Web OS Switch Software



7.2 Command Reference

Part Number: 050155, Revision A, June 2001

Alteon WebSystems
Web Speed for e-Business

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

Copyright 2001 Alteon WebSystems, Inc., 50 Great Oaks Boulevard, San Jose, California 95119, USA. All rights reserved. Part Number: 050155, 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 WebSystems, Inc. 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 WebSystems, Inc. reserves the right to change any products described herein at any time, and without notice. Alteon WebSystems, Inc. assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Alteon WebSystems, 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 WebSystems, Inc.

Web OS is a trademarks of Alteon WebSystems, 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

Preface 11
Who Should Use This Guide 11
How This Book Is Organized 11
Typographic Conventions 12
Contacting Alteon WebSystems 13

Chapter 1: The Command Line Interface 15
Connecting to the Switch 16
- Establishing a Console Connection 16
- Establishing a Telnet Connection 17
- Establishing an SSH Connection 18

Accessing the Switch 21
- Disabling Access Levels 22

CLI Versus Setup 23
Idle Time-Out 23

Chapter 2: First Time Configuration 25
Using the Setup Utility 25
- Information Needed For Setup 25
- Starting Setup Under Factory Default Conditions 26
- Stopping and Restarting Setup Manually 27
- Setup Part 1: Basic System Configuration 27
- Setup Part 2: Port Configuration 30
- Setup Part 3: VLANs 32
- Setup Part 4: IP Configuration 34
- Setup Part 5: Final Steps 36

Setting Passwords 37
- Changing Access Passwords 37

Chapter 3: Menu Basics 41
The Main Menu 41
Menu Summary 42
Global Commands 43
Command Line History and Editing 44
Command Line Interface Shortcuts 44
Command Stacking 44
Command Abbreviation 44
Tab Completion 45

Chapter 4: The Information Menu 47
Information Menu 47
System Information 51
Environment Information 52
Show Last 10 Syslog Messages 53
Link Status Information 54
Port Information 55
Slot Status Information 56
Interpreting SP Index Numbers 57
VLAN Information 58
Spanning Tree Information 59
Trunk Group Information 61
Address Resolution Protocol Information 62
Show All ARP Entry Information 63
ARP Address List Information 63
FDB Information 64
Show All FDB Information 65
Display IP Information 66
Border Gateway Protocol Information 66
BGP Peer information 67
BGP Summary information 67
Dump BGP Information 68
IP Information 69
IP Routing Information 71
Show All IP Route Information 72

OSPF Information Menu 74
General OSPF Information. 75
OSPF Area Information 75
OSPF Interface Information 76
OSPF Neighbor Information 76
OSPF Retransmission List 77
OSPF Routes 77
OSPF Database Menu 78
Summary of All 79
SLB Information Menu 80
Show Session Table Information 81
Show All Layer 4 Information 82
VRRP Information 83
Software Enabled Keys 84
Information Dump 84

Chapter 5: The Statistics Menu 85

Statistics Menu 85
Port Statistics Menu 87
Bridging (*dot1*) Statistics 88
Ethernet (dot3) Statistics 89
Interface (IF) Statistics 93
Link Statistics 95
RMON Statistics 95
Server Load Balancing Statistics 99
Real Server SLB Statistics 101
Real Server Group Statistics 102
Virtual Server SLB Statistics 102
Filter SLB Statistics 103
Switch Processor SLB Statistics 103
SP Real Server SLB Statistics 104
SP Real Server Group SLB Statistics 104
SP Virtual Server SLB Statistics 105
SP Filter SLB Statistics 105
SP Maintenance SLB Statistics 106
Global SLB Statistics 108
Real Server Global SLB Statistics 109
Real Server Group Global SLB Statistics 109
Virtual Server Global SLB Statistics 110
Global SLB Maintenance Statistics 111
SLB Maintenance Statistics 111

Management Processor Statistics 112
IP Interface Statistics for Management Processor 114
IP Statistics for Management Processor 115
ICMP Statistics for Management Processor 118
TCP Statistics for Management Processor 120
UDP Statistics for Management Processor 122
SNMP Statistics for Management Processor 122
Address Resolution Protocol Statistics for MP 125
Route Statistics for Management Processor 126
DNS Statistics for Management Processor 127
OSPF Statistics Menu for Management Processor 127
VRRP Statistics for Management Processor 128
Maintenance Statistics for Management Processor 128
Letter Statistics for Management Processor 130
STEM Message Statistics for Management Processor 131
All STEM Memory Allocated Blocks for Management Processor 132
STEM Thread Statistics for Management Processor 133
Packet Counts for Management Processor 134
All TCB Allocated Control Blocks for Management Processor 134
UDP Allocated Control Blocks for Management Processor 134
UART Statistics for Management Processor 135
CPU Statistics for Management Processor 135
Switch Processor Statistics 136
IP Statistics for Switch Processor 136
Maintenance Statistics for Switch Processor 138
Letter Counts for Switch Processor 139
STEM Message Statistics for Switch Processor 140
STEM Memory Statistics for Switch Processor 140
STEM Thread Statistics for Switch Processor 141
CPU Statistics for Switch Processor 141
Statistics Dump 142

Chapter 6: The Configuration Menu 143
Viewing, Applying, and Saving Changes 143
Viewing Pending Changes 143
Applying Pending Changes 144
Saving the Configuration 144
Configuration Menu 145
System Configuration 147
System Host Log Configuration 149
Radius Server Configuration 150
User Access Control Configuration 151
SSH Server Configuration 153
User-Defined Defaults Configuration 154
User-Defined Link Defaults Configuration 155
Ethertype Defaults Configuration 156
Port Configuration 157
Port Link Configuration 158
Port VLAN ID (PVID) Configuration 160
User PVID Port Configuration 161
VLAN Configuration 162
Spanning Tree Configuration 164
Bridge Spanning Tree Configuration 165
Spanning Tree Port Configuration 167
SNMP Configuration 168
Trap Host SNMP Configuration 170
Trunk Configuration 171
Port VLAN ID Trunk Configuration 173
User PVID Port Configuration Menu 174
IP Configuration 175
IP Interface Configuration 177
Default IP Gateway Configuration 178
IP Static Route Configuration 179
IP Forwarding Configuration 180
IP Network Filters Configuration 181
Routing Map Configuration 182
IP Access List 184
Autonomous System Filter Path 185
Routing Information Protocol Configuration 186
Routing Information Protocol Interface Configuration 187

Open Shortest Path First Configuration 189
OSPF Area Configuration 191
OSPF Summary Range Configuration 192
OSPF Interface Configuration 194
OSPF Virtual Link Configuration 196
OSPF Host Entry Configuration 198
OSPF Host Entry Configuration 199
Border Gateway Protocol Configuration 200
BGP Peer Configuration 201
BGP Redistribution Configuration 203
Aggregate Routing Configuration 204
IP Forwarding Port Configuration 205
Domain Name System Configuration 206
Bootstrap Protocol Relay Configuration 207
Default Gateway Metrics Configuration 208
SLB Configuration 209
Real Server SLB Configuration 211
Real Server Group SLB Configuration 214
Virtual Server SLB Configuration 218
Virtual Server Service Configuration 220
SLB Filter Configuration 222
Advanced Filter Configuration 227
Port SLB Configuration 228
Global SLB Configuration 230
GSLB Remote Site Configuration 232
Global Server Load Balance Lookup Configuration 234
GSLB Internet Lookup Network Preference Configuration 234
Synchronizing Peer Switch Configurations 236
Peer Switch Configuration 237
Advanced Layer 4 Configuration 238
Scriptable Health Checks Configuration 240
VRRP Configuration 242
Virtual Router Configuration 243
Virtual Router Priority Tracking Configuration 245
Virtual Router Group Configuration 247
Virtual Router Group Priority Tracking Configuration 249
VRRP Interface Configuration 250
VRRP Tracking Configuration 251

Chapter 7: The Operations Menu 253
Operations Menu 253
Operations-Level Port Options 255
Operations-Level Switch Processor Options 256
Operations-Level Line Card Module Options 257
Operations-Level SLB Options 258
Operations-Level VRRP Options 259
Operations-Level IP Options 259
Activating Optional Software 260
Removing Optional Software 261

Chapter 8: The Boot Options Menu 263
Updating the Switch Software Image 264
Downloading New Software to Your Switch 264
Selecting a Software Image to Run 266
Selecting a Configuration Block 267
Resetting the Switch 268

Chapter 9: The Maintenance Menu 269
Maintenance Menu 269
FLASH Dump Manipulation Options 271
Uuencode FLASH Dump 272
TFTP System Dump Put 273
Clearing Dump Information 273
System Maintenance Options 274
Software Build Information 275
Switch Processor Maintenance Options 275
Forwarding Database Options 275
Debugging Options 277
SP Xtrace Buffer 279
ARP Cache Options 279
IP Route Manipulation 281
Panic Command 282
Tech Support Dump 283
Unscheduled System Dumps 283

Chapter 10: Troubleshooting 285
Definitions 285
System Problems 286
Switch Management Problems 286
Link Problems 287
SNAP Traces 288
Switch Lockup Problems 288
Switch Boot Failure 289
Switching Problems 291
Connectivity Problems 291
Spanning Tree Protocol Problems 292
Switch Receives Its Own Spanning Tree BPDU Message 292
Spanning Tree Recalculation 293
Server Load Balancing Configurations 293
General 293
Service Problems 294

Appendix A: Defining IP Address Ranges for Filters 295

Index 297

Preface

This manual describes how to configure and use the Web OS 7.2 software included in the Alteon 708 Web switch.

For documentation on installing the switches physically, see the hardware installation guide.

Who Should Use This Guide

This manual is intended for network installers and system administrators engaged in configuring and maintaining a network. The network administrators should be familiar with Ethernet concepts, IP addressing, the IEEE 802.1d Spanning Tree Protocol, and SNMP configuration parameters.

How This Book Is Organized

Chapter 1, "The Command Line Interface," describes how to connect to the switch and access the information and configuration menus.

Chapter 2, "First Time Configuration," describes how to use the Setup utility for initial switch configuration and how to change the system passwords.

Chapter 3, "Menu Basics," provides an overview of the menu system, including a menu map, global commands, and menu shortcuts.

Chapter 4, "The Information Menu," shows how to view switch configuration parameters.

Chapter 5, "The Statistics Menu," shows how to view switch performance statistics.

Chapter 6, "The Configuration Menu," shows how to configure switch system parameters, ports, VLANs, Spanning Tree Protocol, SNMP, IP Routing, Port Trunking, Server Load Balancing, Filtering, and more.

Chapter 7, "The Operations Menu," shows how to use commands that affect switch performance immediately, but do not alter permanent switch configurations, such as temporarily disabling ports. This chapter also describes how to activate or deactivate optional software features.

Chapter 8, "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 9, "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.

Chapter 10, "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. Do not type the brackets.	To establish a Telnet session, enter: `host#` **`telnet`** *<IP address>*
	This also shows book titles, special terms, or words to be emphasized.	Read your *User's Guide* thoroughly.
[ ]	Command items shown inside brackets are optional and can be used or excluded as the situation demands. Do not type the brackets.	`host#` **`ls`** [**`-a`**]

Contacting Alteon WebSystems

Use the following information to access Alteon WebSystems' support and sales.

- URL for Alteon WebSystems Online:

  `http://www.alteonwebsystems.com`

  This website includes product information, software updates, release notes, and white papers. The website also includes access to Alteon WebSystems Customer Support for accounts under warranty or covered by a maintenance contract.

- E-mail access:

  `support@alteon.com`

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

- Telephone access to Alteon WebSystems Customer Support:

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

  Telephone access to Alteon WebSystems 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. Pacific Standard Time.

- Telephone access to Alteon WebSystems 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.

CHAPTER 1

The Command Line Interface

Your Alteon Web 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 Web OS 7.2 software included in your switch provides a variety of options for accessing and configuring the switch:

- The Command Line Interface (CLI) provides a built-in, text-based menu system for access via local terminal or remote Telnet session.
- The Browser-Based Interface (BBI) provides a Web-based management interface for interactive network access through your Web browser.
- SNMP support provides access through network management software such as HP-OpenView.

The 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 you can use to view information and statistics about the switch and perform any necessary configuration.

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

Connecting to the Switch

You can access the CLI in three ways:

- Using a console connection via the console port
- Using a Telnet connection over the network
- Using an SSH connection to securely log into another computer 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 1-1 Console Configuration Parameters

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

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

Procedure

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

   Next, you will be required to enter a password for access to the switch.

NOTE – It is recommended that you change default switch passwords after initial configuration and as regularly as required under your network security policies. For more information, see "Setting Passwords" on page 38.

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 4: IP Configuration" on page 34).

Using a BOOTP Server

By default, the Web OS 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 51).

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 10, "Troubleshooting."

Establishing an SSH Connection

Although a remote network administrator can manage the configuration of an Alteon Web switch via Telnet, this method does not provide a secure connection. The Secure Shell (SSH) protocol enables you to securely log into another computer over a network to execute commands remotely. As a secure alternative to using Telnet to manage switch configuration, SSH ensures that all data sent over the network is encrypted and secure.

SCP (Secure Copy) is used on an Alteon Web switch to download and upload the switch configuration via secure channels. SCP uses SSH for encryption of data on the network.

NOTE – When the SSH feature is enabled, SCP is also enabled.

The supported encryption and authentication methods for both SSH and SCP are listed below.

- Server Host Authentication: Client RSA—authenticates the switch in the beginning of every connection
- Key Exchange: RSA
- Encryption: 3DES-CBC, DES
- User Authentication: local password authentication, Radius, SecurID (via Radius for SSH only—does not apply to SCP)

NOTE – The Web OS implementation of SSH is based on SSH version 1.5 and supports SSH-1.5-1.X.XX. SSH clients of other versions (especially Version 2) will not be supported.

The switch can do only one session of key/cipher generation at a time. Thus, an SSH/SCP client will not be able to log in if the switch is doing key generation at that time, or if another client has just logged in before this client. Similarly, the system will fail to do the key generation if a SSH/SCP client is logging in at that time.

Configuring SSH/SCP

NOTE – SSH/SCP can only be configured via the console port, using the CLI. You cannot configure SSH/SCP via telnet.

To enable/disable the SSH/SCP feature, use this command:

```
>> # /cfg/sys/sshd on
```

To set the interval of RSA server key autogeneration, use this command:

```
>> # /cfg/sys/sshd/intrval <n>
```

where *n* (number of hours) must be in the range (from 0 to 24) and a value of 0 denotes that RSA server key autogeneration is disabled. When *n* is greater than 0, the switch will autogenerate the RSA server key every *n* hours. However, RSA server key generation will be skipped if the switch is busy doing other key or cipher generation when the timer expires.

