group 1)
>> Real server group 1# add 2               (Add real server 2 to group 1)
```

5. **On the primary switch, define a Virtual Server.**

   All client requests will be addressed to a virtual IP on a virtual server defined on the switch. Clients acquire the virtual IP through normal DNS resolution. HTTP uses well-known TCP port 80. In this example, HTTP is configured as the only service running on this virtual IP, and is associated with our real server group. For example:

```
>> Real server group 1# /cfg/slb/virt 1     (Select virtual server 1)
>> Virtual server 1# vip 200.200.200.1      (Assign a virtual server IP address)
>> Virtual server 1# add http 1             (Associate virtual port to real group)
>> Virtual server 1# ena                    (Enable the virtual server)
```

NOTE – This configuration is not limited to HTTP web service. Other TCP/IP services can be configured in a similar fashion. For a list of other well-known services and ports, see the command option information on page 7-39.

6. On the primary switch, define the Port States.

In this example, the following ports are being used on both ACEswitch 180s:

Table 15-5 Web Host Example: ACEswitch 180 Port Usage

Port	Host	Port Setting
1	Server A	`server`
2	Server B	`server`
3	Back-end NFS server. All three real servers get their web content from this machine. This port does not require Layer 4 switching.	`none`
4	Client switch A. This connects the ACEswitch to client requests.	`client`
5	Client switch B. This also connects the ACEswitch to client requests.	`client`
6	Connection to secondary ACEswitch.	`failover`

NOTE – Physical port arrangement on the primary and secondary ACEswitches must be identical.

Configure the primary switch as follows:

```
>> Virtual server 1# /cfg/slb/port 1          (Select physical switch port 1)
>> SLB port 1# state server                   (Assign port 1 to server traffic)
>> SLB port 1# ../port 2                      (Select physical switch port 2)
>> SLB port 2# state server                   (Assign port 2 to server traffic)
>> SLB port 2# ../port 3                      (Select physical switch port 3)
>> SLB port 3# state none                     (Assign port 3 to no Layer 4 for NFS)
>> SLB port 3# ../port 4                      (Select physical switch port 5)
>> SLB port 4# state client                   (Assign port 4 to client traffic)
>> SLB port 4# ../port 5                      (Select physical switch port 6)
>> SLB port 5# state client                   (Assign port 5 to client traffic)
>> SLB port 5# ../port 6                      (Select physical switch port 6)
>> SLB port 6# state failover                 (Assign inter-switch protocol)
```

7. On the primary switch, set the failover parameters.

```
>> SLB port 6# ../fail                        (Select failover parameters)
>> SLB failover menu# prima 200.200.200.100   (Define primary switch)
>> SLB failover menu# secon 200.200.200.101   (Define secondary switch)
>> SLB failover menu# on                      (Enable hot-standby failover)
```

8. **On the primary switch, enable, apply, and verify the configuration.**

```
>> SLB failover menu# ..                    (Select the SLB Menu)
>> Server Load Balancing# on                (Turn Server Load Balancing on)
>> Server Load Balancing# apply             (Make your changes active)
>> Server Load Balancing# cur               (View current settings)
```

Examine the resulting information. If any settings are incorrect, make any appropriate changes.

9. **On the primary switch, save your new configuration changes.**

```
>> Server Load Balancing# save              (Save for restore after reboot)
```

10. **On the secondary switch, set the failover parameters.**

```
>> Main# /cfg/slb/fail                      (Select failover parameters)
>> SLB failover menu# prima 200.200.200.100(Define primary switch)
>> SLB failover menu# secon 200.200.200.101(Define secondary switch)
>> SLB failover menu# on                    (Enable hot-standby failover)
>> SLB failover menu# ..                    (Select SLB Menu)
>> Server Load Balancing# on                (Enable Layer 4 processing)
>> Server Load Balancing# apply             (Make your changes active)
```

11. **On each switch, check the Server Load Balancing information.**

```
>> Server Load Balancing# /info/slb         (View SLB information)
```

Check that all Server Load Balancing parameters are working according to expectation. If necessary, make any appropriate configuration changes and then check the information again.

Additional Server Load Balancing Options

In the examples above, many of the Server Load Balancing options are left to their default values. The following configuration options can be used to tune the system.

NOTE – You must `apply` any changes in order for them to take effect, and you must `save` them if you wish them remain in effect after switch reboot.

Metrics for Real Server Groups

Metrics are used for selecting which real server will receive the next client connection (see page 7-36). The default metric is Least Connections (`leastconns`). To change a real server group metric, to `minmisses` for example, enter:

```
>> # /cfg/slb/group group-number                (Select the real server group)
>> Real server group# metrc minmisses           (Use round robin metric)
```

Weights for Real Servers

Weights can be assigned to each real server. These weights bias load balancing to give the fastest real servers a bigger share of connections during load balancing. Weight is specified as a number from 1 (the default) to 48. Each increment increases the number of connections the real server gets. By default, each real server is given a weight setting of 1. A setting of 10 would assign the server roughly 10 times the number of connections as a server with a weight of 1. To set weights, enter the following commands:

```
>> # /cfg/slb/real real-server-number           (Select the real server)
>> Real server# wght 10                         (8 times the number of connections)
```

Connection Time-outs for Real Servers

In some cases, open TCP/IP sessions are not closed properly (for example, the switch receives the `SYN` for the session, but no `FIN` is sent). If a session is inactive for 10 minutes (the default), it is released from the switch. To change the time-out period, enter the following:

```
>> # /cfg/slb/real real-server-number           (Select the real server)
>> Real server# tmout 4                         (Specify an even numbered interval)
```

Maximum Connections for Real Servers

You can set the number of open connections each real server is allowed to handle for Server Load Balancing. To set the connection limit, enter the following:

```
>> # /cfg/slb/real real-server-number          (Select the real server)
>> Real server# mcon 1600                      (Allow 1600 connections maximum)
```

Values average between about 500 HTTP connections for slower servers to 1,500 for quicker, multi-processor servers. The appropriate value also depends on the duration each session lasts, as well as how much CPU capacity is occupied by processing each session. Connections that use a lot of Java or `CGI-bin` scripts for forms or searches require more server resources and thus a lower `mcon` limit. You may wish to use a performance bench-mark tool to determine how many connections your real servers can handle.

Health-Check Parameters for Real Servers

By default, the switch checks the status of each service on each real server every two seconds. Sometimes, the real server may be too busy processing connections to respond to health checks. By default, if a service does not respond to four consecutive health checks, the switch declares the service unavailable. Both the health check interval and the number of retries can be changed:

```
>> # /cfg/slb/real real-server-number          (Select the real server)
>> Real server# intr 4                         (Check real server every 4 seconds)
>> Real server# retry 6                        (If 6 consecutive health checks fail,
                                                declare real server down)
```

Backup/Overflow Servers

A real server can backup other real servers, and can handle overflow traffic when the maximum connection limit is reached. Each backup real server must be assigned a real server number and real IP address. It must then be enabled. Finally, the backup must be assigned to each real server it will backup. The following defines Real Server #4 as a backup for Real Servers #1 and #2:

```
>> # /cfg/slb/real 4                           (Select real server #4 as backup)
>> Real server 4 # rip 200.200.200.5           (Assign backup IP address)
>> Real server 4 # ena                         (Enable real server #4)
>> Real server 4 # ../real 1                   (Select real server #1)
>> Real server 1 # bkup 4                      (Real server #4 is backup for #1)
>> Real server 1 # ../real 2                   (Select real server #2)
>> Real server 2 # bkup 4                      (Real server #4 is backup for #2)
```

In a similar fashion, a backup/overflow server can be assigned to a real server group. If all real servers in a real server group fail or overflow, the backup comes online.

```
>> # /cfg/slb/group real-server-group-number      (Select real server group)
>> Real server group# bkup 4                      (Assign real server #4 as backup)
```

IP Proxy Addresses for Complex Networks

For proper Server Load Balancing, all client-to-server requests to a particular virtual server and all related server-to-client responses *must* pass through the *same* Layer 4 switch.

In complex network topologies, routers and other devices can create alternate paths around the switch managing Layer 4 functions (see Figure 15-2 on page 15-4). Under such conditions, the client switch ports must use a proxy IP address.

When the client requests services from the switch's virtual server, the client sends its own IP address for use as a return address. If a proxy IP address is configured for the client port on the switch, the switch replaces the client's source IP address with the switch's own proxy IP address before sending the request to the real server. This creates the illusion that the switch originated the request. The real server uses the switch's proxy IP address as the destination address for any response. This forces the Layer 4 traffic to return through the proper switch, regardless of alternate paths. Once the switch receives the proxied data, it puts the original client IP address into the destination address and sends the packet to the client.

NOTE – Because requests appear to come from the switch proxy IP address rather than the client source IP address, use of proxy addresses can generate misleading access information for network statistics or debugging.

The proxy IP address can also be used for direct access to the real servers (see "Direct Client Access to Real Servers" on page 7-41).

When implementing proxies, server ports should be reconfigured to use the "pnone" state, rather than "server." Re-examining example #1, the following port states are used:

Table 15-6 Proxy Example: ACEswitch 180 Port Usage

Port	Host	Port Setting
1	Server A	none
2	Server B	none
3	Server C	none
4	Back-end NFS server. All three real servers get their web content from this machine. This port does not require Layer 4 switching.	none
5	Client router A. This connects the switch to the Internet where all client requests originate.	client
6	Client router B. This also connects the switch to the Internet where all client requests originate.	client

The following commands are used for setting the port states:

```
>> # /cfg/slb/port 1                    (Select switch port #1)
>> SLB port 1# state none                (Set state for port #1)
>> SLB port 1# ../port 2                 (Select switch port #2)
>> SLB port 2# state none                (Set state for port #2)
>> SLB port 2# ../port 3                 (Select switch port #3)
>> SLB port 3# state none                (Set state for port #3)
```

Only the "client" ports require proxy IP addresses. Each proxy IP address must be unique on your network. The following shows commands used to configure proxies for this example:

```
>> # /cfg/slb/port 5                    (Select network port #5)
>> SLB port 5# pip 200.200.200.68        (Set proxy IP address for port #5)
>> SLB port 5# ../port 6                 (Select network port #6)
>> SLB port 6# pip 200.200.200.69        (Set proxy IP address for port #6)
```

The Layer 4 proxies are transparent to the user. No additional client configuration is needed.

NOTE – Remember that you must apply any changes in order for them to take effect, and you must save them if you wish them remain in effect after switch reboot. Also, the /info/slb command is useful for checking the state of Server Load Balancing operation.

CHAPTER 16

Filtering

This chapter describes how to configure and use Filtering for security and redirection applications.

NOTE – For Application Redirection, the optional Layer 4 software must be enabled (see "Filtering and Layer 4" on page 7-30).

Filtering Overview

Benefits

Layer 3 (IP) and Layer 4 (Application) filtering gives the network administrator a powerful tool with the following benefits:

- Filtering increases security for server networks.

  Filters can be configured to allow or deny traffic according to various IP address, protocol, and port criteria. This gives the administrator fine control over the types of traffic permitted through the switch. Optionally, any filter can generate `syslog` messages for increased security visibility.

- General Network Address Translation

  NAT can be used to map the source or destination IP address and port of private network traffic to/from an advertised network IP address and port.

- Application Redirection improves network bandwidth and provides unique network solutions.

  Filters can be created which redirect traffic to cache and application servers. Repeated client access to common web or application content across the Internet can be an inefficient use of network resources. By redirecting client requests to a local web-cache or application server, you increase the speed at which clients access the information and free up valuable network bandwidth.

Filtering Criteria

Up to 224 filters can be configured on the switch. Each filter can be set to allow, deny, redirect, or translate traffic based on any combination of the following criteria:

- Source IP Address or range
- Destination IP Address or range
- Protocol type (for example: IP, UDP, TCP, ICMP, and others)
- Application, source port or range (For example: FTP, HTTP, Telnet, 31000-33000, etc.)
- Application, destination port or range (For example: FTP, HTTP, Telnet, 31000-33000, etc.)
- Inverse: activate the filter whenever the specified conditions are *not* met.

For example, you can create a single filter that blocks external Telnet traffic to your main server, except from a trusted IP address. Another filter could warn you if FTP access is attempted from a specific IP address. Another filter could redirect all incoming e-mail traffic to a master post-office where it can be analyzed for spam. The options are nearly endless.

Below are a list of the well-known protocols and applications.

Table 16-1 Well-Known Protocol Types

Number	Protocol Name
1	icmp
2	igmp
6	tcp
17	udp
89	ospf

Table 16-2 Well-Known Application Ports

Number	TCP/UDP Application	Number	TCP/UDP Application	Number	TCP/UDP Application
20	ftp-data	70	gopher	144	news
21	ftp	79	finger	161	snmp
22	ssh	80	http	162	snmptrap
23	telnet	109	pop2	179	bgp
25	smtp	110	pop3	194	irc
37	time	111	sunrpc	220	imap3
42	name	119	nntp	389	ldap
43	whois	123	ntp	443	https
53	domain	143	imap	520	rip
69	tftp				

Stacking Filters

Once configured, filters are assigned and enabled on a per port basis. Each filter can be used by itself or in combination with any other filter on any given switch port. The filters are numbered 1 through 224. When multiple filters are stacked together on a port, the filter's number determines its order of precedence: the filter with the lowest number is checked first. When traffic is encountered at the switch port, if the filter matches, its configured action takes place and the rest of the filters are ignored. If the filter criteria doesn't match, the next filter is tried.

As long as the filters do not overlap, you can improve filter performance by making sure that the most heavily utilized filters are applied first. For example, consider a filter system where the Internet is divided according to destination IP address:

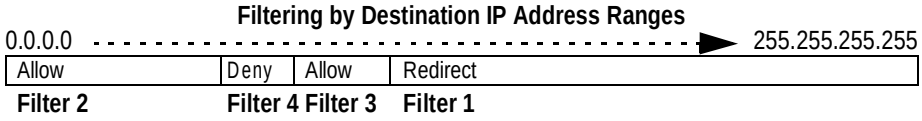

Figure 16-1 Assigning Filters according to Range of Coverage

Assuming that traffic is distributed evenly across the Internet, the largest area would be the most utilized and is assigned to filter 1. The smallest area is assigned to filter 4.

Overlapping Filters

Filters are permitted to overlap, although special care should be taken to ensure the proper order of precedence. When overlapping filters are present, the more specific filters (those that target fewer addresses or ports) should be applied before the generalized filters. For example:

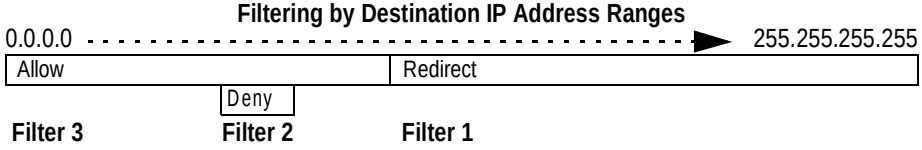

Figure 16-2 Assigning Filters to Overlapping Ranges

In this example, the "deny" filter must be processed prior to the "allow" filter. If the "allow" filter was allowed to take precedence, the "deny" filter could never be triggered.

The Default Filter

Before filtering can be enabled on any given port, a default filter should be configured. This filter handles any traffic not covered by any other filter. All the criteria in the default filter must be set to the full range possible ("`any`"). For example:

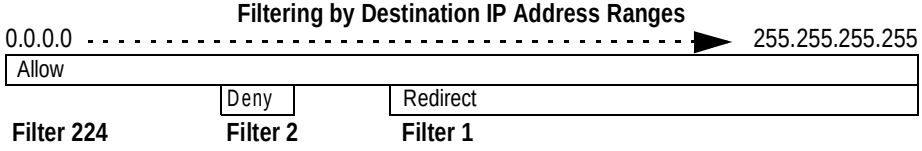

Figure 16-3 Assigning a Default Filter

In this example, filter 224 is the default filter. If no other filter acts on the traffic, filter 224 handles it. All criteria in filter 224 is set to the "**any**" state.

Although recommended when configuring filters for IP traffic control and redirection, default filters are not required. Using default filters can increase session performance, but takes some of the session binding resources. If you experience an unacceptable number of binding failures as shown in the Server Load Balancing Maintenance Statistics (`/stats/slb/maint`), you may wish to remove some of the default filters.

Numbering Filters

You may wish to consider numbering your filters by increments of 5 or 10 (for example: 5, 10, 15, 20, etc.). This allows for filters to be easily inserted between others in the list, if required.

Security Example

Consider the following sample network:

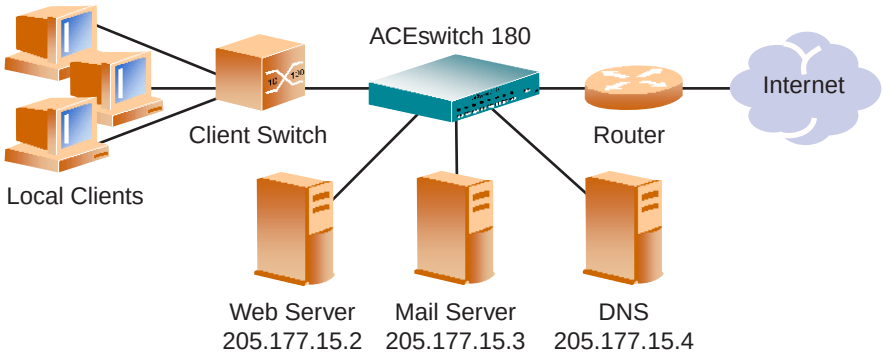

Figure 16-4 Example Security Topology

In this example, the network is made of local clients on a collector switch, a web server, a mail server, a domain name server, and a connection to the Internet. All the local devices are on the same subnet.

For best security, deny everything except for those services you definitely want to allow. In this example, the administrator wishes to install basic security filters to allow only the following traffic:

- External HTTP access to the local web server
- External POP3 (mail) access to the local mail server
- Local clients browsing the World Wide Web
- Local clients using Telnet to access sites outside the intranet
- Domain Name Service

All other traffic will be denied and logged.

NOTE – Since IP address and port information can be manipulated by external sources, filtering does not replace the necessity for a well-constructed network firewall.

Example Configuration for the Security Solution

Prior to configuration, you must be connected to the switch command-line interface as the administrator (see Chapter 2, "The Command-Line Interface").

NOTE – For details about any of the menu commands described in this example, see "The Filter Menu" on page 7-43.

In this example, *all filters will be applied only to the switch port which connects to the Internet.* If intranet restrictions were required, filters could be placed on switch ports connecting to local devices.

Also, filtering is not limited to the few protocols and TCP or UDP applications shown in this example. See the tables on page 16-2 for a list of other well-known protocols and services.

1. **Assign an IP address to each of the network devices.**

   For this example, the network devices have the following IP addresses on the same IP subnet:

   Table 16-3 Web-Cache Example: Real Server IP addresses

Network Device	IP address
Local Subnet	205.177.15.0 - 205.177.15.255
Web Server	205.177.15.2
Mail Server	205.177.15.3
Domain Name Server	205.177.15.4

2. **On the switch, create a default filter that will deny and log unwanted traffic.**

   The default filter is defined as filter 224 in order to give it the lowest order of precedence:

```
>> # /cfg/slb/filt 224                (Select the default filter)
>> Filter 224# sip any                 (From any source IP addresses)
>> Filter 224# dip any                 (To any destination IP addresses)
>> Filter 224# proto any               (For any protocols)
>> Filter 224# actio deny              (Deny matching traffic)
>> Filter 224# log enable              (Log matching traffic to syslog)
>> Filter 224# ena                     (Enable the default filter)
```

NOTE – When the `proto` parameter is *not* `tcp` or `udp`, then `sport` and `dport` are ignored.

3. **On the switch, create a filter that will allow external HTTP requests to reach the web server.**

   The filter must recognize and allow TCP traffic with the web-server's destination IP address and HTTP destination port:

```
>> Filter 224# ../filt 1                    (Select the menu for Filter #1)
>> Filter 1# sip any                        (From any source IP address)
>> Filter 1# dip 205.177.15.2               (To web-server dest. IP address)
>> Filter 1# dmask 255.255.255.255          (Fill mask for exact dest. address)
>> Filter 1# proto tcp                      (For TCP protocol traffic)
>> Filter 1# sport any                      (From any source port)
>> Filter 1# dport http                     (To an HTTP destination port)
>> Filter 1# actio allow                    (Allow matching traffic to pass)
>> Filter 1# ena                            (Enable the filter)
```

4. **On the switch, create a pair of filters to allow incoming and outgoing mail to and from the mail server.**

   Filter 2 allows incoming mail to reach the mail server, and filter 3 allows outgoing mail to reach the Internet:

```
>> Filter 1# ../filt 2                      (Select the menu for Filter #2)
>> Filter 2# sip any                        (From any source IP address)
>> Filter 2# dip 205.177.15.3               (To mail-server dest. IP address)
>> Filter 2# dmask 255.255.255.255          (Fill mask for exact dest. address)
>> Filter 2# proto tcp                      (For TCP protocol traffic)
>> Filter 2# sport any                      (From any source port)
>> Filter 2# dport pop3                     (To a POP3 destination port)
>> Filter 2# actio allow                    (Allow matching traffic to pass)
>> Filter 2# ena                            (Enable the filter)
>> Filter 2# ../filt 3                      (Select the menu for Filter #3)
>> Filter 3# sip 205.177.15.3               (From mail-server source IP address)
>> Filter 3# smask 255.255.255.255          (Fill mask for exact source address)
>> Filter 3# dip any                        (To any destination IP address)
>> Filter 3# proto tcp                      (For TCP protocol traffic)
>> Filter 3# sport pop3                     (From a POP3 port)
>> Filter 3# dport any                      (To any destination port)
>> Filter 3# actio allow                    (Allow matching traffic to pass)
>> Filter 3# ena                            (Enable the filter)
```

5. **On the switch, create a filter that will allow local clients to browse the web.**

   The filter must recognize and allow TCP traffic to reach the local client destination IP addresses if originating from any HTTP source port:

```
>> Filter 3# ../filt 4                    (Select the menu for Filter #4)
>> Filter 4# sip any                      (From any source IP address)
>> Filter 4# dip 205.177.15.0             (To base local network dest. address)
>> Filter 4# dmask 255.255.255.0          (For entire subnet range)
>> Filter 4# proto tcp                    (For TCP protocol traffic)
>> Filter 4# sport http                   (From any source HTTP port)
>> Filter 4# dport any                    (To any destination port)
>> Filter 4# actio allow                  (Allow matching traffic to pass)
>> Filter 4# ena                          (Enable the filter)
```

6. **On the switch, create a filter that will allow local clients to Telnet anywhere outside the local intranet.**

   The filter must recognize and allow TCP traffic to reach the local client destination IP addresses if originating from a Telnet source port:

```
>> Filter 4# ../filt 5                    (Select the menu for Filter #5)
>> Filter 5# sip any                      (From any source IP address)
>> Filter 5# dip 205.177.15.0             (To base local network dest. address)
>> Filter 5# dmask 255.255.255.0          (For entire subnet range)
>> Filter 5# proto tcp                    (For TCP protocol traffic)
>> Filter 5# sport telnet                 (From a Telnet port)
>> Filter 5# dport any                    (To any destination port)
>> Filter 5# actio allow                  (Allow matching traffic to pass)
>> Filter 5# ena                          (Enable the filter)
```

7. **On the switch, create a series of filters to allow Domain Name System (DNS) traffic.**

   DNS traffic requires four filters. One pair is needed for UDP traffic: incoming and outgoing. Another pair is needed for TCP traffic: incoming and outgoing.

For UDP:

```
>> Filter 5# ../filt 6                    (Select the menu for Filter #6)
>> Filter 6# sip any                      (From any source IP address)
>> Filter 6# dip 205.177.15.4             (To local DNS Server)
>> Filter 6# dmask 255.255.255.255        (Fill mask for exact dest. address)
>> Filter 6# proto udp                    (For UDP protocol traffic)
>> Filter 6# sport any                    (From any source port)
>> Filter 6# dport domain                 (To any DNS destination port)
>> Filter 6# actio allow                  (Allow matching traffic to pass)
>> Filter 6# ena                          (Enable the filter)
>> Filter 6# ../filt 7                    (Select the menu for Filter #7)
>> Filter 7# sip 205.177.15.4             (From local DNS Server)
>> Filter 7# smask 255.255.255.255        (Fill mask for exact source address)
>> Filter 7# dip any                      (To any destination IP address)
>> Filter 7# proto udp                    (For UDP protocol traffic)
>> Filter 7# sport domain                 (From a DNS source port)
>> Filter 7# dport any                    (To any destination port)
>> Filter 7# actio allow                  (Allow matching traffic to pass)
>> Filter 7# ena                          (Enable the filter)
```

Similarly, for TCP:

```
>> Filter 7# ../filt 8                    (Select the menu for Filter #8)
>> Filter 8# sip any                      (From any source IP address)
>> Filter 8# dip 205.177.15.4             (To local DNS Server)
>> Filter 8# dmask 255.255.255.255        (Fill mask for exact dest. address)
>> Filter 8# proto tcp                    (For TCP protocol traffic)
>> Filter 8# sport any                    (From any source port)
>> Filter 8# dport domain                 (To any DNS destination port)
>> Filter 8# actio allow                  (Allow matching traffic to pass)
>> Filter 8# ena                          (Enable the filter)
>> Filter 8# ../filt 9                    (Select the menu for Filter #9)
>> Filter 9# sip 205.177.15.4             (From local DNS Server)
>> Filter 9# smask 255.255.255.255        (Fill mask for exact source address)
>> Filter 9# dip any                      (To any destination IP address)
>> Filter 9# proto tcp                    (For TCP protocol traffic)
>> Filter 9# sport domain                 (From a DNS source port)
>> Filter 9# dport any                    (To any destination port)
>> Filter 9# actio allow                  (Allow matching traffic to pass)
>> Filter 9# ena                          (Enable the filter)
```

8. **On the switch, assign the filters to the switch port that connects to the Internet:**

```
>> Filter 9# ../port 5                    (Select the SLB port 5 to the Internet)
>> SLB Port 5 # add 1                     (Add filter 1 to port 5)
>> SLB Port 5 # add 2                     (Add filter 2 to port 5)
>> SLB Port 5 # add 3                     (Add filter 3 to port 5)
>> SLB Port 5 # add 4                     (Add filter 4 to port 5)
>> SLB Port 5 # add 5                     (Add filter 5 to port 5)
>> SLB Port 5 # add 6                     (Add filter 6 to port 5)
>> SLB Port 5 # add 7                     (Add filter 7 to port 5)
>> SLB Port 5 # add 8                     (Add filter 8 to port 5)
>> SLB Port 5 # add 9                     (Add filter 9 to port 5)
>> SLB Port 5 # add 224                   (Add the default filter to port 5)
>> SLB Port 5 # filt enable               (Enable filtering for port 5)
```

9. **On the switch, apply and verify the configuration.**

```
>> SLB Port 5 # ..                        (Select Server Load Balancing Menu)
>> Server Load Balancing# apply           (Make your changes active)
>> Server Load Balancing# cur             (View current settings)
```

Examine the resulting information. If any settings are incorrect, make appropriate changes.

10. **On the switch, save your new configuration changes.**

```
>> Server Load Balancing# save            (Save for restore after reboot)
```

11. **On the switch, check the Server Load Balancing information.**

```
>> Server Load Balancing# /info/slb       (View SLB information)
```

Check that all Server Load Balancing parameters are working according to expectation. If necessary, make any appropriate configuration changes and then check the information again.

NOTE – Changes to filters on a given port do not take effect until the port's session information is updated (every two minutes or so). To make filter changes take effect immediately, clear the session binding table for the port (see the `clear` command under Table 8-3 on page 8-5).

Web-Cache Redirection Example

For many companies, the Internet is an indispensable source for business and technical information. Much of the information brought into your company from the Internet, however, is not unique. Often, clients will access the same information many times as they return to a web-page for additional information or to explore other links.

Duplicate information may be requested more inadvertently as the myriad components that make up Internet data (pictures, buttons, frame, text, and so on) are reloaded from page to page. Add multiple clients to the picture, and the amount of repeated data that comes in through your Internet router can account for a great deal of its congestion. Redundant requests also decrease the amount of your available bandwidth to the Internet.

Web-cache redirection can help alleviate the congestion seen at your Internet router. When Application Redirection filters are properly configured for your ACElerate powered switch, outbound client requests for Internet data are intercepted and redirected to a group of web-cache servers on your network. The web-cache servers duplicate and store inbound Internet data that has been requested by your clients. If the web-cache servers recognize a client's outbound request as one that can be filled with cached information, the web-cache servers will supply the information, rather than sending the request out across the Internet.

In addition to increasing the efficiency of your network, access to locally cached information can be granted much faster than by pulling the same information across the Internet.

Web-Cache Redirection Environment

Consider a network where client HTTP requests begin to regularly overload the Internet router.

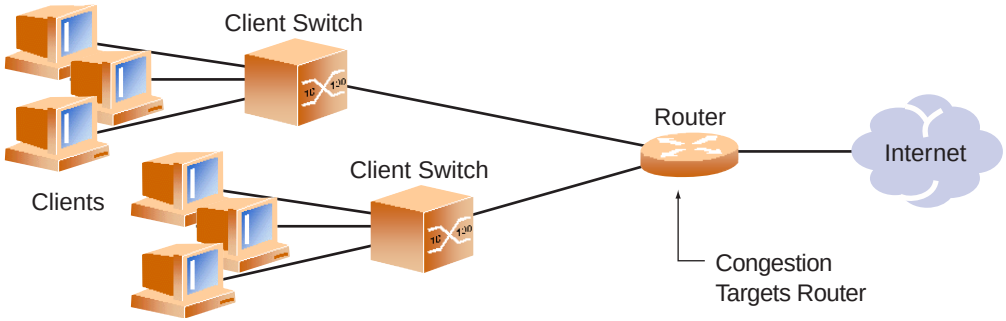

Figure 16-5 Traditional network without Web Cache Redirection

The network needs a solution that addresses the following key concerns:

- The solution must be readily scalable
- The administrator should not have to reconfigure all the clients' browsers to use Proxy Servers.

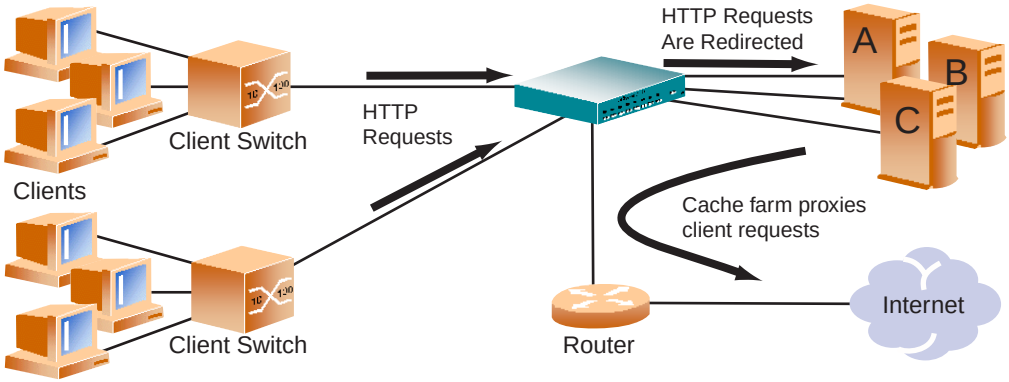

Figure 16-6 Network with Web Cache Redirection

Adding an ACEswitch with optional Layer 4 software addresses these issues:

- Web-cache servers can be added or removed dynamically without interrupting services.
- Performance is improved by balancing the cached web request load across multiple servers. More servers can be added at any time to increase processing power.
- The proxy is transparent to the client.
- Frames that are not associated with HTTP requests are passed normally to the router.

Example Configuration for the Web-Cache Solution

The following is required prior to configuration:

- You must be connected to the switch command-line interface as the administrator (see Chapter 2, "The Command-Line Interface").
- Optional Layer 4 software must be enabled (see "Activating Optional Software" on page 8-6).

NOTE – For details about any of the menu commands described in this example, see "Configuring Server Load Balancing" on page 7-29."

In this example, an ACEswitch 180 is placed between the clients and the border gateway to the Internet. The switch will be configured to intercept all Internet bound HTTP requests (on default TCP port 80), and redirect them to the web-cache servers. The switch will distribute HTTP requests equally to the web-cache servers based on the destination IP address of the requests.

Also, filters are not limited to the few protocols and TCP or UDP applications shown in this example. See the tables on page 16-2 for a list of other well-known protocols and services.

1. **Assign an IP address to each of the web-cache servers.**

   Just as with Server Load Balancing, the web-cache real servers will be assigned an IP address and placed into a real server group. The real servers must be in the same VLAN and must have an IP route to the switch that will perform the web-cache redirection. In addition, the path from the switch to the real servers must not contain a router. The router would stop HTTP requests from reaching the web-cache servers, instead directing them back out to the Internet.

   More complex network topologies can be used if configuring IP proxy addresses (see "IP Proxy Addresses for Transparent Proxies or Complex Networks" on page 16-18).

   For this example, the three web-cache real servers have the following IP addresses on the same IP subnet:

   Table 16-4 Web-Cache Example: Real Server IP addresses

Web Cache Server	IP address
Server A	200.200.200.2
Server B	200.200.200.3
Server C	200.200.200.4

2. **Install web-cache software on all three web-cache servers.**

3. **Full Network Address Translation (NAT) is required.**

   Install transparent proxy software with NAT on all three web-cache servers, or define proxy IP addresses on the switch (see "IP Proxy Addresses for Transparent Proxies or Complex Networks" on page 16-18).

4. **Define an IP interface on the switch.**

   The switch must have an IP route to all of the real servers which receive redirection services. The switch uses this path to determine the level of TCP/IP reachability of the real servers.

   To configure an IP interface for this example, enter this command from the CLI:

```
>> Main# /cfg/ip/if 1                          (Select IP interface #1)
>> IP Interface 1# addr 200.200.200.100       (Assign IP address for the interface)
>> IP Interface 1# ena                        (Enable IP interface #1)
```

NOTE – The IP interface and the real servers must belong to the same VLAN. This example assumes that all ports and IP interfaces use default VLAN #1, requiring no special VLAN configuration for the ports or IP interface.

5. **On the switch, define each Real Server**

   For each web-cache real server, you must assign a real server number, specify its actual IP address, and enable the real server. For example:

```
>> ip# /cfg/slb/real 1                         (Server A is real server 1)
>> Real server 1 # rip 200.200.200.2           (Assign Server A IP address)
>> Real server 1 # ena                         (Enable real server 1)
>> Real server 1 # ../real 2                   (Server B is real server 2)
>> Real server 2 # rip 200.200.200.3           (Assign Server B IP address)
>> Real server 2 # ena                         (Enable real server 2)
>> Real server 2 # ../real 3                   (Server C is real server 3)
>> Real server 3 # rip 200.200.200.4           (Assign Server C IP address)
>> Real server 3 # ena                         (Enable real server 3)
```

6. **On the switch, define a Real Server Group.**

   This places the three web-cache real servers into one service group:

```
>> Real server 3 # /cfg/slb/group 1            (Select real server group 1)
>> Real server group 1 # add 1                 (Add real server 1 to group 1)
>> Real server group 1 # add 2                 (Add real server 2 to group 1)
>> Real server group 1 # add 3                 (Add real server 3 to group 1)
```

7. On the switch, set the Real Server Group metric to `minmisses`.

This helps minimize web-cache misses in the event real servers fail or are taken out of service:

```
>> Real server group 1 # metrc minmisses   (Metric for minimum cache misses.)
```

8. On the switch, define the switch port states.

In this example, the following ports are being used on the ACEswitch 180:

Table 16-5 Web Host Example: ACEswitch 180 Port Usage

Port	Host	Port Setting
1	Server A	`none`
2	Server B	`none`
3	Server C	`none`
4	Internet Router	`none`
5	Client switch A. This connects the switch to a group of clients where client Internet requests originate.	`redir`
6	Client switch B. This connects the switch to a group of clients where client Internet requests originate.	`redir`

NOTE – The switch ports to web-server hosts use the "`none`" setting. Do not use the "`server`" setting with Application Redirection. The "`server`" setting is used only with Server Load Balancing.

The ports are configured as follows:

```
>> Real server group 1 # /cfg/slb/port 1   (Select physical switch port 1)
>> SLB port 1 # state none                 (Set port 1 for no Layer 4 traffic)
>> SLB port 1 # ../port 2                  (Select physical switch port 2)
>> SLB port 2 # state none                 (Set port 2 for no Layer 4 traffic)
>> SLB port 2 # ../port 3                  (Select physical switch port 3)
>> SLB port 3 # state none                 (Set port 3 for no Layer 4 traffic)
>> SLB port 3 # ../port 4                  (Select physical switch port 4)
>> SLB port 4 # state none                 (Set for no Layer 4 traffic for NFS)
>> SLB port 4 # ../port 5                  (Select physical switch port 5)
>> SLB port 5 # state redir                (Set port 5 for cached traffic)
>> SLB port 5 # ../port 6                  (Select physical switch port 6)
>> SLB port 6 # state redir                (Set port 6 for cached traffic)
```

9. **On the switch, create a filter that will intercept and redirect all client HTTP requests.**

   The filter must be able to intercept all TCP traffic for the HTTP destination port, and must redirect it to the proper port on the real server group:

```
>> SLB port 6 # /cfg/slb/filt 2          (Select the menu for Filter #2)
>> Filter 2# sip any                     (From any source IP addresses)
>> Filter 2# dip any                     (To any destination IP addresses)
>> Filter 2# proto tcp                   (For TCP protocol traffic)
>> Filter 2# sport any                   (From any source port)
>> Filter 2# dport http                  (To an HTTP destination port)
>> Filter 2# actio redir                 (Set the action for redirection)
>> Filter 2# rport http                  (Set the redirection port)
>> Filter 2# group 1                     (Select real server group 1)
>> Filter 2# ena                         (Enable the filter)
```

   The `rport` parameter must be configured whenever TCP protocol traffic is redirected. The `rport` parameter defines the real server TCP or UDP port to which redirected traffic will be sent. The port defined by the `rport` parameter is used when performing Layer 4 health checks of TCP services.

   Also, if transparent proxies are used for Network Address Translation (NAT) on the switch (see Step 3. on page 16-14), the `rport` parameter must be configured for all Application Redirection filters. Take care to use the proper port designation with `rport`: if the transparent proxy operation resides on the host, the well-known port (80, or "http") is probably required. If the transparent proxy occurs on the switch, make sure to use the service port required by the specific software package.

   See "IP Proxy Addresses for Transparent Proxies or Complex Networks" on page 16-18 for more about IP proxy addresses.

10. **On the switch, create a default filter.**

    In this case, the default filter will allow all non-cached traffic to proceed normally:

```
>> Filter 2# ../filt 224                 (Select the default filter)
>> Filter 224# sip any                   (From any source IP addresses)
>> Filter 224# dip any                   (To any destination IP addresses)
>> Filter 224# proto any                 (For any protocols)
>> Filter 224# actio allow               (Set the action to allow traffic)
>> Filter 224# ena                       (Enable the default filter)
```

NOTE – When the `proto` parameter is not tcp or udp, then `sport` and `dport` are ignored.

11. On the switch, assign the filters to the client ports.

Recalling Table 16-5 on page 16-15, redirection clients are connected to physical switch ports 5 and 6. Both ports are configured with our filters as follows:

```
>> Filter 224# ../port 5                    (Select the SLB port 5)
>> SLB Port 5 # add 2                       (Add filter 1 to port 5)
>> SLB Port 5 # add 224                     (Add the default filter to port 5)
>> SLB Port 5 # filt enable                 (Enable filtering for port 5)
>> SLB Port 5 # ../port 6                   (Select the SLB port 6)
>> SLB Port 6 # add 2                       (Add filter 1 to port 6)
>> SLB Port 6 # add 224                     (Add the default filter to port 6)
>> SLB Port 6 # filt enable                 (Enable filtering for port 6)
```

12. On the switch, enable, apply, and verify the configuration.

```
>> SLB Port 6 # ..                          (Select Server Load Balancing Menu)
>> Server Load Balancing# on                (Activate Layer 4 software services)
>> Server Load Balancing# apply             (Make your changes active)
>> Server Load Balancing# cur               (View current settings)
```

NOTE – Server Load Balancing must be turned on in order for Application Redirection to work properly. The "`on`" command is valid only if the optional Layer 4 software is enabled on your switch (see "Activating Optional Software" on page 8-6).

Examine the resulting information from the "`cur`" command. If any settings are incorrect, make appropriate changes.

13. On the switch, save your new configuration changes.

```
>> Server Load Balancing# save              (Save for restore after reboot)
```

14. On the switch, check the Server Load Balancing information.

```
>> Server Load Balancing# /info/slb         (View SLB information)
```

Check that all Server Load Balancing parameters are working according to expectation. If necessary, make any appropriate configuration changes and then check the information again.

NOTE – Changes to filters on a given port do not take effect until the port's session information is updated (every two minutes or so). To make filter changes take effect immediately, clear the session binding table for the port (see the `clear` command under Table 8-3 on page 8-5).

IP Proxy Addresses for Transparent Proxies or Complex Networks

Transparent proxies provide the following benefits when used with Application Redirection:

- With proxies IP addresses configured on redirected ports, the switch can redirect client requests to servers located on any subnet, anywhere.
- For HTTP traffic, the switch can perform transparent substitution for all source and destination addresses, including destination port remapping. This provides support for comprehensive, fully-transparent proxies.

Adding proxies requires only a few steps. Re-examining the example above, the port assignments appeared as follows:

Table 16-6 Web Proxy Example: ACEswitch 180 Port Usage

Port	Host	Port Setting
1	Server A	`none`
2	Server B	`none`
3	Server C	`none`
4	Internet Router	`none`
5	Client switch A. This connects the switch to a group of clients where client Internet requests originate.	`redir`
6	Client switch B. This connects the switch to a group of clients where client Internet requests originate.	`redir`

Only the ports set to the "`redir`" state require proxy IP addresses to be configured. Each proxy IP address must be unique on your network. These are configured as follows:

```
>> # /cfg/slb/port 5                    (Select network port #5)
>> SLB port 5# pip 200.200.200.68       (Set proxy IP address for port #5)
>> SLB port 5# ../port 6                (Select network port #6)
>> SLB port 6# pip 200.200.200.69       (Set proxy IP address for port #6)
```

Once proxy IP addresses are established, you need to configure each Application Redirection filter (filter 2 in our example) with the real server TCP or UDP port to which redirected traffic will be sent. In this case, we are mapping the requests to different destination port (8080):

```
>> # /cfg/slb/filt 2                    (Select the menu for Filter #2)
>> Filter 2 # rport 8080                (Set proxy redirection port)
```

The Layer 4 proxies are transparent to the user. No additional client configuration is needed.

Excluding Non-Cacheable Sites

Some web sites provide content which isn't well suited for redirection to cache servers. Such sites might provide browser-based games, applications that keep real-time session information or authenticate by client IP address.

To prevent such sites from being redirected to cache-servers, create a filter which allows this specific traffic to pass normally through the switch. This filter must have a higher precedence (a lower filter number) than the Application Redirection filter.

For example, if you wished to prevent a popular web-based game site on subnet 200.10.10.* from being redirected, you could add the following to the previous example configuration:

```
>> # /cfg/slb/filt 1                    (Select the menu for Filter #1)
>> Filter 1# dip 200.10.10.0             (To the site's destination IP address)
>> Filter 1# dmask 255.255.255.0         (For entire subnet range)
>> Filter 1# sip any                     (From any source IP address)
>> Filter 1# proto tcp                   (For TCP traffic)
>> Filter 1# dport http                  (To an HTTP destination port)
>> Filter 1# sport any                   (From any source port)
>> Filter 1# actio allow                 (Allow matching traffic to pass)
>> Filter 1# ena                         (Enable the filter)
>> Filter 1# ../port 5                   (Select SLB port 5)
>> SLB port 5# add 1                     (Add the filter to port 5)
>> SLB port 5# ../port 6                 (Select SLB port 6)
>> SLB port 6# add 1                     (Add the filter to port 6)
>> SLB port 6# apply                     (Apply configuration changes)
>> SLB port 6# save                      (Save configuration changes)
```

Additional Application Redirection Options

Application Redirection can be used in combination with other Layer 4 options such as load balancing metrics, health checks, real server group backups, and more. See "Additional Server Load Balancing Options" on page 15-17 for details.

Network Address Translation Examples

In the following NAT examples, a company has configured its internal network with "private" IP addresses. A private network is one that is isolated from the global Internet, and is therefor free from the usual restrictions requiring the use of registered, globally unique IP addresses. Private networks can use whatever IP addresses they please, including those that are in use elsewhere on the Internet, or reserved for other purposes.

Private networks serve two main purposes. First, because private IP addresses are not valid or visible outside the private network, they can increase network security. Second, since valid, registered IP addresses are a limited resource, many companies use private IP addresses to create internal networks much larger than they could using only their official addresses.

With Network Address Translation (NAT), private networks are not required to remain isolated. NAT capabilities within the switch allow internal, private network IP addresses to be translated to valid, publicly advertised IP addresses and back again.

Internal Client Access to Internet

In this dynamic NAT example, clients on the internal private network require TCP/UDP access to the Internet:

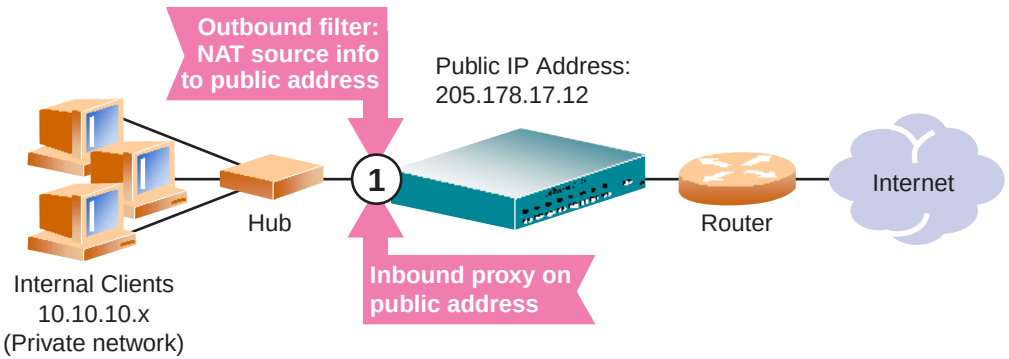

Figure 16-7 Dynamic NAT

This example requires a Network Address Translation (NAT) filter to be configured on the switch port connected to the internal clients. When the NAT filter is triggered by outbound client traffic, the internal private IP address information on the outbound packets is translated to a valid, publicly advertised IP address. In addition, the public IP address must be configured as a proxy IP address on the switch port connected to the internal clients. The proxy performs the reverse translation, restoring the private network addresses on inbound packets.

This is a "many to one" solution: multiple clients on the private subnet take advantage of a single external IP address, thus conserving valid IP addresses.

This example could be configured as follows:

```
>> # /cfg/slb/filt 14                    (Select the menu for client filter)
>> Filter 1# invert ena                   (Invert the filter logic)
>> Filter 1# dip 10.10.10.0               (If the destination is not private)
>> Filter 1# dmask 255.255.255.0          (For the entire private subnet range)
>> Filter 1# sip any                      (From any source IP address)
>> Filter 1# actio nat                    (Perform NAT on matching traffic)
>> Filter 1# nat source                   (Translate source information)
>> Filter 1# proxy enable                 (Allow pip proxy translation)
>> Filter 1# ena                          (Enable the filter)
>> Filter 1# ../port 1                    (Select SLB port 1)
>> SLB port 1# add 14                     (Add the filter to port 1)
>> SLB port 1# pip 205.178.17.12          (Set public IP address proxy)
>> SLB port 1# filt enable                (Enable filtering on port 1)
>> SLB port 1# apply                      (Apply configuration changes)
>> SLB port 1# save                       (Save configuration changes)
```

NOTE – Dynamic NAT solutions apply only to TCP/UDP traffic. Also, filters for dynamic NAT should be placed behind static NAT filters (next example). Dynamic filters should be given higher filter numbers.

External Client Access to Server

In this example, clients on the external Internet require access to a server on the private network:

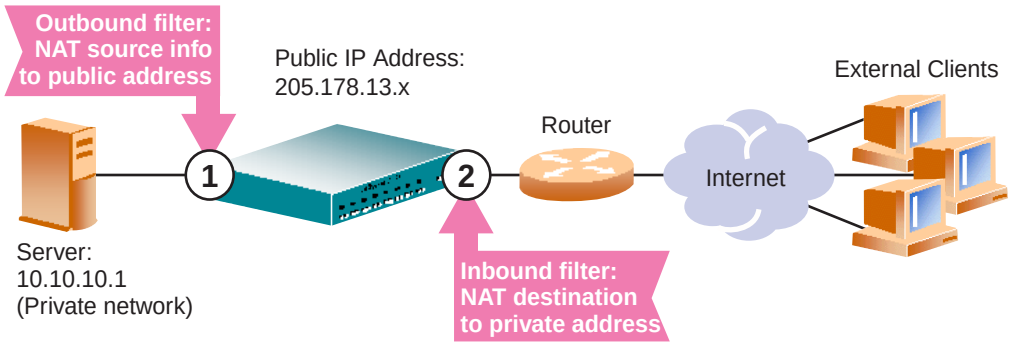

Figure 16-8 Static NAT

This static NAT (non-proxy) example requires two filters: one for the external client-side switch port, and one for the internal, server-side switch port. The client-side filter translates incoming requests for the publicly advertised server IP address to the server's internal private network address. The filter for the server-side switch port reverses the process, translating the server's private address information to a valid public address.

This could be configured as follows:

```
>> # /cfg/slb/filt 10                    (Select the menu for outbound filter)
>> Filter 10# actio nat                  (Perform NAT on matching traffic)
>> Filter 10# nat source                 (Translate source information)
>> Filter 10# sip 10.10.10.0             (From the clients private IP address)
>> Filter 10# smask 255.255.255.0        (For the entire private subnet range)
>> Filter 10# dip 205.178.13.0           (To the public network address)
>> Filter 10# dmask 255.255.255.0        (For the same subnet range)
>> Filter 10# proxy disable              (Override any pip proxy settings)
>> Filter 10# ena                        (Enable the filter)
>> Filter 10# ../filt 11                 (Select the menu for inbound filter)
>> Filter 11# actio nat                  (Use the same settings as outbound)
>> Filter 11# nat dest                   (Reverse the translation direction)
>> Filter 11# sip 10.10.10.0             (Use the same settings as outbound)
>> Filter 11# smask 255.255.255.0        (Use the same settings as outbound)
>> Filter 11# dip 205.178.13.0           (Use the same settings as outbound)
>> Filter 11# dmask 255.255.255.0        (Use the same settings as outbound)
>> Filter 11# proxy disable              (Override any pip proxy settings)
>> Filter 11# ena                        (Enable the filter)
>> Filter 11# ../port 1                  (Select server-side port)
>> SLB port 1# add 10                    (Add the outbound filter)
>> SLB port 1# filt enable               (Enable filtering on port 1)
>> SLB port 1# ../port 2                 (Select the client-side port)
>> SLB port 2# add 11                    (Add the inbound filter)
>> SLB port 2# filt enable               (Enable filtering on port 2)
>> SLB port 2# apply                     (Apply configuration changes)
>> SLB port 2# save                      (Save configuration changes)
```

Note the following imporant poince about this configuration:

- Within each filter, the `smask` and `dmask` values are identical.
- All parameters for both filters are identical except for the NAT direction. For filter #10, `nat source` is used. For filter #11, `nat dest` is used.
- Filters for static (non-proxy) NAT should be placed ahead of dynamic NAT filters (previous example). Static filters should be given lower filter numbers.

CHAPTER 17

Distributed Server Load Balancing

This chapter describes how to configure and use Distributed Server Load Balancing.

NOTE – Both the optional Server Load Balancing and Distributed Server Load Balancing software keys must be enabled (see "Activating Optional Software" on page 8-6).

DSLB Overview

Distributed Server Load Balancing (DSLB), lets you balance server traffic load across multiple physical sites. This allows you to smoothly integrate the resources of a world-wide series of server sites, and to balance web content (or other services) intelligently among them. Alteon Network's DSLB system takes into account individual sites' health, response time, and geographic location for a global performance perspective.

Benefits

DSLB meets the following demands for distributed network services:

- High content availability through distributed content and distributed decision making. If one site becomes disabled, the others become aware of it and take up the load.
- No latency during client connection set up. Instant site hand-off decisions can be made by any distributed switch.
- The best performing sites get a majority of traffic over a given period of time, but are not overwhelmed.
- Switches at different sites regularly exchange information through DSSP (Distributed Site State Protocol), and can trigger exchanges when any site's health status changes. This ensures that each active site has valid state knowledge and statistics.
- Takes geography into account, as well as network topology.

- Gives creative control to the administrator or web-master to build and control content by user, location, target application, and more.
- Easy to deploy, manage, and scale. Switch configuration is straight-forward. There are no complex system topologies involving routers, protocols, etcetera.
- Provides flexible design options.
- Supports all IP protocols.

How DSLB Works

Consider the following sample network:

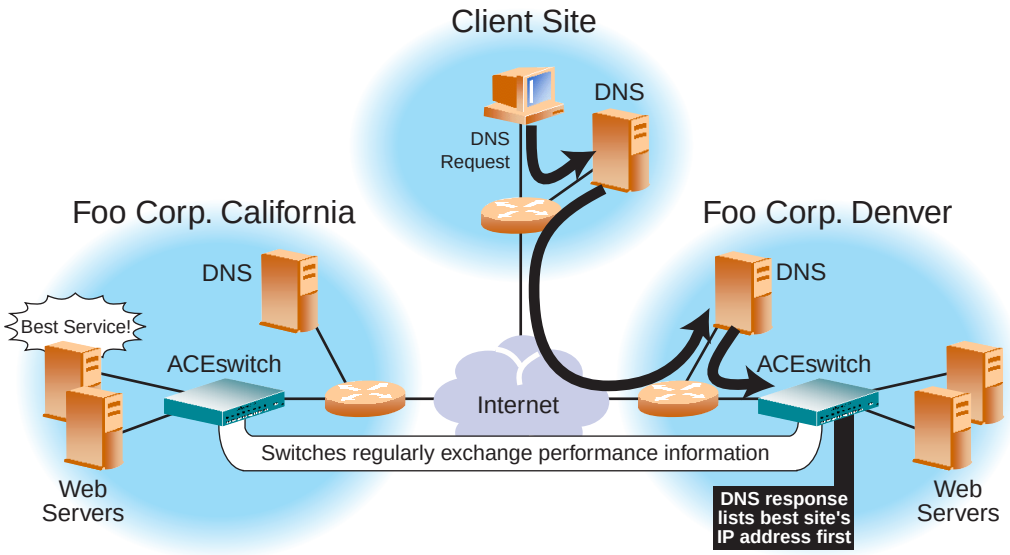

1. Browser requests www.foocorp.com IP address from local DNS.

2. Client's DNS asks its upstream DNS, which in turn asks the next, and so on, until the address is resolved.

3. The Foo Corp. Denver DNS knows that the local ACEswitch is an authoritative name server for www.foocorp.com.

4. The switch DSLB software knows that Foo Corp. California currently provides better service, and responds with Foo Corp. California's virtual IP address listed first.

5. The client connects to Foo Corp. California for the best service.

Figure 17-1 DNS Resolution with Distributed Server Load Balancing

In this example, a client is using their web-browser to view the web-site for the Foo Corporation at www.foocorp.com. The Foo Corporation has two sites: one in California, and one in Denver, each with identical content and services available. Both sites have an Alteon Networks switch configured for Distributed Server Load Balancing. These switches are also configured as the Authoritative Name Servers for www.foocorp.com.

When a client loads their web-browsing software and enters the URL for a web-site such as www.foocorp.com, a query is sent the client's local DNS server, asking for the IP address that represents the domain name entered. If the local DNS server does not have this information cached, it will in turn ask a DNS server further upstream. Eventually, the request will reach an upstream DNS server that has this information on hand, or it will reach one the Foo Corporation's DNS servers. The Foo Corporation's DNS server has been configured to know that the local ACEswitch with Distributed Server Load Balancing software is the authoritative name server for www.foocorp.com.

Each switch with DSLB software is capable of responding to the client's name resolution request. Since each switch regularly checks and communicates health and performance information with its peers, either switch can determine which site (or sites) are best able to serve the client's web-cruising needs, and can respond with a list of IP addresses for the Foo Corporation's distributed sites, prioritized by performance, geography, and other criteria.

The client's web-browser will use the IP address information to open a connection to the best available site. The IP addresses can represent real servers at any site, or they can represent virtual servers at any site, which are in turn locally load balanced according to regular Server Load Balancing configuration.

If the site serving the client HTTP content suddenly experiences a failure (no healthy real servers) or becomes overloaded with traffic (all real servers reach their maximum connection limit), the switch will issue an HTTP Redirect and transparently cause the client to connect to another peer site.

The end result is that the client gets quick, reliable service with no latency and no special client-side configuration.

DSLB Configuration Example

Summary

Configuring Distributed Server Load Balancing is merely an extension of the Server Load Balancing configuration. The process is summarized as follows:

- Use the administrator login to connect to the switch you are configuring.
- Activate Server Load Balancing and Distributed Server Load Balancing software keys.
- Configure the switch at each site with basic attributes
  - Configure the switch IP interface
  - Configure the default gateways
- Configure the switch at each site to act as Domain Name System (DNS) server for each service hosted on its virtual servers. Also, configure the local DNS server to recognize the switch as the authoritative DNS server for the hosted services.
- Configure the switch at each site as usual for local Server Load Balancing.
  - Define each local real server
  - Group local real servers into real server groups
  - Define the local virtual server with its IP address, services, and real server groups
  - Define the switch port states
  - Enable Server Load Balancing
- Finally, make each switch recognize its remote peers.
  - On each switch, configure a remote real server entry for each remote service.
  - Add the remote real server entry to an appropriate real server group.
  - Enable Distributed Server Load Balancing

Example DSLB Configuration Procedure

Consider the following example network:

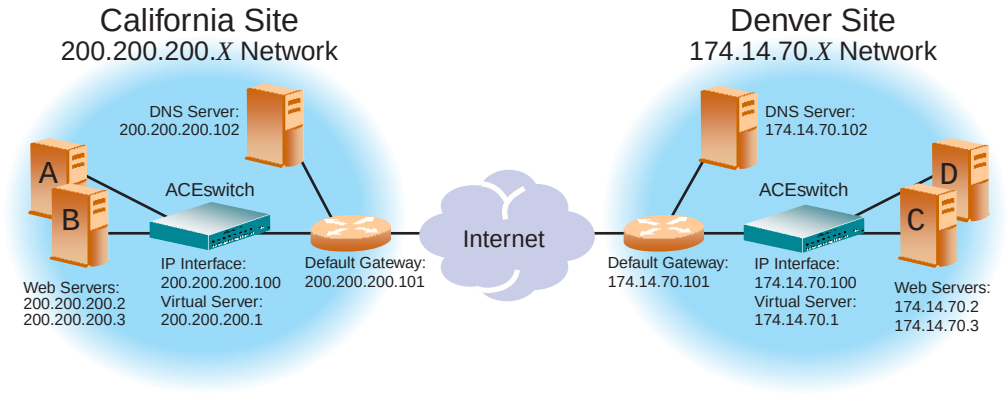

Figure 17-2 Distributed Server Load Balancing Example Topology

In the following examples, many of the options are left to their default values. See "Additional Server Load Balancing Options" on page 15-17 for more options.

The following is required prior to configuration:

- You must be connected to the switch command-line interface as the administrator (see Chapter 2, "The Command-Line Interface").
- Both of the following optional software keys must be activated (see "Activating Optional Software" on page 8-6):
  - Server Load Balancing
  - Distributed Server Load Balancing

NOTE – For details about any of the menu commands described in this example, see Chapter 7, "The Configuration Menu."

Part One: Configure the California Site with Basic System Items

1. **If ACEvision is to be used for managing the California switch, change its service port.**

   Distributed Server Load Balancing uses service port 80 on the IP interface for DSSP updates. By default, ACEvision also uses port 80. Both services cannot use the same port. If ACEvision is enabled (see the `http` command on page 7-5), configure it to use a different port.

   For example, to change the ACEvision port to 8080, enter the following command:

```
>> Main# /cfg/sys                          (Select the System Menu)
>> System# wport 8080                      (Set service port 8080 for ACEvision)
```

2. **On the California switch, define an IP interface.**

   The switch IP interface is the entity that responds when the asked to resolve client DNS requests. The IP interface must have an IP route to the local real servers. The switch uses this path to determine the level of TCP/IP reachability of the real servers.

   To configure an IP interface for this example, enter these commands from the CLI:

```
>> System# /cfg/ip/if 1                    (Select IP interface #1)
>> IP Interface 1# addr 200.200.200.100    (Assign IP address for the interface)
>> IP Interface 1# ena                     (Enable IP interface #1)
```

NOTE – This example assumes that all ports and IP interfaces use default VLAN #1, requiring no special VLAN configuration for the ports or IP interface.

3. **On the California switch, define the default gateway.**

   In this example, a router at the edge of the site acts as the default gateway to the Internet. To configure the default gateway for this example, enter these commands from the CLI:

```
>> IP Interface 1# ../gw 1                 (Select default gateway #1)
>> Default gateway 1# addr 200.200.200.101 (Assign IP address for the gateway)
>> Default gateway 1# ena                  (Enable default gateway #1)
```

4. **Configure the local DNS server to recognize the local DSLB switch as the authoritative name server for the hosted services.**

   Determine the domain name which will be distributed to both sites, and the hostname for each distributed service. In this example, the California DNS server is configured to recognize 200.200.200.100 (the IP interface of the California DSLB switch) as the authoritative name server for www.foocorp.com.

Part Two: Configure the California Switch for Standard SLB

1. **Assign an IP address to each of the real servers in the local California server pool.**

   The real servers in any given real server group must have an IP route to the switch that will perform the Server Load Balancing functions. This is most easily accomplished by placing the switches and servers on the same IP subnet, although advanced routing techniques can be used as long as they do not violate the topology rules outlined in "Network Topology Considerations" on page 15-4.

   For this example, the web-host real servers have the following IP addresses on the same IP subnet:

   Table 17-1 DSLB Example: California Real Server IP Addresses

Real Server	IP address
Server A	200.200.200.2
Server B	200.200.200.3

2. **On the California switch, define each local Real Server.**

   For each local real server, you must assign a real server number, specify its actual IP address, and enable the real server. For example:

```
>> Default gateway 1# /cfg/slb/real 1          (Server A is real server 1)
>> Real server 1 # rip 200.200.200.2           (Assign Server A IP address)
>> Real server 1 # ena                         (Enable real server 1)
>> Real server 1 # ../real 2                   (Server B is real server 2)
>> Real server 2 # rip 200.200.200.3           (Assign Server B IP address)
>> Real server 2 # ena                         (Enable real server 2)
```

3. **On the California switch, define a Real Server Group.**

   This combines the real servers into one service group, and sets the necessary health checking parameters. In this example, HTTP health checking is used to ensure that web content is being served. If the index.html file is not accessible on a real server during health checks, the real server will be marked as down.

   The following commands are entered:

```
>> Real server 2 # /cfg/slb/group 1            (Select real server group 1)
>> Real server group 1# add 1                  (Add real server 1 to group 1)
>> Real server group 1# add 2                  (Add real server 2 to group 1)
>> Real server group 1# healt http             (Use HTTP for health checks)
>> Real server group 1# cntnt index.html       (Set URL content for health checks)
```

4. **On the California switch, define a Virtual Server.**

   All client requests will be addressed to a virtual IP on a virtual server defined on the switch. Clients acquire the virtual IP through normal DNS resolution. HTTP uses well-known TCP port 80. In this example, HTTP is configured as the only service running on this virtual IP, and is associated with our real server group. For example:

```
>> Real server group 1 # /cfg/slb/virt 1      (Select virtual server 1)
>> Virtual server 1# vip 200.200.200.1       (Assign a virtual server IP address)
>> Virtual server 1# add http 1              (Associate virtual port to real group)
>> Virtual server 1# ena                     (Enable the virtual server)
```

NOTE – This configuration is not limited to HTTP web service. Other TCP/IP services can be configured in a similar fashion. For a list of other well-known services and ports, see the command option information on page 7-39.

5. **On the California switch, define the Port States.**

   In this example, the following ports are being used on the ACEswitch 180:

Table 17-2 DSLB Example: California ACEswitch 180 Port Usage

Port	Host	Port Setting
1	Server A	`server`
2	Server B	`server`
6	Default Gateway Router. This connects the switch to the Internet where all client requests originate.	`client`

   The ports are configured as follows:

```
>> Virtual server 1# /cfg/slb/port 1         (Select physical switch port 1)
>> SLB port 1# state server                  (Assign port 1 to server traffic)
>> SLB port 1# ../port 2                     (Select physical switch port 2)
>> SLB port 2# state server                  (Assign port 2 to server traffic)
>> SLB port 2# ../port 6                     (Select physical switch port 6)
>> SLB port 6# state client                  (Set port 6 for client traffic)
```

6. **On the California switch, enable Server Load Balancing.**

```
>> SLB port 6# ..                            (Select the SLB Menu)
>> Server Load Balancing# on                 (Turn Server Load Balancing on)
```

Part Three: Configure the California Site for DSLB

1. **On the California switch, define each remote site.**

   Add and enable the IP address for the IP interface of up to eight remote sites. In this example, there is only one remote site: Denver, with an IP interface address of 174.14.70.100. The following commands are used:

```
>> Server Load Balancing# /cfg/dist/site 1     (Select Remote Site #1)
>> Remote site 1# prima 174.14.70.100         (Define remote IP interface address)
>> Remote site 1# ena                         (Enable remote site #1)
```

   Each additional remote site would be configured in the same manner.

2. **On the California switch, assign each remote distributed service to a local virtual server.**

NOTE – This step can result in improper configuration if not clearly understood. Please take care to note where each configured value originates.

   In this step, we are configuring the local California site to recognize the services offered at the remote Denver site. To do this, configure one real server entry on the California switch for each virtual server located at each remote site. Since there's only one remote site (Denver) with only one virtual server, only one more local real server entry is needed at the California site.

   The new real server entry will be configured with the IP address of the remote virtual server, rather than the usual IP address of a local physical server.

   Also, the "remote" property will be enabled, and the real server entry will be added the real server group under the local virtual server for the intended service. Finally, since the real server health checks will be headed across the Internet, the health checking interval should be increased to 30 or 60 seconds to avoid generating excess traffic. For example:

```
>> Remote site 1# /cfg/slb/real 3          (Create an entry for real server #3)
>> Real server 3# rip 174.14.70.1          (Set remote virtual server IP address)
>> Real server 3# remote enable            (Define the real server as remote)
>> Real server 3# intr 60                  (Set a high health check interval)
>> Real server 3# ena                      (Enable the real server entry)
>> Real server 3# ../group 1               (Select the approp. real server group)
>> Real server group 1# add 3              (Add real server 3 to the group 1)
```

NOTE – The IP address of the real server being added is taken from the virtual server IP address on the remote switch. Do not confuse this value with the IP interface address on the remote switch.

---

3. **On the California switch, define the domain name and hostname for each service hosted on each virtual server.**

   In this example, the domain name for the Foo Corporation is "foocorp.com," and the hostname for the only service (HTTP) is "www." These values are configured as follows:

```
>> Real server group 1# /cfg/slb/virt 1      (Select virtual server #1)
>> Virtual server 1# dname foocorp.com       (Define domain name)
>> Virtual server 1# hname http www          (Define HTTP hostname)
```

   If other services were defined (such as FTP), additional hostname entries would be made.

4. **On the California switch, turn Distributed Server Load Balancing on.**

```
>> Virtual server 1# ../dist                 (Select the DSLB Menu)
>> Distributed SLB menu# on                  (Activate DSLB for the switch)
```

5. **Apply and verify the configuration.**

```
>> Distributed SLB menu# apply               (Make your changes active)
>> Distributed SLB menu# cur                 (View current DSLB settings)
>> Distributed SLB menu# ../cur              (View current SLB settings)
```

   Examine the resulting information. If any settings are incorrect, make and apply any appropriate changes, and then check again.

6. **Save your new configuration changes.**

```
>> Server Load Balancing# save               (Save for restore after reboot)
```

Part Four: Configure the Denver Site with Basic System Items

Following the same procedures as above, configuration the Denver site as follows.

1. **If ACEvision is to be used for managing the Denver switch, change its service port.**

```
>> Main# /cfg/sys                           (Select the System Menu)
>> System# wport 8080                       (Set service port 8080 for ACEvision)
```

2. **On the Denver switch, define an IP interface.**

```
>> Main# /cfg/ip/if 1                       (Select IP interface #1)
>> IP Interface 1# addr 174.14.70.100       (Assign IP address for the interface)
>> IP Interface 1# ena                      (Enable IP interface #1)
```


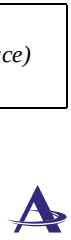

3. **On the Denver switch, define the default gateway.**

```
>> IP Interface 1# ../gw 1                    (Select default gateway #1)
>> Default gateway 1# addr 174.14.70.101      (Assign IP address for the gateway)
>> Default gateway 1# ena                     (Enable default gateway #1)
```

4. **Configure the local DNS server to recognize the local DSLB switch as the authoritative name server for the hosted services.**

   The Denver DNS server is configured to recognize 174.14.70.100 (the IP interface of the Denver DSLB switch) as the authoritative name server for www.foocorp.com).

Part Five: Configure the Denver Switch for Standard SLB

1. **Assign an IP address to each of the real servers in the local Denver server pool.**

Table 17-3 Denver Real Server IP Addresses

Real Server	IP address
Server C	179.14.70.2
Server D	179.14.70.2

2. **On the Denver switch, define each local Real Server.**

```
>> Default gateway 1# /cfg/slb/real 1         (Server C is real server 1)
>> Real server 1 # rip 179.14.70.2            (Assign Server C IP address)
>> Real server 1 # ena                        (Enable real server 1)
>> Real server 1 # ../real 2                  (Server D is real server 2)
>> Real server 2 # rip 179.14.70.3            (Assign Server D IP address)
>> Real server 2 # ena                        (Enable real server 2)
```

3. **On the Denver switch, define a Real Server Group.**

```
>> Real server 2 # /cfg/slb/group 1           (Select real server group 1)
>> Real server group 1# add 1                 (Add real server 1 to group 1)
>> Real server group 1# add 2                 (Add real server 2 to group 1)
>> Real server group 1# healt http            (Use HTTP for health checks)
>> Real server group 1# cntnt index.html      (Set URL content for health checks)
```

4. **On the Denver switch, define a Virtual Server.**

```
>> Real server group 1 # /cfg/slb/virt 1      (Select virtual server 1)
>> Virtual server 1# vip 179.14.70.1          (Assign a virtual server IP address)
>> Virtual server 1# add http 1               (Associate virtual port to real group)
>> Virtual server 1# ena                      (Enable the virtual server)
```

5. **On the Denver switch, define the Port States.**

In this example, the following ports are being used on the ACEswitch 180:

Table 17-4 Web Host Example: ACEswitch 180 Port Usage

Port	Host	Port Setting
3	Server C	`server`
4	Server D	`server`
5	Default Gateway Router. This connects the switch to the Internet where all client requests originate.	`client`

The ports are configured as follows:

```
>> Virtual server 1# /cfg/slb/port 3        (Select physical switch port 3)
>> SLB port 3# state server                 (Assign port 3 to server traffic)
>> SLB port 3# ../port 4                    (Select physical switch port 4)
>> SLB port 4# state server                 (Assign port 4 to server traffic)
>> SLB port 4# ../port 5                    (Select physical switch port 5)
>> SLB port 5# state client                 (Set port 5 for client traffic)
```

6. **On the Denver switch, enable Server Load Balancing.**

```
>> SLB port 5# ..                           (Select the SLB Menu)
>> Server Load Balancing# on                (Turn Server Load Balancing on)
```

Part Six: Configure the Denver Site for DSLB

Following the same procedures as above, here is a summary of the configuration steps for the Denver site.

1. **On the Denver switch, define each remote site.**

Since we are now configuring the Denver site, Denver is local and California is remote. Add and enable the IP address for the IP interface of up to eight remote sites. In this example, there is only one remote site: California, with an IP interface address of 200.200.200.100. The following commands are used:

```
>> Server Load Balancing# /cfg/dist/site 1  (Select Remote Site #1)
>> Remote site 1# prima 200.200.200.100     (Define remote IP interface address)
>> Remote site 1# ena                       (Enable remote site #1)
```

2. **On the Denver switch, assign each remote distributed service to a local virtual server.**

NOTE – This step can result in improper configuration if not clearly understood. Please take care to note where each configured value originates.

In this step, we are configuring the local Denver site to recognize the services offered at the remote California site. As before, configure one real server entry on the Denver switch for each virtual server located at each remote site. Since there's only one remote site (California) with only one virtual server, only one more local real server entry is needed at the Denver site.

The new real server entry will be configured with the IP address of the remote virtual server, rather than the usual IP address of a local physical server.

Also, the "remote" property will be enabled, and the real server entry will be added the real server group under the local virtual server for the intended service. Finally, since the real server health checks will be headed across the Internet, the health checking interval should be increased to 30 or 60 seconds to avoid generating excess traffic. For example:

```
>> Remote site 1# /cfg/slb/real 3          (Create an entry for real server #3)
>> Real server 3# rip 200.200.200.1        (Set remote virtual server IP address)
>> Real server 3# remote enable            (Define the real server as remote)
>> Real server 3# intr 60                  (Set a high health check interval)
>> Real server 3# ena                      (Enable the real server entry)
>> Real server 3# ../group 1               (Select the approp. real server group)
>> Real server group 1# add 3              (Add real server 3 to the group 1)
```

NOTE – The IP address of the real server being added is taken from the virtual server IP address on the remote switch. Do not confuse this value with the IP interface address on the remote switch.

---

3. **On the Denver switch, define the domain name and hostname for each service hosted on each virtual server.**

These will be the same as for the California switch: the domain name is "foocorp.com," and the hostname for the HTTP service is "www." These values are configured as follows:

```
>> Real server group 1# /cfg/slb/virt 1    (Select virtual server #1)
>> Virtual server 1# dname foocorp.com     (Define domain name)
>> Virtual server 1# hname http www        (Define HTTP hostname)
```

4. **On the Denver switch, turn Distributed Server Load Balancing on.**

```
>> Virtual server 1# /cfg/dist/on          (Activate DSLB for the switch)
```

5. **Apply and verify the configuration.**

```
>> Distributed SLB menu# apply          (Make your changes active)
>> Distributed SLB menu# cur            (View current DSLB settings)
>> Distributed SLB menu# ../cur         (View current SLB settings)
```

Examine the resulting information. If any settings are incorrect, make and apply any appropriate changes, and then check again.

6. **Save your new configuration changes.**