To view the current SSH/SCP related configuration, use this command:

```
>> # /cfg/sys/sshd cur
```

To view the difference between the new configuration and the current configuration, use this command:

```
>> # diff
```

To apply the pending changes from the new configuration to the current configuration, use this command:

```
>> # apply
```

NOTE – If SSH/SCP is enabled and an `apply` command is issued, the switch will automatically generate the RSA host and server keys if they are not available. It will take several minutes to complete this process.

To save the current configuration to FLASH memory, use this command:

```
>> # save
```

Usually, there will be no need to manually generate the RSA host and server keys. However, you may still do so by using the following commands:

```
>> # /cfg/sys/sshd/hkeygen (Generates the host key)
>> # /cfg/sys/sshd/skeygen (Generates the server key)
```

NOTE – These two commands will take effect immediately without an **`apply`** command.

Some Supported Client Commands

Up to four simultaneous Telnet/SSH/SCP connections are supported on a switch.

- To log in to the switch:

  ssh <*switch IP*>or **ssh -l** <*username*> <*switch IP*>

- To download the switch configuration using SCP:

  scp <*switch IP*>**:getcfg** <*local filename*>

- To upload the configuration to the switch:

  scp <*local filename*> <*switch IP*>**:putcfg**

Some examples are listed below:

```
>> # ssh 205.178.15.157
>> # ssh -l dleu 205.178.15.157
>> # scp 205.178.15.157:getcfg 708.cfg
>> # scp 708.cfg 205.178.15.157:putcfg
```

where 205.178.15.157 is the IP address of the switch.

Please also note that `apply` and `save` commands are still needed after the last command `(scp 708.cfg 205.178.15.157:putcfg)`is issued. Or, instead, you can use the following three commands to avoid the “apply and save” issue:

```
>> # scp 708.cfg 205.178.15.157:putcfg
>> # scp 708.cfg 205.178.15.157:putcfg_apply
>> # scp 708.cfg 205.178.15.157:putcfg_apply_save
```

Accessing the Switch

To enable better switch management and user accountability, three levels of user access have been implemented on the switch. Access levels are defined as follows:

- User interaction with the switch is passive—nothing can be changed on the switch. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.
- Operators can only effect temporary changes on the switch. These changes will be lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the switch, operators cannot severely impact switch operation. The following types of operators are supported: Server Load Balancing Operator, Layer 4 Operator, and System Operator.
- Administrators are the only ones that may make permanent changes to the switch configuration—changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the switch. Because administrators can also make temporary (operator level) changes as well, they must be aware of the interactions between temporary and permanent changes. The following types of administrators are supported: Server Load Balancing Administrator, Layer 4 Administrator, and System Administrator.

Once you are connected to the switch via local console or Telnet, you are prompted to enter a password. The default password for each access level is listed in the following table.

The default password of some user access levels is NULL. In this case, the user access is disabled until a password is assigned. To disable a user, set the password to NULL.

Table 1-2 User Access Levels

User Account	Description and Tasks Performed	Password
User	The user has no direct responsibility for switch management. He or she can view switch status information and statistics but cannot make any configuration changes.	`user`
SLB Operator	The SLB Operator manages Web servers and other Internet services and their loads. In addition to being able to view all switch information and statistics, the SLB Operator can enable/disable servers using the SLB operation menu.	`NULL`

Table 1-2 User Access Levels (Continued)

User Account	Description and Tasks Performed	Password
Layer 4 Operator	The Layer 4 Operator manages traffic on the lines leading to the shared Internet services. Currently, user has the same level of access as the Layer 4 Administrator.	NULL
System Operator	The Operator manages all functions of the switch. In addition to SLB Operator functions, the Operator can reset ports or the entire switch.	NULL
SLB Administrator	The SLB Administrator configures and manages Web servers and other Internet services and their loads. In addition to SLB Operator functions, the SLB Administrator can configure parameters on the SLB menus, with the exception of not being able to configure filters.	NULL
Layer 4 Administrator	The Layer 4 Administrator configures and manages traffic on the lines leading to the shared Internet services. In addition to SLB Administrator functions, the Layer 4 Administration can configure all switch information and statistics, and all parameters on the Server Load Balancing menus.	NULL
System Administrator	The Administrator has complete access to all menus, information, and configuration commands on the switch, including the ability to change all access passwords.	admin

NOTE – It is recommended that you change default switch passwords after initial configuration and as regularly as required under your network security policies. For more information, see "Setting Passwords" on page 38.

Disabling Access Levels

With the exception of the administrator account, access to a given level can be disabled by setting the password to an empty value. All accounts other than the administrator will be initially disabled (empty password) until they are enabled by the admin user. This is done in order to avoid inadvertently leaving the switch open to unauthorized users.

For more information, see "Setting Passwords" on page 38.

CLI Versus 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 2, "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.

Following is an example of 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]
      revert  - Revert pending or applied changes [global command]
      exit    - Exit [global command, always available]
```

NOTE – If you are the accessing the switch with an account that has less privilege than the Administrator account, some menu options will not be available.

Idle Time-Out

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, which can be set from 1 to 60 minutes. For information on changing this parameter, see "System Configuration" on page 147.

CHAPTER 2

First Time Configuration

To help with the initial process of configuring your switch, the Web OS 7.2 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. This chapter describes how to use the Setup utility and how to change system passwords.

Using the Setup Utility

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 (CLI) any time after logging in.

Information Needed For Setup

Setup requests the following information:

- Basic system information
  - Date and time
  - Whether or not to use BOOTP
  - Whether or not to use Spanning Tree Protocol
- Optional configuration for each port
  - Speed and negotiation mode (as appropriate)
  - Whether or not to use VLAN tagging (as appropriate)
- Optional configuration for each VLAN
  - Name of VLAN
  - 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 or not IP forwarding is enabled
  - Whether or not the Routing Information Protocol supply is enabled

Starting Setup Under Factory Default Conditions

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

1. **Connect to the switch console.**
2. **After connecting, the login prompt will appear as shown below. Enter `admin` as the default administrator password.**

```
Enter Password:
```

3. **If the factory default configuration is detected, the system will show the following prompt. Enter `y` to configure the switch using Setup, or `n` to bypass the Setup utility:**

```
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] y

"Set Up" will walk you through the configuration of System Date
 and Time, BOOTP, IP address and subnet mask, Spanning Tree,
 Port and Link characteristics, and VLANs.
 [Type Ctrl-C to abort "Set Up"]
```

NOTE – If the default `admin` login is unsuccessful, or if the administrator Main Menu (see "The Main Menu" on page 41) 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 267. Otherwise start Setup manually as described below.

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 from the beginning.

Starting Setup

You can start 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, IP address and subnet mask, Spanning Tree,
 Port and Link characteristics, and VLANs.
 [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 VLAN issues, see “Setup Part 3: VLANs” on page 32.

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

2. **Enter the month of the current system date.**

```
System Date:
Enter month [5]:
```

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.**

```
Enter day [5]:
```

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.**

```
Enter year [01]:
```

Enter the last two digits of the year as a number from 00 to 99. “01” is considered 2001. To keep the current year, press <Enter>.

The system displays the date and time settings:

```
System clock set to 13:56:52 Fri May 5, 2001.
```

5. **Enter the hour of the current system time.**

```
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.**

```
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.**

```
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 Fri May 5, 2001.
```

8. **Enable or disable the use of BOOTP.**

```
BootP Option [must DISABLE to configure IP]:
Current BOOTP: disabled
Enter new BOOTP [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. BOOTP must be disabled, however, before the system will prompt for IP parameters.

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.**

```
Spanning Tree:
Current Spanning Tree Group 1 setting: OFF
Turn Spanning Tree Group 1 ON? [y|n]
```

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

Setup Part 2: Port Configuration

NOTE – The port configuration options shown in the following example steps are for a Fast Ethernet port on the Alteon 708. When configuring port options for other switches or port types, some of the prompts and options may be different.

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

```
Port Configuration:
Enter port number (e.g. b12): a1
```

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 32.

2. **If appropriate, configure link type.**

The system prompts for the type of link. Select the appropriate type based on the type of line card module installed for this port.

```
Link Type:
Enter new link type [none|FE|GE-SX|1000Base-T]: fe
New link configuration initialized to FE default.
```

3. **If appropriate, configure link speed.**

The system prompts for the type of link in the port you selected:

```
Link Speed:
Current Port A1 speed setting: none
Pending new speed setting: 10/100
Enter new speed [10/100/any]: any
```

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

4. **The link mode is set to `full duplex`.**

The system reports the fixed configuration and continues to the next step.

```
Link Mode:
Port A1 mode fixed at full duplex.
```

5. **The flow control is set to `none`.**

   The system reports the fixed configuration and continues to the next step.

```
Link Flow Control:
Port A1 flow control fixed at none.
```

6. **If appropriate, configure Ethernet/Fast Ethernet port autonegotiation mode.**

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

```
Link Autonegotiation:
Current Port A1 autonegotiation: none
Pending Port A1 autonegotiation: on
Enter new value [on|off]: on
```

   Enter **`on`** to enable autonegotiation, **`off`** to disable it, or press <Enter> to keep the current setting.

7. **If you have selected to configure VLANs previously in Part 1, the system prompts you to enter the default port-defined VLAN ID (PVID) for IP, ARP, and RARP packet types.**

   To keep the current setting, press <Enter>.

```
Default PVID for IP/ARP/RARP:
Current IP PVID: 1
Enter new PVID:
```

8. **If applicable, enter the PVID for IPX packet types.**

   To keep the current setting, press <Enter>.

```
Default PVID for IPX:
Current IPX PVID: 1
Enter new PVID:
```

9. **Enter the default PVID for all other Ethertypes.**

   To keep the current setting, press <Enter>.

```
Default PVID for all other Ethertypes:
Current other PVID: 1
Enter new PVID:
```

10. **Select whether to forward or discard tagged frames.**

    Enter **d** to discard VLAN tagged frames for the port, or enter **f** to forward them. To keep the current setting, press <Enter>.

```
Forward Tagged frames:
Current tagged frame status:            discard
Enter new tagged frame status [f|d]:    f
Port A1 changed to forward tagged frames.
```

11. **Select whether to forward or discard untagged frames.**

```
Forward Untagged frames:
Current untagged frame status: forward
Enter new untagged frame status [f|d]: d
Port A1 changed to discard untagged frames.
```

12. **The system prompts you to configure the next port.**

```
Enter port (e.g. b12):
```

When you have finished 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 from Part 1, skip to "Setup Part 4: IP Configuration" on page 34.

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

```
VLAN Config:
Enter VLAN number from 1 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 34.

2. **Enter the new VLAN name.**

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

3. **Enter the VLAN port numbers.**

   The system prompts you to define the ports in the VLAN. For example, ports defined for this VLAN could be as follows:

```
Define Ports in VLAN:
Current VLAN 1:  A1-16 B5-16 C1-16 D1-16 E1-16 F1-16 G1-16 H1-16 T1
T2 T3 T4
Enter ports one per line, NULL at end:
```

   Type the first port number (as slot and port together, such as b12) to add to the current VLAN and press <Enter>. The 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.

4. **Assign Spanning Tree Group membership to the VLAN you just configured.**

   The default Spanning Tree group index is 1. All VLANs must belong to a Spanning Tree group.

```
Spanning Tree Group membership:
Current Spanning Tree Group index: 1
Enter new Spanning Tree Group index [1-256]:
```

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 36. Otherwise, if you disabled BOOTP, the system prompts for IP parameters.

IP Interfaces

IP interfaces are used for defining subnets related to the switch.

Up to 1024 IP interfaces can be configured on the switch. The IP address assigned to each IP interface provides 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.**

```
IP Config:
IP interfaces:
Enter interface number: [1-1024]
```

If you wish to configure individual IP interfaces, enter the number of the IP interface you want to configure. To skip IP interface configuration, press <Enter> without typing an interface number and go to "Default Gateways" on page 35.

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

   To keep the current setting, press <Enter>.

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

3. **Enter the IP subnet mask in dotted decimal notation.**

   To keep the current setting, press <Enter>.

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

4. **Enter the broadcast IP address in dotted decimal notation.**

   To keep the current setting, press <Enter>.

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

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

   This prompt appears if you already chose to configure VLANs in Part 1:

```
Current VLAN:     1
Enter new VLAN: N [1-4094]:
```

   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. **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: [1-1024]
```

   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. **Pick a default gateway to configure, 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 36.

2. **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. **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. **Enable or disable forwarding for IP Routing.**

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

Enter **y** to enable IP forwarding, or **n** to disable it. To keep the current setting, press <Enter>.

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.**

```
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.

Only the administrator has the ability to change passwords. Passwords cannot be modified from the other accounts.

Changing Access Passwords

To change the password for any level of access, follow this procedure:

1. **Connect to the switch and log in as the administrator. The default password for the administrator is `admin`, unless changed using this procedure.**

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 configuration menu
      udef    - User-defined defaults menu
      port    - Port configuration menu
      vlan    - VLAN configuration menu
      stp     - Spanning Tree configuration menu
      snmp    - SNMP configuration menu
      trunk   - Trunk Group configuration menu
      ip      - IP configuration menu
      slb     - Layer 4 configuration menu
      vrrp    - VRRP configuration menu
      setup   - Step by step configuration set up
      dump    - Dump current configuration to script file
      putcfg  - Backup current configuration to TFTP server
      getcfg  - Restore current configuration from TFTP server
```

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
      idle    - Set timeout for idle CLI sessions
      bootp   - Enables or disables use of BOOTP
      snmp    - Enables or disables SNMP management access
      Web     - Enables or disables Web management access
      wport   - Set Web server port number
      mnet    - Set management network
      mmask   - Set management netmask
      banner  - Set login banner
      usrpw   - Set user password
      admpw   - Set administrator password
      l4apw   - Set L4 administrator password
      l4opw   - Set L4 operator password
      current - Display current system-wide configuration
```

4. **Select the password you wish to change (`usrpw`, `admpw`, `l4apw`, `l4opw`) at the `System` prompt. For example, change the administrator password as follows:**

```
System# admpw
```

5. **No matter which level of password you are changing, the switch will prompt you to enter the current administrator password to confirm your authority.**

```
Changing ADMINISTRATOR password; validation required...
Enter current administrator password:
```

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

6. **Next, enter the new password for the access level being changed.**

```
Enter new administrator password:
```

7. **Next, enter the new password a second time for confirmation.**

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

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

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

CHAPTER 3

Menu Basics

The switch's Command Line Interface (CLI) is used for viewing switch information and statistics. In addition, an 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 submenus. Each menu displays a list of commands and/or submenus that are available and 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 as shown below. Some menu items are not available under other access levels.

```
[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]
      revert  - Revert pending or applied changes [global command]
      exit    - Exit [global command, always available]
```

Menu Summary

- **Information Menu**

  This menu provides submenus for displaying information about the current status of the switch—from basic system settings, BGP, IP Routing, VLANs, Layer 4 settings, and more.

- **Statistics Menu**

  This menu provides submenus for displaying switch performance statistics. Included are port, IF, IP, ICMP, TCP, UDP, SNMP, routing, ARP, DNS, VRRP, Server SLB, Open Shortest Path First (OSPF), and Layer 4 statistics.

- **Configuration Menu**

  This menu is available only from an administrator login. It includes submenus for configuring every aspect of the switch. Changes to configuration are not active until explicitly applied. Changes can be saved to nonvolatile FLASH 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 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. It is used for upgrading switch software, 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, enabling you to generate a dump of the critical state information in the switch and to clear entries in the forwarding database and the ARP and routing tables.

Global Commands

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

Table 3-1 Global Commands

Command	Action
? <*command*>	Provide 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 nonvolatile 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 host name 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 host names (see "Domain Name System Configuration" on page 206).
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*>]] where *address* is the host name 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 to wait for the response. As with `ping`, the DNS parameters must be configured if specifying host names.
pwd	Display the command path used to reach the current menu.

Table 3-1 Global Commands (Continued)

Command	Action
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 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 History and Editing

Using the command line interface, you can retrieve and modify previously entered commands with just a few keystrokes. The following options are available globally at the command line:

Command Line Interface Shortcuts

Command Stacking

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
```

Command Abbreviation

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

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

Tab Completion

By entering the first letter of a command at any menu prompt and hitting <Tab>, the CLI will display all commands or options in that menu that begin with that letter. Entering additional letters will further refine the list of commands or options displayed. If only one command fits the input text when <Tab> is pressed, that command will be supplied on the command line which can then be applied. If the <Tab> key is pressed without any input on the command line, the currently active menu will be displayed.

CHAPTER 4

The Information Menu

This chapter discusses how to use the Command Line Interface (CLI) to display switch information. You can view configuration information for the switch when logged-in using most access levels.

/info
Information Menu

```
[Information Menu]
      sys     - Show system information
      env     - Show temperature and fan status
      log     - Show last 10 syslog messages
      link    - Show link status
      port    - Show port status
      slot    - Show slot status
      vlan    - Show VLAN information
      stp     - Show STP information
      trunk   - Show Trunk Group information
      arp     - ARP information menu
      fdb     - Forwarding Database information menu
      ip      - IP information menu
      route   - IP routing information menu
      ospf    - OSPF routing information menu
      slb     - Layer 4 information menu
      vrrp    - Show VRRP information
      swkey   - Show enabled software features
      dump    - Dump all information
```

The information provided by each menu option is briefly described in Table 4-1 on page 48, followed by examples of each information screen.

Table 4-1 Information Menu Options (/info)

Command Syntax and Usage
sys Displays system information, including the following: ■ System date and time ■ Switch model name and number ■ Switch name and location ■ Time of last boot ■ Base 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 ■ Login banner, if one is configured To view a sample, see page 51.
env Displays the switch temperature and fan status. To view a sample, see page 52.
log Displays 10 most recent syslog messages. To view a sample, see page 53.
link Displays configuration information about each port, including: ■ Slot letter (A-D) ■ Port number ■ Port speed (10, 100, 10/100, or 1000) ■ Duplex mode (full duplex only—this parameter is fixed) ■ Flow control for transmit and receive (no flow control only—this parameter is fixed) ■ Link status (up or down) For more information, see page 54.
port Displays port status information, including: ■ Slot letter (A-D) ■ Port number ■ Whether the port forwards or discards tagged frames ■ Whether the port forwards or discards untagged frames ■ Whether the port forwards or discards priority-tagged frames ■ Whether Remote Monitoring (RMON) is enabled (e) or disabled (d) ■ Port VLAN ID (PVID) ■ Port name For more information, see page 55.
slot Displays the slot status for Line Card Module (LCM) and Switching Fabric Module (SFM) slots. For more information, see page 56.

Table 4-1 Information Menu Options (/info) (Continued)

Command Syntax and Usage
vlan Displays VLAN configuration information, including: ■ VLAN Number ■ VLAN Name ■ Status ■ Port membership of the VLAN For more information, see page 58.
stp Displays Spanning Tree Protocol (STP) status and the following STP bridge information: ■ Priority ■ Hello interval ■ Maximum age value ■ Forwarding delay ■ Aging time Also displays the following port-specific STP information: ■ Port number and priority ■ Cost ■ State For more information, see page 59.
trunk When trunk groups are configured, you can view the state of each port in the various trunk groups. For more information, see page 61.
arp Displays the Address Resolution Protocol (ARP) Information Menu. For more information, see page 62.
fdb Displays the Forwarding Database Information Menu. For more information, see page 64.
ip Displays the IP Information Menu. IP information includes: ■ IP interface information: Interface number, IP address, subnet mask, broadcast address, VLAN number, class of service, 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 ■ Border Gateway Protocol status For more information, see page 66.

Table 4-1 Information Menu Options (/info) (Continued)

Command Syntax and Usage
route Displays the IP Routing Menu. Using the options of this menu, the system 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 Routing Information Protocol (RIP) tagged routes, specifying the number of hops to the destination (1-15 hops or 16 for infinite hops) ■ Metric for BGP, showing Multi-Exit Discriminator (MED) ■ The IP interface that the route uses For more information, see page 71.
ospf Displays the Open Shortest Path First (OSPF) Information menu, used for gathering information about OSPF configuration and performance. For menu options, see page 74.
slb Displays the Server Load Balancing Information Menu. For more information, see page 80.
vrrp Displays Virtual Router Redundancy Protocol (VRRP) information. For more information, see page 83.
swkey Displays a list of all the optional software packages that have been activated or installed on your switch. For more information, see page 84.
dump Dumps all switch information available from the Information Menu (10K or more, depending on your configuration). If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump commands. For more information, see page 84.

/info/sys
System Information

```
System Information at 17:20:30 Wed Apr 17, 2001

Alteon 708
sysName:     Main-switch
sysLocation: IT Lab
Last boot: 11:11:56 Fri Dec 22, 2000 (watchdog timer)

Switch base MAC address: 00:22:33:44:55:00
Interface 1 MAC address: 00:22:33:44:55:05  IP address:
123.255.224.10
Hardware Revision: A
Hardware Part No: 210019P1
Software Version 7.2.14 from FLASH image2, active configuration.
```

System information includes the following:

- System date and time
- Switch model name and number
- Switch name and location
- Time of last boot
- Base MAC address of the switch management processor
- MAC address and IP address of IP interface #1
- Hardware revision and part number
- Software version number and image file
- Configuration name
- Login banner, if one is configured

/info/env
Environment Information

```
Thermal MP:     36C

Thermal Fan:    20C

Thermal SFM  1: 26C
             2: 24C

Thermal LCM  1: 33.5C
             2: 37C

Power 1: Present
      2: Not Present

Fan: Present @ approx. 2050 RPM
```

Environment information includes:

- Temperature of the thermal MP, fan, switch fabric modules, line card modules
- Presence of power supply modules
- Presence and speed of fans in the fan tray

/info/log
Show Last 10 Syslog Messages

```
May 17 10:22:45 NOTICE  IP: default gateway 123.255.224.101 enabled
May 17 10:42:41 NOTICE  telnet1: admin login from host
123.255.224.230
May 17 10:43:50 NOTICE  telnet1: administrator password changed
May 17 10:43:50 INFO    telnet1: new configuration applied
May 17 10:44:01 INFO    telnet1: new configuration saved
May 17 10:50:50 NOTICE  telnet1: admin logout
May 17 13:24:34 NOTICE  telnet1: admin login from host
123.255.224.123
May 17 13:33:52 NOTICE  telnet1: admin connection closed
May 17 13:34:17 NOTICE  telnet1: admin login from host
123.255.224.123
```

Each syslog message has a criticality level associated with it, included in text form as a prefix to the log message. One of eight different prefixes is used, as shown below:

Table 4-2 Syslog Message Conditions (/info/log)

Parameter	Description
EMERG	Indicates the system is unusable
ALERT	Indicates action should be taken immediately
CRIT	Indicates critical conditions
ERR	Indicates error conditions or errored operations
WARNING	Indicates warning conditions
NOTICE	Indicates a normal but significant condition
INFO	Indicates an information message
DEBUG	Indicates a debug-level message

/info/link
Link Status Information

```
Slot  Port     Speed         Duplex      Flow Ctrl      Link
----  ----  ------------   --------   --TX-----RX--   ------
  A     1        100         full        no     no       up
        2      10/100        full        no     no      down
        3      10/100        full        no     no      down
        4      10/100        full        no     no      down
        5      10/100        full        no     no      down
        6      10/100        full        no     no      down
        7      10/100        full        no     no      down
        8      10/100        full        no     no      down
        9      10/100        full        no     no      down
       10      10/100        full        no     no      down
       11      10/100        full        no     no       up
       12      10/100        full        no     no      down
       13      10/100        full        no     no      down
       14      10/100        full        no     no      down
       15      10/100        full        no     no      down
       16      10/100        full        no     no      down
  B     1       1000         full        no     no      down
        2       1000         full        no     no      down
        3       1000         full        no     no      down
        4       1000         full        no     no       up
        :
```

Use this command to display link status information about each port on an Alteon 708 switch slot.

Table 4-3 Link Status Information Fields

Parameter	Description
`Slot`	LCM Slot (A-D)
`Port`	Port number (1-4 for Gigabit-SX, 1-16 for Fast Ethernet, 1-4 for 1000Base-T)
`Speed`	Port speed (10, 100, 10/100, or 10/100/1000)
`Duplex`	Duplex mode (full duplex only—this parameter is fixed)
`Flow Ctrl`	Flow control for transmit (Tx) and receive (Rx) (this paraemeter is fixed to “no” flow control only)
`Link`	Link status (up or down)

/info/port
Port Information

```
Slot  Port  Tag   UnTag  PriTag  RMON            VLAN(s)
----  ----  ----  -----  ------  ----  ------------------------------
  A     1   disc   frwd   frwd     d      1
        2   disc   frwd   frwd     d      1
        3   disc   frwd   frwd     d      1
        4   disc   frwd   frwd     d      1
  B     1   disc   frwd   frwd     d      1
        2   disc   frwd   frwd     d      1
        3   disc   frwd   frwd     d      1
        4   disc   frwd   frwd     d      1
        5   disc   frwd   frwd     d      1
        6   disc   frwd   frwd     d      1
        7   disc   frwd   frwd     d      1
        8   disc   frwd   frwd     d      1
        9   disc   frwd   frwd     d      1
       10   disc   frwd   frwd     d      1
       11   disc   frwd   frwd     d      1
       12   disc   frwd   frwd     d      1
       13   disc   frwd   frwd     d      1
       14   disc   frwd   frwd     d      1
       15   disc   frwd   frwd     d      1
       16   disc   frwd   frwd     d      1
        :
```

Port information includes the following:

Table 4-4 Port Information Fields

Parameter	Description
`Slot`	LCM slot (A-D).
`Port`	Port number (1-4 for Gigabit-SX, 1-16 for Fast Ethernet, 1-4 for 1000Base-T).
`Tag`	Shows whether tagged frames are discarded (disc) or forwarded (frwd).
`Untag`	Shows whether untagged frames are discarded (`disc`) or forwarded (`frwd`).
`PriTag`	Priority tagged frames are those with a tag on the priority part of a tag, not the VLAN part. This shows whether priority tagged frames are discarded (`disc`) or forwarded (`frwd`).
`RMON`	RMON status: d = disabled, e = enabled.
`VLAN`	Indicates VLAN number(s) to which the port belongs.

/info/slot
Slot Status Information

```
 >> Information# slot
------------------------------------------------------------------
 Slot    State         LCM type          SP States
-----  ---------  ------------------  ---------------
  A     Running    GIGABIT-SX         SP x0:  Running
                                      SP x1:  Running
                                      SP x2:  Running
                                      SP x3:  Running
  B     Running    GIGABIT-SX         SP x4:  Running
                                      SP x5:  Running
                                      SP x6:  Running
                                      SP x7:  Running
  C     Running    FAST-ETHERNET      SP x8:  Running
                                      SP x9:  Running
  D     Running    FAST-ETHERNET      SP x12: Running
                                      SP x13: Running

SFM slot    Usage     LCM slots
--------   -------   -----------
    1         GE       A B
    2         FE           C D
```

Slot status shows the devices inserted in the Line Card Module (LCM) slots and the usage of the Switching Fabric Module (SFM) slots. Slot status information includes the following:

Table 4-5 Slot Status Information Fields

Parameter	Description
`Slot`	LCM slot (A-D).
`State`	Shows the status of the LCM.
`LCM type`	Gigabit-SX, Fast Ethernet, 1000Base-T.
`SP states`	Shows the states of the switch processors (SPs) that are in use. For information about SP index numbers, see "Interpreting SP Index Numbers" below.
`SFM slots`	Shows the SFM slot number(s).
`Usage`	Shows the type of LCM that is using the SFM. SFM types are GE (Gigabit-SX or 1000Base-T), or FE (Fast Ethernet).
`LCM slots`	Shows the slot letter of the LCM using the SFM.

Interpreting SP Index Numbers

Switch processors (SP) are numbered by SP Index, x0 through x15.

- On the Alteon 708 SPs are assigned to slots as follows:
  - slot A= x0 through x3
  - slot B=x4 through x7
  - slot C=x8 through x11
  - slot D=x12 through x15.
- Each LCM slot is assigned four SP index numbers, regardless of the actual number of SPs present on the *installed* LCM:
  - A 4-port Gigabit-SX or 1000Base-T LCM requires four SPs; therefore, each SP index number corresponds to a port number. The `/info/slot` command will show four SP indexes for each LCM slot.
  - An 8-port Fast Ethernet LCM requires only two SPs; each SP controls two ports. Therefore, even though there are four possible SP indexes assigned to each LCM slot, the `/info/slot` command will show only two SP index numbers for a Fast Ethernet LCM slot.

/info/vlan
VLAN Information

```
VLAN                Name                Status   Ports
---- -------------------------------- ------   --------------
1     Default VLAN                       ena    a1 b12 c4-c9
2     VLAN 2                             ena    d3
```

This information display includes all configured VLANs and all member ports that have an active link state.

VLAN information includes:

- VLAN number
- VLAN name
- Status: enabled (`ena`), disabled (`dis`)
- Port membership of the VLAN. Port membership is represented as *slot* (letter) and *port* (number) together, as in “b12.”