```
>> Server Load Balancing# save          (Save for restore after reboot)
```

Basic Tests for DSLB Operation

- Execute a browser request to the configured service (www.foocorp.com in the previous example).
- On each switch, examine the following statistics:
  - `/stats/slb/dist/virt` *virtual-server-number*
  - `/stats/slb/dist/group` *real-server-group-number*
  - `/stats/slb/maint`
- On each switch, examine the `/info/slb` information.

  Check that all Server Load Balancing parameters are working according to expectation. If necessary, make any appropriate configuration changes and then check the information again.

CHAPTER 18

Troubleshooting

This chapter describes the most common problems that might occur with the switch, lists the probable causes for the problems, and defines possible solutions.

Definitions

- **Management Processor (MP)**

  The processor that handles management of the switch. It processes the CLI, Telnet, SNMP operation, and Spanning-Tree.

- **Switch Processor (SP)**

  The switch processor that processes both switched user frames and switched management frames.

- **Forwarding Database (FDB)**

  This is the database of learned and being-learned MAC addresses.

- **Spanning-Tree Protocol (STP)**

  The IEEE 802.1d specified loop prevention protocol widely used in Ethernet bridge networks.

- **Bridge Protocol Data Unit (BPDU)**

  Frames used to convey Spanning-Tree information to form a loop-free network topology.

System Problems

Switch Management Problems

Cannot `ping` a switch IP interface. Cannot Telnet to a switch IP interface. MIB Browser cannot discover the switch. The switch does not send SNMP traps.

Possible Causes

- Incorrect switch IP interface configuration
- Link state of the port the `ping`ing station is connected to is the in the "down" state
- Spanning-Tree port state is not in "forwarding" state
- Incorrect SNMP community strings
- Trap server is not configured
- Switch IP interface address is used by some other device in the network

Actions

- Check `/cfg/ip/cur` to be sure the switch IP interface addresses, subnet masks, and default gateways are correctly configured, and that the IP interfaces are enabled.
- Check `/info/link` to be sure the management port link is in the "up" state.
- Check `/info/stp` to be sure port Spanning-Tree is in "forwarding" state.
- Check `/cfg/snmp/cur` to be sure SNMP community strings are correct.
- Check `/cfg/snmp/cur` to be sure the Trap server is specified.
- Check for duplicate IP address and correct if necessary.

Link Problems

Green link LED does not come on. Link state is in "down" state from the CLI (`/info/link`).

Potential Causes

- Port Configuration mismatch between the switch and the other device
- Different version of Link Negotiation used between the switch and the other device
- Bad or incorrect cable

Actions

- If ports are configured with specific values such as 100Mbs speed, then make sure the other device is configured the same way.

- Port Configuration: Make sure both the switch port and the other device are configured with the same negotiation mode. If the switch port is configured with either Speed or Duplex mode in "auto," the other device must have the same configuration.
- Check the cabling between the switch and the other device. If the other device is a workstation, straight through cable should be used. However, if it is either another switch or a hub, a cross-over cable should be used unless there is an "uplink" enable/disable switch used instead on the switch or hub.

Table 18-1 Pin-outs for Crossover cable

pin 1 ---------- pin 3
pin 2 ---------- pin 6
pin 3 ---------- pin 1
pin 6 ---------- pin 2

NOTE – These pin-outs are for the 10/100 Mbps physical ports only.

- Check link status in `/info/link`. If link state is "up", then the problem is a bad LED.

SNAP Traces

If a console is hooked up to the switch, a message will indicate that the switch had taken a "snap trace."

Possible Causes

- Watchdog Timer: If the MP fails to refresh the on-board timer, this will reset the processor, initiating a snap trace and reset of the switch.
- Different software resets: When encountering certain error conditions or anomalies, the software will trigger a panic which in turn will generate a snap trace, coredump, and reset the switch.

Actions

- Messages: Any message(s) on the console should be recorded and sent to Alteon Networks Customer Support.
- Coredump: Retrieve the coredump (if available) by accessing the Maintenance menu and invoking the `uudmp` option. Alternately, you can enter **`/maint/uudmp`** to retrieve the coredump. Any coredump should be sent to Alteon Networks Customer Support.

Switch Boot Failure

The switch will not boot.

Possible Causes

- Corrupted firmware
- Firmware and configuration was corrupted when rebooting with an older firmware image. This can occur when replacing Release 4 software with Release 3.0.20 or earlier software without first resetting the switch to factory default configuration.

Actions

Replace the corrupted firmware by performing a serial download of a new binary firmware image.

NOTE – The procedure for serial download is different from the procedure for TFTP download.

This procedure requires the following:

- A computer running terminal emulation software
- A standard serial cable with a male DB9 connector (see your switch hardware installation guide for specifics)
- A *binary* switch firmware image (*not* the `tfpt` file used for TFTP download)

Procedure

1. **Using the serial cable, connect the computer to the switch Console port (Serial Port on some models).**
2. **Make sure that the new binary firmware file is available on the computer.**
3. **Start your terminal emulation software and set the communication parameters:**

Table 18-2 Console Configuration Parameters

Parameter	Value
Baud Rate	9600
Data Bits	8
Parity	None
Stop Bits	1

4. **Turn on the switch power and press <Shift-F> while the switch is first attempting to boot.**

   When performed correctly, the following message appears:

```
Xmodem flash download 1.0.5
To download to flash use xmodem at 57600 baud
Power cycle to end xmodem.
```

5. **Reconfigure your terminal emulation software for the following parameters:**

Parameter	Value
Baud Rate	57,800
Data Bits	8
Parity	None
Stop Bits	1

6. **Set the file transfer mode to Xmodem.**

7. **Transfer the binary firmware image file to the switch.**

   This process can take three or four minutes to complete. When finished, the message "done" will appear on your terminal.

8. **Disconnect the terminal emulation session and reconfigure your terminal emulation software for normal switch connection parameters:**

Parameter	Value
Baud Rate	9600
Data Bits	8
Parity	None
Stop Bits	1

9. **Reconnect the terminal session to the switch.**

10. **Turn the switch power off, and then back on again.**

    The switch should now boot normally.

Switching Problems

This section lists the most common switching problems, their causes, and solutions.

Connectivity Problems

Client "A" on port 1 cannot connect to server "B" on port 2.

Potential Causes

- Incorrect configuration of client/server machines: the IP address is wrong.
- Ports 1 or 2 may be down (link down).
- Spanning-Tree Port State is not in "forwarding" state.
- Frames from either "A" or "B" are received with errors or not transmitted due to error conditions on outgoing port.
- MAC Address of either "A" or "B" is learned incorrectly from ports other than 1 and 2.

Actions

- Check `/info/link` to be sure link state is up.
- Check `/info/stp` to be sure Spanning-Tree Port is in "forwarding" state.
- Check port interface statistics (`/stats/port` *port-number*`/if`) to see whether `ifInErrors`, `ifInDiscards`, `ifOutErrors`, or `ifOutDiscards` are incrementing.
  - `ifInErrors`: MAC errors
  - `ifInDiscards`: STP blocking state, filtering, frame errors, PCI busy
  - `ifOutErrors`: not used
  - `ifOutDiscards`: due to backup on link
- Check port `dot3` statistics (`/stats/port 1/ether`) for Ethernet specific errors.
- Search MAC addresses for "A" and "B" from the FDB. For example, if A's MAC address is 00:00:00:00:00:01 and B's is 00:00:00:00:00:02, search for A's MAC address by typing the following from the CLI: **`/info/fdb/find 00:00:00:00:00:01`**

  Output similar to the following example should be displayed.

```
MAC Address Port State Referenced from Ports...
00:00:00:00:00:01 1 FWD
```

Spanning-Tree Protocol Problems

The topology in the following figure is used to illustrate the STP problems in this section.

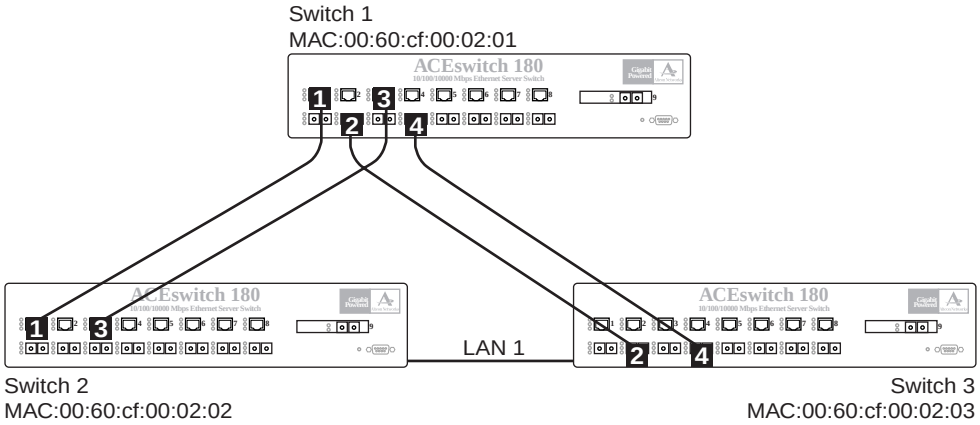

Figure 18-1 Spanning-Tree Topology

All switches have the default STP parameters except the following:

- Switch 1 MAC: 00:60:cf:00:02:01
- Switch 2 MAC: 00:60:cf:00:02:02, Path cost for port 1 (to Switch 1) is 10. Path cost for port 3 (to Switch 1) is 5.
- Switch 3 MAC: 00:60:cf:00:02:03, Path cost for port 2 and port 4 (to Switch 1) is 1.

Switch Receives its own Spanning-Tree BPDU Message

If the switch software receives its own bridge protocol data unit (BPDU) message, the switch port will be disabled. As an example, this could occur when the switch transmits the BPDU message out switch port 1 to a hub that has two hub ports connected together in a loop.

You must remove the loop from the port and manually re-enable the switch port. To manually re-enable the switch port, enter the following command:

```
Main# /oper/port port-number/enable
```

Spanning-Tree Recalculation

The IEEE 802.1d Spanning-Tree algorithm can take up to 45 seconds from the time it detects a topology change to the time it transitions from "spanning-tree port" state to "forwarding" state. During Spanning-Tree recalculation, frame forwarding from the port will stop and interrupt normal network traffic flow. Unlike shared media environments, in a switched network environment when the end station directly connected to a switch port is rebooted, it causes the switch port link state to change, resulting in recalculation of the "spanning-tree port" state. This is seen by loss of connection upon end station reboot.

Server Load Balancing Configurations

General

The following checklist will help you resolve the most common difficulties configuring Server Load Balancing.

- Check the Server Load Balancing maintenance statistics (page 6-13) and the Server Load Balancing information (page 5-8) for anything unexpected.
- On the switch, check that the real servers, real server groups, virtual servers, etc. have been *enabled*.
- Check that the real servers are physically functioning.
- Check that all the services which the switch is expecting to find on each real server are installed, configured, and running properly.
- On the switch, make sure that you used `apply` and `save` to activate your configuration changes (see "Viewing, Applying, and Saving Changes" on page 7-2).
- On the switch, make sure that the real servers were added to the proper real server groups and that the real server groups are associated with a virtual server.
- Make sure that you are not violating any of the network topology restrictions, such as by connecting clients and servers to the same switch port (see "Network Topology Considerations" on page 15-4).
- Make sure that the port state for each switch port is properly configured as `client`, `server`, or `none` (see "The SLB Port Menu" on page 7-47).
- Make sure that the switch is configured to accept the TCP/UDP port numbers on which each particular service is expected to run.
- If a service on a real server runs on a different port number than typical (such as using TCP port 8000 for HTTP, instead of TCP port 80), make sure that the virtual port and real port are properly mapped ("Mapping Virtual Ports to Real Ports" on page 7-42).

Service Problems

Periodic loss of a configured TCP service (such as HTTP). Real server does not come into service, or comes into service and fails periodically.

Possible Causes

- Invalid topology or port state: the real server is connected to the switch through a port configured in the `client` or `redir` state.
- There may be a health-check failure between the switch and the real server.
- One of the real servers of the real server group does not respond to the service request.

Actions

- Monitor the health checks. At Layer 4, there should be a 3-way TCP handshake for opening a TCP connection, followed by a 4-way TCP handshake to close a TCP connection.
- Verify that the real server has a default gateway or a route back to the client.
- Verify that the requested HTTP object is present on every real server in the real server group.

Miscellaneous

LED Patterns on Gigabit Ethernet Ports

LED patterns on Gigabit Ethernet Ports 9 and 10 are different upon switch reset. Both LEDs on Port 10 (in the switch I/O Module slot) come up OFF, and all LEDs on Port 9 come up ON.

Lost Character Output on Console Port

Characters written to the console are sometimes lost. This problem occurs rarely, but it can be seen as misaligned output or missing prompts. A missing prompt might appear to be a switch hang, but pressing Return or Control-C (^C) will cause the prompt to be repeated, returning the switch to normal operation.

Index

Symbols

Numerics

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

Q

R

S

T

U

V

W

X