/info/stp
Spanning Tree Information

```
Spanning Tree Group 1: On

Current Root:            Path-Cost  Port Hello MaxAge FwdDel Aging
 7fff 00:22:33:44:55:00        4      A1    2     20     15     15

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

Slot  Port  VLAN  Priority   Cost      State
----  ----  ----  --------   ----    ----------
  A     1     1       128      1    FORWARDING
        2     1       128      0     DISABLED
        3     1       128      0     DISABLED
        4     1       128      0     DISABLED
  B     1     1       128      0     DISABLED
        2     1       128      0     DISABLED
        3     1       128      0     DISABLED
        4     1       128      0     DISABLED
        5     1       128      0     DISABLED
        6     1       128     10    FORWARDING
        7     1       128      0     DISABLED
        8     1       128      0     DISABLED
        9     1       128      0     DISABLED
       10     1       128      0     DISABLED
       11     1       128      0     DISABLED
       12     1       128    100    FORWARDING
       13     1       128      0     DISABLED
       14     1       128      0     DISABLED
       15     1       128      0     DISABLED
       16     1       128      0     DISABLED
        :
        :
```

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:

- Slot number
- Port number
- Priority
- VLAN
- Cost
- State

The following table describes the Spanning Tree Protocol (STP) parameters.

Table 4-6 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`.

/info/trunk
Trunk Group Information

```
 Group    Slot  Port  STG     State
-------   ----  ----  ---  -----------
  1         A     1     1     DOWN
  1         A     2     1     DOWN
  1         A     3     1     DOWN
  1         A     4     1     DOWN
  2         B     1     1     forwarding
  2         B     2     1     forwarding
  2         B     3     1     forwarding
  2         B     4     1     forwarding
```

When trunk groups are configured, you can view the state of each port in the various trunk groups.

NOTE – If Spanning Tree Protocol on any port in the trunk group is set to `forwarding`, all active ports in the trunk group will have the same state.

/info/arp
Address Resolution Protocol Information

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

The ARP information includes the following:

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

Table 4-7 ARP Information Menu Options (/info/arp)

Command Syntax and Usage
find *<IP address, such as 192.4.17.101>* Displays a single ARP entry by IP address.
port *<port—must be slot and port together, such as b12>* Displays the ARP entries on a single port.
refsp *<port—must be slot and port together, such as b12>* Displays the ARP entries referenced by a single switch processor (SP).
vlan *<VLAN number (1-4094)>* Displays the ARP entries on a single VLAN.
dump Displays all ARP entries, including: ■ IP address and MAC address of each entry ■ Address status flag (see below) ■ The VLAN and port to which the address belongs ■ The ports that have referenced the address (empty if no port has routed traffic to the IP address shown) For more information, see page 63.

Table 4-7 ARP Information Menu Options (/info/arp)

Command Syntax and Usage
addr Displays the ARP address list, which is the list of IP addresses that receive responses to ARP requests. For more information, see page 63.

/info/arp/dump

Show All ARP Entry Information

```
  IP address     Flags    MAC address     VLAN   Port  Referenced SPs
--------------- ----- ----------------- ----  ------ -------------
123.255.224.101        00:11:22:33:44:b0    1    A1  1 2 5-8
123.255.224.104        00:11:22:33:44:ee    1    A1  empty
123.255.224.111        00:11:22:33:44:40    1    A1  empty
123.255.224.230        00:11:22:33:44:8d    1    B12 empty
123.255.224.240        00:11:22:33:44:fd    1    A1  empty
```

`Flag` field is interpreted as follows:

Table 4-8 ARP Flag Parameter

Flag	Description
U	Unresolved ARP entry. The MAC address has not been learned.

/info/arp/addr

ARP Address List Information

```
   IP address        IP mask        MAC address    VLAN Flags
  --------------- --------------- ----------------- ---- -----
  123.255.224.66   255.255.255.255 00:70:cf:03:20:04        P
  123.255.224.1    255.255.255.255 00:70:cf:03:20:06    1
  123.255.224.64   255.255.255.255 00:70:cf:03:20:05    1
```

/info/fdb
FDB Information

```
[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
      dump    - Show all FDB entries
      count   - Show count of FDB entries
```

The Forwarding Database (FDB) contains information that maps the Media Access Control (MAC) address and VLAN address of each known device to the switch port where the device address was learned. The FDB also shows which other ports have seen frames destined for a particular MAC address and VLAN.

Table 4-9 FDB Information Menu Options (/info/fdb)

Command Syntax and Usage
find *<MAC address>* **[***VLAN number (1-4094)***]** Displays a single database entry by its MAC address and, optionally, its VLAN. You are prompted to enter the MAC address of the device. Enter the MAC address using the hex format, such as `08:00:20:12:34:56`. You can also enter the MAC address using the format `xxxxxxxxxxxx`, such as `080020123456`.
port *<port defined as slot and port, such as b12>* Displays all FDB entries for a particular port.
vlan *<1-4094>* Displays all FDB entries on the selected VLAN.
dump Displays all entries in the FDB. For more information, see page 65.
count Displays the number of FDB entries.

/info/fdb/dump

Show All FDB Information

```
   MAC Address      VLAN   Port   State
00:11:22:33:44:45     1     A1     FWD
00:11:22:33:44:7a     1     A1     FWD
00:11:22:33:44:3d     1     A1     FWD
00:11:22:33:44:a4     1     A1     FWD
00:11:22:33:44:24     1     A1     FWD
00:11:22:33:44:4c     1     A1     FWD
00:11:22:33:44:b0     1     A1     FWD
00:11:22:33:44:40     1     A1     FWD
00:11:22:33:44:00     1     A1     FWD
00:11:22:33:44:05     1     A1     FWD
00:11:22:33:44:10     1     A1     FWD
00:11:22:33:44:28     1     A1     FWD
00:11:22:33:44:c1     1     B6     FWD
00:11:22:33:44:15     1     A1     FWD
00:11:22:33:44:5b     1     A1     FWD
      :
```

The following port states may be seen:

Table 4-10 FDB Information Parameters (/info/fdb/dump)

Parameter	Description
FWD	The address is in the *forwarding* state, meaning that it has been learned by the switch.
TRK	The address is in the *trunking* state. The port field represents the trunk group number.
UNK	The address is in the *unknown* state, meaning 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 ports*.
IF	The entry is for an *IP interface*. The MAC address is for a standard VRRP virtual router.
VIP	The entry is for a *virtual server*. The MAC address is for a virtual server router—a virtual router with the same IP address as a virtual server.

Clearing Entries from the Forwarding Database

To delete a MAC address from the FDB or to clear the entire FDB, refer to "Forwarding Database Options" on page 275.

/info/ip
Display IP Information

```
[IP Menu]
      bgp     - BGP information menu
      dump    - Show all IP information
```

Table 4-11 IP Information Menu Options (/info/ip)

Command Syntax and Usage
bgp Shows the Border Gateway Protocol menu. To view the menu options, see page 66.
dump Shows all IP information. To view an example, see "/info/ip/dump" on page 69.

/info/ip/bgp
Border Gateway Protocol Information

```
[BGP Menu]
      peer    - Show all BGP peers
      summary - Show all BGP peers in summary
      dump    - Show all BGP routes
```

Table 4-12 Border Gateway Protocol Menu Options (/info/ip/bgp)

Command Syntax and Usage
peer Shows all Border Gateway Protocol (BGP) peers. To view an example, see page 67.
summary Shows all BGP peers in summary. To view an example, see page 67.
dump Shows all BGP routes. To view an example, see page 68.

/info/ip/bgp/peer

BGP Peer information

Following is an example of the information that `/info/ip/bgp/peer` provides.

```
BGP Peer Information:
  1: 205.178.23.142  , version 4, TTL 1
    Remote AS: 142, Local AS: 101, Link type: EBGP
    Remote router ID: 205.178.23.142,    Local router ID:
     205.178.23.101
    BGP status: established, Old status: established
    Total received packets: 114, Total sent packets: 121
    Received updates: 1, Sent updates: 8
    Keepalive: 30, Holdtime: 90, MinAdvTime: 30
    LastErrorCode: cease(6), LastErrorSubcode: unspecified(0)

  2: 205.178.15.148  , version 0, TTL 2
    Remote AS: 148, Local AS: 101, Link type: EBGP
    Remote router ID: 0.0.0.0,    Local router ID: 205.178.23.101
    BGP status: connect, Old status: connect
    Total received packets: 0, Total sent packets: 0
    Received updates: 1, Sent updates: 8
    Keepalive: 0, Holdtime: 0, MinAdvTime: 30
    LastErrorCode: unknown(0), LastErrorSubcode: unspecified(0)
```

/info/ip/bgp/summary

BGP Summary information

Following is an example of the information that `/info/ip/bgp/summary` provides.

```
   BGP Peer Summary Information:
        Peer          V    AS     MsgRcvd  MsgSent Up/Down   State
     --------------- - -------- -------- -------- -------- ----------
  1: 205.178.23.142  4      142      113      121 00:00:28 established
  2: 205.178.15.148  0      148        0        0 never     connect
```

/info/ip/bgp/dump

Dump BGP Information

Following is an example of the information that `/info/ip/bgp/dump` provides.

```
 >> BGP# dump
Status codes: * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
    Network          Next Hop        Metr  LcPrf Wght  Path
   --------------- ---------------  ----- ---- ----- --------------
*> 10.0.0.0         205.178.21.147      1        256 147 148 i
*>i205.178.15.0     0.0.0.0                      0 i
*                   205.178.21.147      1        128 147 i
*> 205.178.17.0     205.178.21.147      1        128 147 i
   13.0.0.0         205.178.21.147      1        256 147 {35} ?

The 13.0.0.0 is filtered out by rrmap; or, a loop detected.
```

/info/ip/dump

IP Information

Following is an example of IP information provided by `/info/ip/dump`.

```
IP information:

Interface information:
   1: 172.25.1.28,     255.255.0.0,     172.25.255.255,  vlan 1, up
   2: 172.25.1.11,     255.255.0.0,     172.25.255.255,  vlan 1, DIS-
ABLED

Default gateway information: metric strict
  1: 172.25.1.1,      up

Current forwarding setting: ON, dirbr disabled

Current port settings: all ports ON
Current network filter settings:
  none

Current route map settings:

Current OSPF settings: ON
  Default route none
  lsdb limit 0
  enabled traps 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15,
16,

Current OSPF area settings:
  0: 0.0.0.0,         type transit, auth none, metric 1, spf 10,
enabled
  1: 0.0.0.1,         type stub, auth none, metric 1, spf 10, enabled

Current OSPF interface settings:
  1: 172.25.1.28,     area index  0, prio 5, cost 1, enabled
    hello 10, dead 40, trans 1, retra 5
    key
  2: 172.25.1.11,     area index  1, prio 5, cost 1, enabled
    hello 10, dead 40, trans 1, retra 5
    key
```

IP information includes:

- Router ID, IP address, Autonomous System (AS) number
- 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
- Current forwarding status
- Current port settings
- Current Routing Information Protocol (RIP) settings enable status, update period, and the following active modes

Table 4-13 RIP Settings (/info/ip/dump)

Mode	Description
`listen`	The switch learns routes from other routers.
`supply`	The switch sends routes to other routers.
`default`	The switch accepts Routing Information Protocol default routes from other routers and gives them priority over configured default gateways.
`poison`	The switch uses split horizon with poisoned reverse.
`trigg`	The switch supplies routes when the routing table has been updated. The table is updated when routes are added or removed.
`mcast`	RIPv2 compatible, send route advertisements via multicast. In multicast, information is supplied only to specific addresses.
`metric`	The hop count the switch delivers to the router. The hop count is the number of hops to the final destination.
`auth`	Authorization.

/info/route
IP Routing Information

```
[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
```

Using the commands listed below, you can display all or a portion of the IP routes currently held in the switch.

Table 4-14 Route Information Menu Options (/info/route)

Command Syntax and Usage
find *<IP address, such as 192.4.17.101>* Displays a single route by destination IP address.
gw *<default gateway address, such as 192.4.17.44>* Displays routes to a single gateway.
type indirect|direct|local|broadcast|martian|multicast Displays routes of a single type. For more information, see Table 4-15 on page 72.
tag fixed|static|snmp|addr|rip|ospf|bgp|icmp|broadcast|martian|multicast|dynamic Displays routes of a single tag. For more information, see Table 4-16 on page 73.
if *<(1-1024)>* Displays routes on the selected interface.
dump Displays all routes known to the switch. For more information, see page 72.

/info/route/dump

Show All IP Route Information

Following is an example of the information that `/info/route/dump` provides.

```
  Destination         Mask          Gateway        Type      Tag     If Metr
------------- --------------- --------------- --------- --------- --- --
0.0.0.0       0.0.0.0         123.178.13.101  indirect  static     1
10.0.0.0      255.0.0.0       123.178.13.111  indirect  rip        1    2
127.0.0.0     255.0.0.0       0.0.0.0         martian   martian
123.17.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    2
123.18.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    4
123.19.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    2
123.20.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    2
123.21.0.0    255.255.0.0     123.21.1.1      direct    fixed      2
123.21.1.1    255.255.255.    255 123.21.1.1  local     addr       2
123.21.255.255 255.255.255.255 123.21.255.255 broadcast broadcast  2
123.23.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    3
123.24.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    3
123.25.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    4
123.26.0.0    255.255.0.0     123.178.13.39   indirect  rip        1    2
123.27.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    5
123.28.0.0    255.255.0.0     123.178.13.101  indirect  rip        1    2
123.29.0.0    255.255.0.0     123.178.13.111  indirect  rip        1    5
```

Table 4-15 describes the routing `Type` parameters.

Table 4-15 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 that is filtered out. Packets sent to this destination are discarded.
`multicast`	Indicates a multicast route.

Table 4-16 describes the routing `Tag` parameters.

Table 4-16 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`	The 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).
`ospf`	The address was learned by Open Shortest Path First (OSPF).
`bgp`	The address was learned by the Border Gateway Protocol (BGP).
`broadcast`	Indicates a broadcast address.
`martian`	The address belongs to a filtered group.
`multicast`	Indicates a multicast address.

/info/ospf
OSPF Information Menu

```
[OSPF Information Menu]
      general - Show general information
      aindex  - Show area(s) information
      if      - Show interface(s) information
      virtual - Show details of virtual links
      nbr     - Show neighbor(s) information
      dbase   - Database Menu
      reqlist - Show request list of this neighbor
      retlist - Show retransmission list of this neighbor
      sumaddr - Show summary address list
      nsumadd - Show NSAA summary address list
      routes  - Show OSPF routes
```

This menu allows you to display information about Open Shortest Path First (OSPF).

Table 4-17 OSPF Information Options (/info/ospf)

Command Syntax and Usage
general Displays general Open Shortest Path First (OSPF) information. To view an example, see page 75.
aindex *<area index (0-7)>* Displays area information. To view an example, see page 75.
if **<***interface number (1-1024)>* Displays area information. To view an example, see page 76.
virtual Displays virtual link information.
nbr *<neighbor router id (A.B.C.D.)>* Displays neighbor information. To view an example, see page 76.
dbase *<interface number (1-1024)>* Displays the Database Menu. To view the menu options, see page 78.
reqlist *<neighbor router id (A.B.C.D)>* Displays the request list of this neighbor.
retlist *<neighbor router id (A.B.C.D)>* Shows the retransmission list of this neighbor. To view an example, see page 77.

Table 4-17 OSPF Information Options (/info/ospf) (Continued)

Command Syntax and Usage
sumaddr *<area index (0-7)>* Displays the summary address list
nsumadd *<area index (0-7)>* Displays the NSAA summary address list.
routes Displays the OSPF routes. To view an example, see page 77.

/info/ospf/general

General OSPF Information.

```
>> OSPF Information# general
OSPF is disabled.
```

/info/ospf/aindex <*index*>

OSPF Area Information

NOTE – The `index number` (0-7) represents the OSPF area that you wish to display.

```
Area (index) 1 with ID 0.0.0.1
   Number of interfaces in this area is 2 of which 2 are UP
   Area authentication: none
   LS types accepted by this area 63
   SPF algorithm executed 2 times
   Number of LSDB entries of this area 10
   Count of ASBR local to this area 1
   Count of ABR local to this area 2
   Total neighbors are 2 of which
                2 are >=INIT state,
                2 are >=EXCH state,
                2 are =FULL state
```

/info/ospf/if <*interface number*>

OSPF Interface Information

NOTE – The *interface numbe*r (1-1024) represents the OSPF interface that you wish to display.

Following is an example of the information that `/info/ospf/if` provides.

```
Ip Address 10.10.10.1, Area 0.0.0.1, Admin Status UP
   Router ID 10.10.12.1, State Backup DR, Priority 5
   Designated Router (ID) 10.10.13.2, Ip Address 10.10.10.2
   Backup Designated Router (ID) 10.10.12.1, Ip Address 10.10.10.1
   Timer intervals, Hello 10, Dead 40, Wait 0, Retransmit 5,
   Poll interval 0, Transit delay 1
   Neighbor count is 1   If Event 8, Authentication type None
```

/info/ospf/nbr

OSPF Neighbor Information

```
Intf  NeighborID       Prio  State         Address
----  ----------       ----  -----         -------
   1  10.10.13.2          1  Full           10.10.10.2
   2  10.10.9.2           1  Full           10.10.9.2
```

/info/ospf/retlist

OSPF Retransmission List

```
> OSPF Information# retlist
Neighbor : 10.0.0.1
 Link ID       ADV Router      Options Age   Seq#         Checksum
 13.0.0.2       13.0.0.2        3       0    0     0x80000001 0x75CD
Neighbor : 13.0.0.3
 Link ID         ADV Router    Options Age   Seq#         Checksum
 13.0.0.2         13.0.0.2       3      0    0     0x80000001 0x75CD
Neighbor : 10.0.0.1
 Link ID         ADV Router    Options Age   Seq#         Checksum
 13.0.0.2         13.0.0.2       3      0    0     0x80000001  0x75CD
Neighbor : 13.0.0.3
 Link ID         ADV Router    Options Age   Seq#         Checksum
 13.0.0.2         13.0.0.2       3     0    0      0x80000001 0x75CD
Neighbor : 13.0.0.3
 Link ID         ADV Router     Options Age   Seq#         Checksum
 13.0.0.2         13.0.0.2       3      0    3600 0x80000002 0x7DC3
 13.0.0.2         13.0.0.2       1      2    0    0x80000007 0x9878
Neighbor : 13.0.0.3
 Link ID         ADV Router      Options Age   Seq#         Checksum
 13.0.0.2         13.0.0.2       3       0    0   0x80000003   0x71CF
```

/info/ospf/routes

OSPF Routes

```
Codes: IA - OSPF inter area,
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
```

/info/ospf/dbase

OSPF Database Menu

```
[OSPF Database Menu]
      advrtr  - LS Database info for an Advertising Router
      asbrsum - ASBR Summary LS Database info
      dbsumm  - LS Database summary
      ext     - External LS Database info
      nw      - Network LS Database info
      nssa    - NSSA External LS Database info
      rtr     - Router LS Database info
      self    - Self Originated LS Database info
      summ    - Network-Summary LS Database info
      all     - All
```

Table 4-18 OSPF Database Options (/info/ospf/dbase)

Command Syntax and Usage
advrtr **<***router ID (A.B.C.D.)***>** Displays the Link State (LS) Database information advertised by the router with router ID A.B.C.D.
asbrsum *<adv rtr (A.B.C.D.)>* | *<link state id (A.B.C.D.>* | *<self>* Displays the Autonomous System Boundary Router (ASBR) Summary LS Database information. For more information, see Note – on page 79.
dbsumm *<adv rtr (A.B.C.D.)>* | *<link state id (A.B.C.D.>* | *<self>* Displays the LS Database summary. For more information, see Note – on page 79.
ext *<adv rtr (A.B.C.D)>* | *<link state id (A.B.C.D.>* | *<self>* Displays the External LS Database information. For more information, see Note – on page 79.
nw *<adv rtr (A.B.C.D)>* | *<link state id (A.B.C.D.>* | *<self>* Displays the Network LS Database information. For more information, see Note – on page 79.
nssa *<adv rtr (A.B.C.D)>* | *<link state id (A.B.C.D.>* | *<self>* Displays the Not-So-Stubby Area (NSSA) External LS Database information. For more information, see Note – on page 79.
rtr *<adv rtr (A.B.C.D)>* | *<link state id (A.B.C.D.>* | *<self>* Displays the Router LS Database information. For more information, see Note – on page 79.
self *<adv rtr (A.B.C.D)>* | *<link state id (A.B.C.D.>* | *<self>* Displays the Self Originated LS Database information. For more information, see Note – on page 79.

Table 4-18 OSPF Database Options (/info/ospf/dbase) (Continued)

Command Syntax and Usage
summs Displays the Network Summary LS Database information.
all Displays a summary of all information. To view an example, see page 79.

NOTE – The Link State (LS) database information can be searched based on the following criteria:

1. **Advertising Router:** LS database information advertised by the router with router ID A.B.C.D.
2. **Link State ID:** LSAa with link state ID A.B.C.D.
3. **Self:** LS Database information advertised by this router (self-originated LSAs).

/info/ospf/dbase/all

Summary of All

```
                Router LSAs (Area 0.0.0.0)
  Link ID         ADV Router      Options Age   Seq#         Checksum
  0.0.0.3         0.0.0.3          0x2     1082  0x80000003   0x2D56
```

/info/slb
SLB Information Menu

```
[Server Load Balancing Information Menu]
      sess    - Session table information menu
      real    - Show real server information
      virt    - Show virtual server information
      filt    - Show redirect filter information
      port    - Show port information
      gslb    - Show GSLB information
      dump    - Show all layer 4 information
```

This menu provides access to Layer 4 information.

Table 4-19 Layer 4 Information Menu Options (/info/slb)

Command Syntax and Usage
sess Displays the Session Table Information Menu. To view menu options, see page 81.
real *<real server number (1-1024)>* Displays Real server number, real IP address, MAC address, VLAN, physical switch port, layer where health check is performed, and health check result.
virt *<virtual server number (1-1024)>* ■ Displays Virtual Server State: Virtual server number, virtual server IP address, virtual MAC address ■ Virtual Port State: Virtual service or port, server port mapping, real server group, group backup server
filt *<filter ID (1-1024)>* Displays the filter number, destination port, real server port, real server group, health check layer, group backup server, and real server group, IP address, backup server, and status.
port *<port defined as slot and port, such as b12>* Displays the physical port number, proxy IP address, filter status, a list of applied filters, and client and/or server Layer 4 activity.
gslb Displays remote switch number, remote switch IP address, IP subnet mask, and health status.
dump Displays all Layer 4 information for the switch. For details, see page 82.

/info/slb/sess

Show Session Table Information

```
[Session Table Information Menu]
      find    - Show all session entries with source IP address
      sp      - Show all session entries on SP
      dump    - Show all session entries
```

Table 4-20 Session Information Menu Options (/info/slb/sess)

Command Syntax and Usage
find *<IP address, such as 192.4.17.101>* Displays all session entries with source IP address.
sp *<SP as port such as b12, or index such as x5>* Displays all session entries on the Switch Processor (SP).
dump Displays all session entries.

/info/slb/dump

Show All Layer 4 Information

```
Real server state:
   1: 123.178.13.12, 00:60:cf:20:2f:cf, vlan 1, port B4, health 4, up
   2: 123.178.13.21, 00:60:cf:20:2f:be, vlan 1, port B3, health 4, up

Virtual server state:
   1: 123.178.13.1,    00:60:cf:50:ac:04
    virtual ports:
    http: rport http, group 1, backup none
        real servers:
           1: 123.178.13.12, backup none, up
           2: 123.178.13.21, backup none, up

Redirect filter state:

Port state:
 A1: client  enabled, server disabled
     pip 0.0.0.0, submac disabled
     filt disabled, filters: empty
 A2: client  enabled, server  enabled
     pip 0.0.0.0, submac disabled
     filt  enabled, filters: 1 2 224
 A3: client  enabled, server disabled
     pip 0.0.0.0, submac disabled
     filt disabled, filters: empty
 B3: client disabled, server  enabled
     pip 0.0.0.0, submac disabled
     filt disabled, filters: empty
 B4: client disabled, server  enabled
     pip 0.0.0.0, submac disabled
     filt disabled, filters: empty

Note: All undisplayed ports are configured with the following:
     client disabled, server disabled
     pip 0.0.0.0, submac disabled
     filt disabled, filters empty
```

/info/vrrp
VRRP Information

Virtual Router Redundancy Protocol (VRRP) support on Alteon WebSystems' switches provides redundancy between routers in a LAN. This redundancy is accomplished by configuring the same virtual router IP address and ID number on each participating VRRP-capable routing device. One of the virtual routers is then elected as the master, based on a number of priority criteria, and assumes control of the shared virtual router IP address. If the master fails, one of the backup virtual routers will assume routing authority and take control of the virtual router IP address.

```
VRRP information:
 1: vrid  1, 172.21.11.102,  if  1, renter, prio 120, master
 2: vrid  2, 37.21.1.1,      if  2, renter, prio 120, master
 3: vrid  3, 172.21.11.100,  if  1, renter, prio 120, master, server
```

When virtual routers are configured, you can view the status of each virtual router using this command. VRRP information includes:

- Virtual router number
- Virtual router ID and IP address
- Interface number
- Ownership status:
  - *Owner* identifies the preferred master virtual router. A virtual router is the owner when the IP address of the virtual router and its IP interface are the same.
  - *Renter* identifies virtual routers which are not owned by this device.
- Priority value: During the election process, the virtual router with the highest priority becomes master.
- Activity status:
  - *Master* identifies the elected master virtual router.
  - *Backup* identifies that the virtual router is in backup mode.
- Server status: The server state identifies virtual routers that support Layer 4 services. These are known as *virtual server routers*—any virtual router whose IP address is the same as any configured virtual server IP address.
- Proxy status: The proxy state identifies *virtual proxy routers*, where the virtual router shares the same IP address as a proxy IP address. The use of virtual proxy routers enables redundant switches to share the same IP address, minimizing the number of unique IP addresses that must be configured.

/info/swkey
Software Enabled Keys

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

```
>> Information# swkey
Enabled Software features:  Layer 4: SLB + WCR
```

Software key information includes a list of all the optional software packages which have been activated or installed on your switch.

info/dump
Information Dump

Use the dump command to dump all switch information available from the Information Menu (10K or more, depending on your configuration). The system will ask you to confirm dumping all information. Enter **y** to confirm, **n** to cancel.

```
>> Information# dump
Confirm dumping all information [y|n]: y
```

This data is useful for tuning and debugging switch performance.

If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump commands.

CHAPTER 5

The Statistics Menu

This chapter discusses how to use the Command Line Interface (CLI) to display switch statistics. You can view the statistics in both the user and Administrator command modes.

/stats
Statistics Menu

```
[Statistics Menu]
      port    - Statistics Menu for one port
      slb     - Layer 4 Statistics Menu
      mp      - Management Processor Statistics Menu
      sp      - Switch Processor Statistics Menu
      dump    - Dump all statistics
```

The statistics provided by each menu option is briefly described in Table 5-1, followed by examples of each information screen.

Table 5-1 Statistics Menu Options (/stats)

Command Syntax and Usage
port *<port as slot and port, such as b12>* Displays the Port Statistics Menu for the specified port. Use this command to display traffic statistics on a port-by-port basis. Traffic statistics include SNMP Management Information Base (MIB) objects. To view menu options, see page 87.
slb Displays the Server Load Balancing (SLB) Statistics Menu. To view menu options, see page 99.
mp Displays the Management Processor Statistics Menu. Use this command to view information on how switch management processes and resources are currently being allocated. To view menu options, see page 112.

Table 5-1 Statistics Menu Options (/stats) (Continued)

Command Syntax and Usage
sp *<port as slot and port, such as b12>* Displays the Switch Processor Statistics Menu. To view menu options, see page 136.
dump After you enter **y** to confirm, this command dumps all switch statistics. Use this command to gather data for tuning and debugging switch performance. If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump command. To view menu options, see page 142.

/stats/port <slot port>
Port Statistics Menu

This menu displays traffic statistics on a port-by-port basis. Traffic statistics include SNMP Management Information Base (MIB) objects.

```
[Port Statistics Menu]
      brg     - Bridging ("dot1") statistics
      ether   - Ethernet ("dot3") statistics
      if      - Interface ("if") statistics
      link    - Link statistics
      rmon    - RMON statistics
      clear   - Clear statistics
```

Table 5-2 Port Statistics Menu Options (/stats/port)

Command Syntax and Usage
brg Displays bridging IEEE 802.1 (“dot1”) statistics for the port. For more information, see page 88.
ether Displays Ethernet IEEE 802.3 (“dot3”) statistics for the port. For more information, see page 89.
if Displays interface statistics for the port. For an example, see page 93.
link Displays link statistics for the port. For an example, see page 95.
rmon Displays Remote Monitoring (RMON) statistics for the port—Etherstats, History, and Alarms and Events. For more information, see page 95
clear Clears all switch statistics for the port.

/stats/port <slot port>/brg

Bridging (*dot1*) Statistics

```
Bridging statistics for Port A1:
dot1PortInFrames:                           0
dot1PortOutFrames:                          0
dot1PortInDiscards:                     16833
dot1TpLearnedEntryDiscards:                 0
dot1BasePortDelayExceededDiscards:          0
dot1BasePortMtuExceededDiscards:            0
dot1StpPortForwardTransitions:              0
```

Table 5-3 Bridging *dot1* Statistics* (/stats/port/brg)

Statistic	Description
dot1PortInFrames	The number of frames that have been received by this port from its segment. Note that a frame received on the interface corresponding to this port is only counted by this object if and only if it is for a protocol being processed by the local bridging function, including bridge management frames.
dot1PortOutFrames	The number of frames that have been transmitted by this port to its segment. Note that a frame transmitted on the interface corresponding to this port is only counted by this object if, and only if, it is for a protocol being processed by the local bridging function, including bridge management frames.
dot1PortInDiscards	The count of valid frames received that were discarded (i.e., filtered) by the Forwarding Process.
dot1TpLearnedEntryDiscards	The total number of Forwarding Database entries, which have been or would have been learnt, but have been discarded due to a lack of storage space in the Forwarding Database. If this counter is increasing, it indicates that the Forwarding Database is regularly becoming full (a condition which has unpleasant performance effects on the subnetwork). If this counter has a significant value but is not presently increasing, it indicates that the problem has been occurring but is not persistent.
dot1BasePortDelayExceededDiscards	The number of frames discarded by this port due to excessive transit delay through the bridge. It is incremented by both transparent and source route bridges.

*** Extract from RFC-1493**

Table 5-3 Bridging *dot1* Statistics* (/stats/port/brg) (Continued)

Statistic	Description
dot1BasePortMtuExceed-edDiscards	The number of frames discarded by this port due to an excessive size. It is incremented by both transparent and source route bridges.
dot1StpPortForward-Transitions	The number of times this port has transitioned from the Learning state to the Forwarding state.

*** Extract from RFC-1493**

/stats/port <*slot-letter port-number*>/ether

Ethernet (dot3) Statistics

```
Ethernet statistics for Port A1:
dot3StatsAlignmentErrors:                  0
dot3StatsFCSErrors:                        0
dot3StatsSingleCollisionFrames:            0
dot3StatsMultipleCollisionFrames:          0
dot3StatsSQETestErrors:                    0
dot3StatsDeferredTransmissions:            0
dot3StatsLateCollisions:                   0
dot3StatsExcessiveCollisions:              0
dot3StatsInternalMacTransmitErrors:        0
dot3StatsCarrierSenseErrors:               0
dot3StatsFrameTooLongs:                    0
dot3StatsInternalMacReceiveErrors:         0
dot3CollFrequencies [1-15]:
```

Table 5-4 Ethernet (“dot3”) statistics* (/stats/port/ether)

Statistic	Description
dot3StatsAlignmentErrors	A count of frames received on a particular interface that are not an integral number of octets in length and do not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object is incremented when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtain are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.

*** Extract from RFC-1643**

Table 5-4 Ethernet (“dot3”) statistics* (/stats/port/ether) (Continued)

Statistic	Description
dot3StatsFCSErrors	A count of frames received on a particular interface that are an integral number of octets in length but do not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object is incremented when the frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtain are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
dot3StatsSingleCollisionFrames	A count of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts, and is not counted by the corresponding instance of the dot3StatsMultipleCollisionFrames object.
dot3StatsMultipleCollisionFrames	A count of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts, and is not counted by the corresponding instance of the dot3StatsSingleCollisionFrames object.
dot3StatsSQETestErrors	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.
dot3StatsDeferredTransmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.

*** Extract from RFC-1643**

Table 5-4 Ethernet (“dot3”) statistics* (/stats/port/ether) (Continued)

Statistic	Description
dot3StatsLateCollisions	The number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet. Five hundred and twelve bit-times corresponds to 51.2 micro-seconds on a 10 Mbit/s system. A (late) collision included in a count represented by an instance of this object is also consid-ered as a (generic) collision for purposes of other collision-related statistics.
dot3StatsExcessiveCol-lisions	A count of frames for which transmission on a particular inter-face fails due to excessive collisions.
dot3StatsInternalMac-TransmitErrors	A count of frames for which transmission on a particular inter-face fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object. The precise meaning of the count represented by an instance of this object is implementation- specific. In particular, an instance of this object may represent a count of transmission errors on a particular interface that are not otherwise counted.
dot3StatsCarrierSense-Errors	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a partic-ular interface. The count represented by an instance of this object is incre-mented at most once per transmission attempt, even if the car-rier sense condition fluctuates during a transmission attempt.
dot3StatsFrameTooLongs	A count of frames received on a particular interface that exceed the maximum permitted frame size. The count represented by an instance of this object is incre-mented when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtain are, according to the conventions of IEEE 802.3 Layer Management, counted exclu-sively according to the error status presented to the LLC.

*** Extract from RFC-1643**

Table 5-4 Ethernet (“dot3”) statistics* (/stats/port/ether) (Continued)

Statistic	Description
`dot3StatsInternalMac-ReceiveErrors`	A count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsFrameTooLongs object, the dot3StatsAlignmentErrors object, or the dot3StatsFCSErrors object. The precise meaning of the count represented by an instance of this object is implementation- specific. In particular, an instance of this object may represent a count of receive errors on a particular interface that are not otherwise counted.
`dot3CollFrequencies [1-15]`	A count of individual MAC frames for which the transmission (successful or otherwise) on a particular interface occurs after the frame has experienced exactly the number of collisions in the associated dot3CollCount object. For example, a frame which is transmitted on interface 77 after experiencing exactly 4 collisions would be indicated by incrementing only dot3CollFrequencies.77.4. No other instance of dot3CollFrequencies would be incremented in this example.

*** Extract from RFC-1643**

/stats/port <slot-letter port-number>/if

Interface (IF) Statistics

```
Interface statistics for Port A1:          ifHCIn Counters      ifH-
COut Counters
Octets:                      112090947             7296755
UcastPkts:                      350762               53493
BroadcastPkts:                  148133               10593
MulticastPkts:                   77805                   1
Discards:                        16833                   0
Errors:                              0                   0
ifInUnknownProtos:                   0
```

Table 5-5 Interface (IF) Statistics (/stats/port/if)

Statistic	Description
ifInOctets	The total number of octets received on the interface, including framing characters.
ifInUcastPkts	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were not addressed to a multicast or broadcast address at this sub-layer.
ifInBroadcastPkts	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a broadcast address at this sub-layer.
ifInMulticastPkts	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
ifInDiscards	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.
ifInErrors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.

*** Extract from RFC-1213**

Table 5-5 Interface (IF) Statistics (/stats/port/if) (Continued)

Statistic	Description
ifInUnknownProtos	For packet-oriented interfaces, the number of packets received via the interface which were discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces which support protocol multiplexing the number of transmission units received via the interface which were discarded because of an unknown or unsupported protocol. For any interface which does not support protocol multiplexing, this counter will always be 0.
ifOutOctets	The total number of octets transmitted out of the interface, including framing characters.
ifOutUcastPkts	The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.
ifOutBroadcastPkts	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a broadcast address at this sub-layer, including those that were discarded or not sent. This object is a 64-bit version of ifOutBroadcastPkts.
ifOutMulticastPkts	The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses. This object is a 64-bit version of ifOutMulticastPkts.
ifOutDiscards	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
ifOutErrors	For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.

*** Extract from RFC-1213**

/stats/port <slot-letter port-number>/link

Link Statistics

```
Link statistics for Port A1:
linkStateChange:             1
```

Table 5-6 Link Statistics (/stat/sport/link)

Statistic	Description
linkStateChange	The total number of changes in link state.

/stats/port <slot-port number>/rmon

RMON Statistics

```
RMON statistics for port A1:
etherStatsDropEvents:                      0
etherStatsOctets:                       2665
etherStatsPkts:                           33
etherStatsBroadcastPkts:                  19
etherStatsMulticastPkts:                   7
etherStatsCRCAlignErrors:                  0
etherStatsUndersizePkts:                   0
etherStatsOversizePkts:                    0
etherStatsFragments:                       0
etherStatsJabbers:                         0
etherStatsCollisions:                      0
etherStatsPkts64Octets:                   25
etherStatsPkts65to127Octets:               5
etherStatsPkts128to255Octets:              2
etherStatsPkts256to511Octets:              1
etherStatsPkts512to1023Octets:             0
etherStatsPkts1024to1518Octets:            0
```

Table 5-7 RMON Statistics* (/stats/port/rmon)

Statistic	Description
etherStatsDropEvents	The total number of events in which packets were dropped by the probe because of lack of resources. Note that this number is not necessarily the number of packets dropped; it is just the number of times this condition has been detected.

*** Extract from RFC-1757**

Table 5-7 RMON Statistics* (/stats/port/rmon) (Continued)

Statistic	Description
`etherStatsOctets`	Total number of data octets (including those in bad packets) received on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval. The differences in the sampled values are Pkts and Octets, respectively, and the number of seconds in the interval is Interval. These values are used to calculate the Utilization as follows: $$\text{Utilization} = \frac{\text{Pkts} \times (9.6 + 6.4) + (\text{Octets} \times 0.8)}{\text{Interval} \times 10,000}$$ The result of this equation is the value "Utilization," which is the percent utilization of the ethernet segment on a scale of 0 to 100 percent.
`etherStatsPkts`	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
`etherStatsBroadcastPkts`	The total number of good packets received that were directed to the broadcast address. Note that this does not include multicast packets.
`etherStatsMulticastPkts`	The total number of good packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
`etherStatsCRCAlignErrors`	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
`etherStatsUndersizePkts`	The total number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
`etherStatsOversizePkts`	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.

*** Extract from RFC-1757**

Table 5-7 RMON Statistics* (/stats/port/rmon) (Continued)

Statistic	Description
etherStatsFragments	The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that it is entirely normal for etherStatsFragments to increment. This is because it counts both runts (which are normal occurrences because of collisions) and noise hits.
etherStatsJabbers	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
etherStatsCollisions	The best estimate of the total number of collisions on this Ethernet segment. The value returned will depend on the location of the RMON probe. Section 8.2.1.3 (10BASE-5) and section 10.3.1.3 (10BASE-2) of IEEE standard 802.3 states that a station must detect a collision, in the receive mode, if three or more stations are transmitting simultaneously. A repeater port must detect a collision when two or more stations are transmitting simultaneously. Thus a probe placed on a repeater port could record more collisions than a probe connected to a station on the same segment would. Probe location plays a much smaller role when considering 10BASE-T. 14.2.1.4 (10BASE-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10BASE-T station can only detect collisions when it is transmitting. Thus probes placed on a station and a repeater, should report the same number of collisions. Note also that an RMON probe inside a repeater should ideally report collisions between the repeater and one or more other hosts (transmit collisions as defined by IEEE 802.3k) plus receiver collisions observed on any coax segments to which the repeater is connected.

*** Extract from RFC-1757**

Table 5-7 RMON Statistics* (/stats/port/rmon) (Continued)

Statistic	Description
etherStatsPkts64Octets	The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including FCS octets).
etherStatsPkts65to127-Octets	The total number of packets (including bad packets) received that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
etherStatsPkts128to255-Octets	The total number of packets (including bad packets) received that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
etherStatsPkts256to511-Octets	The total number of packets (including bad packets) received that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
etherStatsPkts512to1023-Octets	The total number of packets (including bad packets) received that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
etherStatsPkts1024to1518-Octets	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).

*** Extract from RFC-1757**

/stats/slb
Server Load Balancing Statistics

```
[Server Load Balancing Statistics Menu]
      real    - Real server statistics
      group   - Real server group statistics
      virt    - Virtual server statistics
      filt    - Filter statistics
      sp      - SLB statistics menu for one Switch Processor
      gslb    - Global SLB statistics
      maint   - Maintenance statistics
      clear   - Clear all SLB statistics
      dump    - Dump all SLB statistics
```

Table 5-8 SLB Statistics Menu Options (/stats/slb)

Command Syntax and Usage
real *<real server number (1-1024)>* Displays the following real server statistics: ■ Number of times the real server has failed its health checks. ■ Number of sessions currently open on the real server. ■ Total sessions on the real server. ■ Highest number of simultaneous sessions recorded for each real server. ■ Real server transmit/receive octets. To view an example of what is displayed on the screen, see page 101.
group *<real server group number (1-256)>* Displays the following real server group statistics: ■ 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. ■ Highest number of simultaneous sessions recorded for each real server. ■ Real server transmit/receive octets. For more information, see page 102.
virt *<virtual server number (1-1024)>* Displays the following virtual server statistics: ■ 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. ■ Highest number of simultaneous sessions recorded for each real server. ■ Real server transmit/receive octets. For per server octet counters, see the procedure on page 101. For more information, see page 102.

Table 5-8 SLB Statistics Menu Options (/stats/slb) (Continued)

Command Syntax and Usage
filt *<filter ID (1-1024)>* Displays the total number of times the specified filter has been used. For more information, see page 103.
sp *<Enter SP as port (such as b12) or SP index (such as x5)>* Displays the Switch Processor SLB Statistics Menu. Enter the SP as port or index. For more information, see page 103.
gslb Displays the Global SLB Statistics Menu. For more information, see page 108.
maint Displays the SLB maintenance statistics. To see an example of what is displayed on-screen, see page 111.
clear Clears all switch SLB statistics. The values are reset to zero.
dump Dumps all switch SLB statistics. Use this command to gather data for tuning and debugging switch performance. If you want to capture dump data to a file, set your communication software on your workstation to capture session data before you enter the `dump` command.

/stats/slb/real <real server number>

Real Server SLB Statistics

This command displays the SLB statistics of the selected real server (1-1024). A sample follows:

```
Real server 1 stats:
Health check failures:                0
Current sessions:                     0
Total sessions:                     745
Highest sessions:                    15
Octets:                               0
```

Per Service Octet Counters

For each load-balanced real server, the octet counters represent the combined number of transmit and receive bytes (octets). These counters are then added to report the total octets for each virtual server.

The octet counters are provided per server—not per service. If you need octet counters on a per-service basis, you can accomplish this through the following configuration:

1. **Configure a separate IP address for each service on each server being load balanced.**

   For instance, you can configure IP address 10.1.1.20 for HTTP services, and 10.1.1.21 for FTP services on the same physical server.

2. **On the switch, configure a real server with a real IP address for each service above.**

   Continuing the previous example (from step 1), two real servers would be configured for the physical server (representing each real service). If there were five physical servers providing the two services (HTTP and FTP), 10 real servers would have to be configured: five for the HTTP services on each physical server, and five for the FTP services on each physical server.

3. **On the switch, configure one real server group for each type of service, and group each appropriate real server IP address into the group that handles the specific service.**

   Thus, in keeping with our example, two groups would be configured: one for handling HTTP and one for handling FTP.

4. **Configure a virtual server and add the appropriate services to that virtual server.**

/stats/slb/group <real server group number>

Real Server Group Statistics

This item displays the statistics of *Current Sessions*, *Total Sessions*, *Highest Sessions*, and *Octets* for all real servers that belong to the selected real server group (1-256) on a switch port.

```
Real server group 1 stats:
                      Current      Total  Highest
Real IP address      Sessions   Sessions Sessions           Octets
---- --------------- -------- ---------- --------  ---------------
   1 123.255.13.12          0        745       15                0
   2 123.255.13.21          0        739       17                0
---- --------------- -------- ---------- --------  ---------------
                            0       1484       32                0
```

/stats/slb/virt <virtual server number>

Virtual Server SLB Statistics

This command displays the statistics of *Current sessions*, *Total Sessions*, *Highest Sessions*, and *Octets* for all virtual servers that belong to the selected real server group (1-256).

```
Virtual server 1 stats:
                      Current      Total  Highest
Real IP address      Sessions   Sessions Sessions           Octets
---- --------------- -------- ---------- --------  ---------------
   1 123.255.13.12          0        745       15                0
   2 123.255.13.21          0        739       17                0
---- --------------- -------- ---------- --------  ---------------
     205.178.13.1           0       1484       32                0
```

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

/stats/slb/filt <filter number>

Filter SLB Statistics

Total firings displays the number of filter hits on the switch port for the selected filter (1-1024).

```
Filter 1 stats:
Total firings:                          1011
```

/stats/slb/sp <sp number>

Switch Processor SLB Statistics

```
[Server Load Balancing SP Statistics Menu]
      real    - Real server statistics
      group   - Real server group statistics
      virt    - Virtual server statistics
      filt    - Filter statistics
      maint   - Maintenance statistics
      clear   - Clear SP statistics
```

Table 5-9 Switch Processor SLB Statistics Menu Options (/stats/slb/sp)

Command Syntax and Usage
real *<1-1024>* Displays real server statistics of the switch processor. For more information, see page 104.
group *<1-256>* Displays real server group statistics of the switch processor. For more information, see page 104.
virt *<1-1024>* Displays virtual server statistics of the switch processor. For more information, see page 105.
filt *<1-1024>* Displays filter statistics of the switch processor. For more information, see page 105.
maint Displays maintenance statistics for the selected switch processor. For more information, see page 106.
clear Clears all switch processor SLB statistics for this switch processor, resetting them to zero.

/stats/slb/sp <sp number>/real <server number>

SP Real Server SLB Statistics

```
SP index 0 (port A1) Real server 1 stats:
Current sessions:                  677
Total sessions:                 7313257
Octets:                               0
```

/stats/slb/sp <sp number>/group <group number>

SP Real Server Group SLB Statistics

This command displays statics of the selected real server group.

```
SP index 0 (port A1) Real server group 1 stats:
Real server IP address     Current Sessions   Total Sessions         Octets
---------- -------------- ----------------  --------------- -------------
       1     123.255.24.6            460         7317741                 0
       2     123.255.24.133          341         5527415                 0
  ---------- -------------- ----------------  --------------- -----------
     Totals                          801        12845156                 0
```

/stats/slb/sp <sp number>/virt <server number>

SP Virtual Server SLB Statistics

This command displays statistics of the selected virtual server.

```
SP index 0 (port A1) Virtual server 1 stats:
Real server IP address     Current Sessions   Total Sessions         Octets
  ----------- -------------- ----------------  --------------- -----------
        1     123.255.24.6           649         7323544                0
        2     123.255.24.133         488         5531652                0
  ----------- -------------- ----------------  --------------- -----------
     Totals   172.21.11.200        1137         12855196               0
```

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

/stats/slb/sp <sp number>/filter <filter number>

SP Filter SLB Statistics

This menu option displays the total number of times a filter has been fired on the selected switch processor.

```
SP index 0 (port A1) Filter 1024 stats:
Total firings:                        0
```

/stats/slb/sp <sp number>/maint <server number>

SP Maintenance SLB Statistics

This menu provides information that may be requested by an Alteon WebSystems support engineer for troubleshooting purposes.

```
SP index 0 (port A1) SLB Maintenance stats:
Current sessions:                0   Highest sessions:                0
  4 second average:              0    64 second average:              0
Allocation failures:             0   L4 disabled:                     0
Overflows:                       0   Not ready:                       0
Lookup failures:                 0   Non TCP/IP frames:               0
IP options:                      0   TCP fragments:                   0
UDP datagrams:                   0   Filtered (denied) frames:        0
Incorrect VIPs:                  0   Incorrect Vports:                0
No available real server:        0   Timed-out sessions:              0
Client slowpath frames:          0   Client sessions created:         0
Server slowpath frames:          0   ICMP slowpath frames:            0
DAM slowpath frames:             0   DAM sessions created:            0
SW alloc failures:               0
```

Table 5-10 SP Server Load Balancing Maintenance Statistics (/stats/slb/sp/maint)

Statistic	Description
`Current sessions`	Number of session bindings currently in use for the specified SP.
`Allocation failures`	The number of times a session entry could not be created because the table for the specified SP was full.
`L4 disabled`	The number of frames addressed to a virtual server for which client processing is disabled for the specified SP.
`Overflows`	The number of hash chain overflows for the specified SP.
`Not ready`	The number of times that a session entry for the specified SP was not ready for second or subsequent frames.
`Lookup failures`	The number of frames for the specified SP for which a Layer 3 (MAC address) or Layer 4 (IP address) lookup failed.
`TCP fragments`	The number of TCP fragments that were dropped for the specified Switch Processor.
`UDP datagrams`	The number of UDP datagrams that were dropped because the virtual service for the specified SP did not support UDP.

Table 5-10 SP Server Load Balancing Maintenance Statistics (/stats/slb/sp/maint) (Continued)

Statistic	Description
IP options	The number of frames for the specified SP that could not be processed completely because they contain IP options.
Non TCP/IP frames	The number of non-IP based frames for the specified SP that were received by the virtual server.
Incorrect VIPs	The number of times that a frame was dropped for the specified SP because it was addressed to a VMAC but an unknown VIP. The switch has received a Layer 4 request for a virtual server that was not configured.
Incorrect Vports	The number of frames for the specified SP that were dropped to a valid VIP but not to a valid VPORT (destination port). 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	The number of times a frame for the specified SP was dropped because there were no real servers available to accept the session. All real servers are either out of service or at their mcon limit.
Filtered (denied) Frames	The number of frames that where dropped because they matched an active filter with the "deny" action set.
Client slowpath frames	The number of client frames processed by the control processor(s).
Server slowpath frames	The number of server frames processed by the control processor(s).
ICMP slowpath frames	Number of client ICMP frames processed by the control processor(s).
DAM slowpath frames	The number of DAM frames processed by the control processor(s).
DAM sessions created	The number of DAM sessions added to the session table by the control processor(s).
Slowpath sessions created	The number of sessions added to session table by the control processor(s).
SW alloc failures	The number of instances where the switch ran out of available bindings for a port.

/stats/slb/gslb

Global SLB Statistics

```
[Global SLB Statistics Menu]
      real    - Real server statistics
      group   - Real server group statistics
      virt    - Virtual server statistics
      maint   - Global SLB maintenance statistics
      clear   - Clear SLB statistics
```

Table 5-11 Global SLB Statistics Menu Options (/stats/slb/gslb)

Command Syntax and Usage
real *<real server number (1-1024)>* Where the real server number represents the real server ID on this switch, under which the remote server is configured. For more information, see page 111.
group *<real server group number (1-256)>* Displays real server group global statistics. For more information, see page 109.
virt *<virtual server number (1-1024)>* Displays virtual server group global statistics. For more information, see page 110.
maint Displays global SLB maintenance statistics. For more information, see page 111.
clear Clears all global SLB statistics for the switch.

/stats/slb/gslb/real <real server number>

Real Server Global SLB Statistics

NOTE – The *real server number* (1-1024) represents the real server you wish to observe.

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

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

- Number of DNS hand-offs to the remote server
- Number of HTTP redirects to the remote server

/stats/slb/gslb/group <group number>

Real Server Group Global SLB Statistics

NOTE – The *group number* (1-256) represents the server group you wish to observe.

```
Real server group 1 Global SLB stats:
  Real server IP address          DNS Handoffs  HTTP Redirects
  ----------- --------------- ----------------  ---------------
        1     123.255.224.54           1240              30
        2     123.255.224.223           608              12
  ----------- --------------- ----------------  ---------------
     Totals                           1848              42
```

Real server group global statistics include the following information:

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

/stats/slb/gslb/virt

<virtual server number>

Virtual Server Global SLB Statistics

```
Virtual server 1 Global SLB stats:
  Service Server IP address      Response time Min sessions avail
  ------- ------ --------------- ------------- ------------------
  http     v1    123.255.224.55          16            21190
  http     r1    123.255.224.54          10            24120
  telnet   v1    123.255.224.55           4            31032
```

Virtual server global statistics include the following information:

- 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. On the switch, each remote site is configured as if it were a real server. The number (for example, r1) represents the real server ID on this switch. The remote server is configured with this number.
- 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.
- Minimum 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.

/stats/slb/gslb/maint

Global SLB Maintenance Statistics

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

Global SLB maintenance statistics include the following information.

- 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 GSLB switch configuration problem. If bad updates occur, check your `syslog` for configuration error messages.

/stats/slb/maint

SLB Maintenance Statistics

This menu contains information that may be requested by an Alteon WebSystems support engineer for troubleshooting purposes.

```
SLB Maintenance stats:
Current sessions:              717   Highest sessions:             1657
  4 second average:              0    64 second average:             0
Allocation failures:             0   L4 disabled:                   0
Overflows:                       0   Not ready:                     0
Lookup failures:                 0   Non TCP/IP frames:             0
IP options:                      0   TCP fragments:                 0
UDP datagrams:                   0   Filtered (denied) frames:   2514
Incorrect VIPs:                  0   Incorrect Vports:              0
No available real server:        0   Timed-out sessions:         6850
Backup server activated:         0   Ovrflo server activated:       0
Client slowpath frames:      13220   Client sessions created:       5
Server slowpath frames:          0   ICMP slowpath frames:          0
DAM slowpath frames:             0   DAM sessions created:          0
Slowpath alloc failures:         0
```

/stats/mp
Management Processor Statistics

```
[MP Statistics Menu]
      if      - IP interface ("if") statistics
      ip      - IP statistics
      icmp    - ICMP statistics
      tcp     - TCP statistics
      udp     - UDP statistics
      snmp    - SNMP statistics
      arp     - ARP statistics
      route   - Route statistics
      dns     - DNS server statistics
      ospf    - OSPF Statistics Menu
      vrrp    - VRRP statistics
      maint   - MP Maintenance Statistics Menu
      clear   - Clear statistics
```

This menu is used to display management processor (MP) statistics. The operations the MP processes include SNMP, CLI, Telnet, and Spanning Tree Protocol.

Table 5-12 Management Processor Statistics Menu Options (/stats/mp)

Command Syntax and Usage
if *<1-1024>* Displays IP interface statistics for the management processors. The range of interface numbers is 1 through 1024.
ip Displays IP statistics for the management processors.
icmp Displays ICMP statistics for the management processors. For more information see page 118.
tcp Displays TCP statistics for the management processors. For more information see page 120.
udp Displays UDP statistics for the management processors. For more information see page 122.
snmp Displays SNMP statistics for the management processors. For more information see page 122.

Table 5-12 Management Processor Statistics Menu Options (/stats/mp)

Command Syntax and Usage
arp Displays Address Resolution Protocol (ARP) statistics for the management processors. For more information see page 125.
route Displays route statistics for the management processors. For more information see page 126.
dns Displays Domain Name Server (DNS) statistics for the management processors. For more information see page 127.
ospf Displays the OSPF Statistics menu. For more information see page 127.
vrrp When virtual routers are configured for the management processors, you can display the following protocol statistics for VRRP: ■ Advertisements received (`vrrpInAdvers`) ■ Advertisements transmitted (`vrrpOutAdvers`) ■ Advertisements received, but ignored (`vrrpBadAdvers`) For more information see page 128.
maint Displays Maintenance Statistics menu for the management processor. This submenu provides information that may be requested by an Alteon WebSystems support engineer for troubleshooting purposes. For more information see page 128.
clear Clears all management processor statistics for the switch.

/stats/mp/if <interface number>

IP Interface Statistics for Management Processor

```
IP interface 1 statistics:
ifInOctets:           21969155   ifInUcastPkts:            60325
ifInNUCastPkts:         227832   ifInDiscards:              4723
ifInErrors:                  0   ifInUnknownProtos:            0
ifOutOctets:           3039417   ifOutUcastPkts:           86834
ifOutNUcastPkts:           162   ifOutDiscards:                0
ifOutErrors:                 1   ifStateChanges                1
```

Table 5-13 IP Interface Statistics* (/stats/mp/if)

Statistic	Description
ifInOctets	The total number of octets received on the interface, including framing characters.
ifInUcastPkts	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
ifInNUCastPkts	The number of packets, delivered by this sub-layer to a higher (sub-)layer, which were not addressed to a multicast or broadcast address at this sub-layer.
ifInDiscards	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.
ifInErrors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.
ifInUnknownProtos	The number of packets received via the interface which were discarded because of an unknown or unsupported protocol
ifOutOctets	The total number of octets transmitted out of the interface, including framing characters.
ifOutUcastPkts	The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.

*** extract from RFC-1213**

Table 5-13 IP Interface Statistics* (/stats/mp/if) (Continued)

Statistic	Description
ifOutNUcastPkts	The total number of packets that higher-level protocols requested be transmitted to a non- unicast (i.e., a subnetwork-broadcast or subnetwork-multicast) address, including those that were discarded or not sent.
ifOutUcastPkts	The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent
ifOutDiscards	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space
ifOutErrors	For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.
ifStateChanges	The number of times the interface state changes.

*** extract from RFC-1213**

/stats/mp/ip

IP Statistics for Management Processor

```
IP statistics:
ipInReceives:          154588   ipInHdrErrors:              0
ipInAddrErrors:             0   ipForwDatagrams:            0
ipInUnknownProtos:          8   ipInDiscards:               0
ipInDelivers:           76217   ipOutRequests:          86865
ipOutDiscards:             24   ipOutNoRoutes:             24
ipReasmReqds:               0   ipReasmOKs:                 0
ipReasmFails:               0   ipFragOKs:                  0
ipFragFails:                0   ipFragCreates:              0
ipRoutingDiscards:          0   ipDefaultTTL:             255
ipReasmTimeout:             5
```

Table 5-14 IP Statistics (/stats/mp/ip)

Statistic	Description
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInHdrErrors	The number of input datagrams discarded because of errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
ipInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this device. This count includes invalid addresses (such as 0.0.0.0) and addresses of unsupported Classes (such as Class E). For entities that are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
ipForwDatagrams	The number of input datagrams for which this device was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter will include only those packets which were Source-Routed via this device, and the Source- Route option processing was successful.
ipInUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
ipInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (such as for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
ipInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
ipOutRequests	The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.

*** Extract from RFC-1213**

Table 5-14 IP Statistics (/stats/mp/ip) (Continued)

Statistic	Description
ipOutDiscards	The number of output IP datagrams, for which no problems were encountered to prevent transmission to their destination, but which were discarded (for example, lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
ipOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this no-route criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.
ipReasmReqds	The number of IP fragments received that needed to be reassembled at this device.
ipReasmOKs	The number of IP datagrams successfully reassembled.
ipReasmFails	The number of failures detected by the IP re- assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.
ipFragOKs	The number of IP datagrams that have been successfully fragmented at this device.
ipFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this device but could not. For example: because the *Don't Fragment* flag was set.
ipFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this device.
ipRoutingDiscards	The number of routing entries that were discarded although they were valid. One possible reason for discarding such an entry could be to free up buffer space for other routing entries.
ipDefaultTTL	The default value inserted into the Time-To-Live field of the IP header of datagrams originated at this device, whenever a TTL value is not supplied by the transport layer protocol.
ipReasmTimeout	The maximum number of seconds that received fragments are held while they are awaiting reassembly at this device.

*** Extract from RFC-1213**

/stats/mp/icmp

ICMP Statistics for Management Processor

```
ICMP statistics:
icmpInMsgs:              78683   icmpInErrors:                70
icmpInDestUnreachs:          0   icmpInTimeExcds:              0
icmpInParmProbs:             0   icmpInSrcQuenchs:             0
icmpInRedirects:            34   icmpInEchos:              26908
icmpInEchoReps:          51671   icmpInTimestamps:             0
icmpInTimestampReps:         0   icmpInAddrMasks:              0
icmpInAddrMaskReps:          0   icmpOutMsgs:              78960
icmpOutErrors:               0   icmpOutDestUnreachs:         11
icmpOutTimeExcds:            0   icmpOutParmProbs:             0
icmpOutSrcQuenchs:           0   icmpOutRedirects:             0
icmpOutEchos:            52041   icmpOutEchoReps:          26908
icmpOutTimestamps:           0   icmpOutTimestampReps:         0
icmpOutAddrMasks:            0   icmpOutAddrMaskReps:          0
```

Table 5-15 ICMP Statistics* (/stats/mp/icmp)

Statistic	Description
`icmpInMsgs`	The total number of ICMP messages which the device received. Note that this counter includes all those counted by `icmpInErrors`.
`icmpInErrors`	The number of ICMP messages that the device received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
`icmpInDestUnreachs`	The number of ICMP Destination Unreachable messages received.
`icmpInTimeExcds`	The number of ICMP Time Exceeded messages received.
`icmpInParmProbs`	The number of ICMP Parameter Problem messages received.
`icmpInSrcQuenchs`	The number of ICMP Source Quench messages received.
`icmpInRedirect`	The number of ICMP Redirect messages received.
`icmpInEchos`	The number of ICMP Echo (request) messages received.
`icmpInEchoReps`	The number of ICMP Echo Reply messages received.
`icmpInTimestamps`	The number of ICMP Timestamp (request) messages received.
`icmpInTimestampReps`	The number of ICMP Timestamp Reply messages received.

*** Extract from RFC-1213**

Table 5-15 ICMP Statistics* (/stats/mp/icmp) (Continued)

Statistic	Description
icmpInAddrMasks	The number of ICMP Address Mask Request messages received.
icmpInAddrMaskReps	The number of ICMP Address Mask Reply messages received.
icmpOutMsgs	The total number of ICMP messages that this switch attempted to send. This count includes all those counted by icmpOutErrors.
icmpOutErrors	The number of ICMP messages which this device did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer, such as the inability of IP to route the resultant datagram. In some implementations, there may be no types of error which contribute to this counter's value.
icmpOutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
icmpOutTimeExcds	The number of ICMP Time Exceeded messages sent.
icmpOutParmProbs	The number of ICMP Parameter Problem messages sent.
icmpOutSrcQuenchs	The number of ICMP Source Quench messages sent.
icmpOutRedirects	The number of ICMP Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.
icmpOutEchos	The number of ICMP Echo (request) messages sent.
icmpOutEchoReps	The number of ICMP Echo Reply messages sent.
icmpOutTimestamps	The number of ICMP Timestamp (request) messages sent.
cmpOutTimestampReps	The number of ICMP Timestamp Reply messages sent.
icmpOutAddrMasks	The number of ICMP Address Mask Request messages sent.
icmpOutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.

*** Extract from RFC-1213**

/stats/mp/tcp

TCP Statistics for Management Processor

```
TCP statistics:
tcpRtoAlgorithm:              4   tcpRtoMin:                    0
tcpRtoMax:               240000   tcpMaxConn:                1024
tcpActiveOpens:               0   tcpPassiveOpens:            352
tcpAttemptFails:              0   tcpEstabResets:             325
tcpInSegs:                 7637   tcpOutSegs:                5834
tcpRetransSegs:              34   tcpInErrs:                    0
tcpCurBuff:                   0   tcpCurConn:                   4
tcpOutRsts:                1368
```

Table 5-16 TCP Statistics (/stats/mp/tcp)

Statistic	Description
tcpRtoAlgorithm	The algorithm used to determine the time-out value used for retransmitting unacknowledged octets.
tcpRtoMin	The minimum value permitted by a TCP implementation for the retransmission time-out, measured in milliseconds. More refined semantics for objects of this type depend upon the algorithm used to determine the retransmission time-out. In particular, when the time-out algorithm is rsre(3), an object of this type has the semantics of the LBOUND quantity described in RFC 793.
tcpRtoMax	The maximum value permitted by a TCP implementation for the retransmission time-out, measured in milliseconds. More refined semantics for objects of this type depend upon the algorithm used to determine the retransmission time-out. In particular, when the time-out algorithm is rsre(3), an object of this type has the semantics of the UBOUND quantity described in RFC 793.
tcpMaxConn	The limit on the total number of TCP connections the device can support. In entities where the maximum number of connections is dynamic, this object should contain the value *-1*.
tcpActiveOpens	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state.
tcpPassiveOpens	The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state.

*** Extract from RFC-1213**

Table 5-16 TCP Statistics (/stats/mp/tcp) (Continued)

Statistic	Description
tcpAttemptFails	The number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
tcpEstabResets	The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
tcpInSegs	The total number of segments received, including those received in error. This count includes segments received on currently established connections.
tcpOutSegs	The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets.
tcpRetransSegs	The total number of segments retransmitted—that is, the number of TCP segments transmitted containing one or more previously transmitted octets.
tcpInErrs	The total number of segments received in error (such as bad TCP checksums).
tcpCurBuff	The total number of outstanding memory allocations from heap by TCP protocol stack.
tcpCurConn	The total number of outstanding TCP sessions that are currently open.
tcpOutRsts	The number of delivered TCP segments that contain the RST flag.

*** Extract from RFC-1213**

/stats/mp/udp

UDP Statistics for Management Processor

```
UDP statistics:
udpInDatagrams:          25471   udpOutDatagrams:          1240
udpInErrors:                 0   udpNoPorts:              43873
```

Table 5-17 UDP Statistics (/stats/mp/udp)

Statistic	Description
udpInDatagrams	The total number of UDP datagrams delivered to UDP users.
udpOutDatagrams	The total number of UDP datagrams sent from this device.
udpInErrors	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
udpNoPorts	The total number of received UDP datagrams for which there was no application at the destination port.

/stats/mp/snmp

SNMP Statistics for Management Processor

```
SNMP statistics:
snmpInPkts:              25485   snmpInBadVersions:           0
snmpInBadC'tyNames:          0   snmpInBadC'tyUses:           0
snmpInASNParseErrs:          0   snmpEnableAuthTraps:         0
snmpOutPkts:             25485   snmpInBadTypes:              0
snmpInTooBigs:               0   snmpInNoSuchNames:           0
snmpInBadValues:             0   snmpInReadOnlys:             0
snmpInGenErrs:               0   snmpInTotalReqVars:      50661
snmpInTotalSetVars:          0   snmpInGetRequests:         405
snmpInGetNexts:          25080   snmpInSetRequests:           0
snmpInGetResponses:          0   snmpInTraps:                 0
snmpOutTooBigs:              0   snmpOutNoSuchNames:          0
snmpOutBadValues:            0   snmpOutReadOnlys:            0
snmpOutGenErrs:              0   snmpOutGetRequests:          0
snmpOutGetNexts:             0   snmpOutSetRequests:          0
snmpOutGetResponses:     25485   snmpOutTraps:                0
```

Table 5-18 SNMP Statistics (/stats/mp/snmp)

Statistic	Description
snmpInPkts	The total number of Messages delivered to the SNMP device from the transport service.
snmpInBadVersions	The total number of SNMP Messages that were delivered to the SNMP protocol device and were for an unsupported SNMP version.
snmpInBadC'tyNames	The total number of SNMP Messages delivered to the SNMP protocol device that used a SNMP community name not known to said device.
snmpInBadC'tyUses	The total number of SNMP Messages delivered to the SNMP protocol device that represented an SNMP operation which was not allowed by the SNMP community named in the Message.
snmpInASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP protocol device when decoding received SNMP Messages.
snmpEnableAuthTraps	Indicates whether the SNMP agent process is permitted to generate authentication-failure traps. The value of this object overrides any configuration information; as such, it provides a means whereby all authentication-failure traps may be disabled. Note that it is strongly recommended that this object be stored in non-volatile memory so that it remains constant between re-initializations of the network management system.
snmpOutPkts	The total number of SNMP Messages that were passed from the SNMP protocol device to the transport service.
snmpInBadTypes	The total number of SNMP Messages that failed ASN parsing.
snmpInTooBigs	The total number of SNMP PDUs that were delivered to the SNMP protocol device and for that the value of the error-status field is `tooBig`.
snmpInNoSuchNames	The total number of SNMP PDUs that were delivered to the SNMP protocol device and for that the value of the error-status field is `noSuchName`.
snmpInBadValues	The total number of SNMP PDUs that were delivered to the SNMP protocol device and for that the value of the error-status field is `badValue`.

Table 5-18 SNMP Statistics (/stats/mp/snmp) (Continued)

Statistic	Description
snmpInReadOnlys	The total number of valid SNMP PDUs that were delivered to the SNMP protocol device and for which the value of the error-status field is 'readOnly'. It should be noted that it is a protocol error to generate an SNMP PDU that contains the value `readOnly' in the error-status field, as such this object is provided as a means of detecting incorrect implementations of the SNMP.
snmpInGenErrs	The total number of SNMP PDUs that were delivered to the SNMP protocol device and for which the value of the error-status field is genErr.
snmpInTotalReqVars	The total number of MIB objects that have been retrieved successfully by the SNMP protocol device as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
snmpInTotalSetVars	The total number of MIB objects that have been altered successfully by the SNMP protocol device as the result of receiving valid SNMP Set-Request PDUs.
snmpInGetRequests	The total number of SNMP Get-Request PDUs that have been accepted and processed by the SNMP protocol device.
snmpInGetNexts	The total number of SNMP Get-Next PDUs that have been accepted and processed by the SNMP protocol device.
snmpInSetRequests	The total number of SNMP Set-Request PDUs that have been accepted and processed by the SNMP protocol device.
snmpInGetResponses	The total number of SNMP Get-Response PDUs that have been accepted and processed by the SNMP protocol device.
snmpInTraps	The total number of SNMP Trap PDUs that have been accepted and processed by the SNMP protocol device.
snmpOutTooBigs	The total number of SNMP PDUs that were generated by the SNMP protocol device and for that the value of the error-status field is tooBig.
snmpOutNoSuchNames	The total number of SNMP PDUs that were generated by the SNMP protocol device and for which the value of the error-status is noSuchName.
snmpOutBadValues	The total number of SNMP PDUs that were generated by the SNMP protocol device and for which the value of the error-status field is badValue.
snmpOutReadOnlys	Not in use.

Table 5-18 SNMP Statistics (/stats/mp/snmp) (Continued)

Statistic	Description
snmpOutGenErrs	The total number of SNMP PDUs that were generated by the SNMP protocol device and for which the value of the error-status field is genErr.
snmpOutGetRequests	The total number of SNMP Get-Request PDUs that have been generated by the SNMP protocol device.
snmpOutGetNexts	The total number of SNMP Get-Next PDUs that have been generated by the SNMP protocol device.
snmpOutSetRequests	The total number of SNMP Set-Request PDUs that have been generated by the SNMP protocol device.
snmpOutGetResponses	The total number of SNMP Get-Response PDUs that have been generated by the SNMP protocol device.
snmpOutTraps	The total number of SNMP Trap PDUs that have been generated by the SNMP protocol device.

/stats/mp/arp

Address Resolution Protocol Statistics for MP

```
ARP statistics:
arpEntriesCur:                5   arpEntriesHighWater:          6
arpEntriesMax:             8192
```

Table 5-19 ARP Statistics (/stats/mp/arp)

Statistic	Description
arpEntriesCur	The total number of outstanding ARP entries in the ARP table.
arpEntriesHighWater	The highest number of ARP entries ever recorded in the ARP table.
arpEntriesMax	The maximum number of ARP entries that are supported.

/stats/mp/route

Route Statistics for Management Processor

```
Route statistics:
ripInPkts:                    0   ripOutPkts:                   0
ripInDiscards:                0   ripRoutesAgedOut:            0
ipRoutesCur:                  8   ipRoutesHighWater:           8
ipRoutesMax:             8192
```

Table 5-20 Route Statistics (/stats/mp/route)

Statistic	Description
ripInPkts	The total number of good Routing Information Protocol (RIP) advertisement packets that have been received.
ripOutPkts	The total number of RIP advertisement packets that have been sent.
ripInDiscards	The total number of received RIP advertisement packets that have been discarded.
ripRoutesAgedOut	The total number of route learned via RIP that have aged out.
ipRoutesCur	The total number of outstanding routes in the route table.
ipRoutesHighWater	The highest number of routes ever recorded in the route table.
ipRoutesMax	The maximum number of routes that are supported.

/stats/mp/dns

DNS Statistics for Management Processor

```
DNS statistics:
dnsInRequests:                0   dnsOutRequests:               2
dnsBadRequests:               0
```

Table 5-21 DNS Statistics (/stats/mp/dns)

Statistic	Description
dnsInRequests	The total number of DNS request packets that have been received.
dnsOutRequests	The total number of DNS response packets that have been transmitted.
dnsBadRequests	The total number of DNS request packets received that were dropped.

/stats/mp/ospf

OSPF Statistics Menu for Management Processor

```
[OSPF Statistics Menu]
      general - Show global stats
      aindex  - Show area(s) stats
      if      - Show interface(s) stats
```

Table 5-22 OSPF Statistics (/stats/mp/ospf)

Command Syntax and Usage
general Displays the OSPF global statistics.
aindex Displays the OSPF area statistics.
if Displays the OSPF interface statistics.

/stats/mp/vrrp

VRRP Statistics for Management Processor

```
VRRP statistics:
      vrrpInAdvers:          973614    vrrpBadAdvers:          0
      vrrpOutAdvers:              0
```

Virtual Router Redundancy Protocol (VRRP) support on Alteon Web switches provides redundancy between routers in a LAN. This redundancy is accomplished by configuring the same virtual router IP address and ID number on each participating VRRP-capable routing device. One of the virtual routers is then elected as the master, based on a number of priority criteria, and assumes control of the shared virtual router IP address. If the master fails, one of the backup virtual routers will assume routing authority and take control of the virtual router IP address.

When virtual routers are configured, you can display the following protocol statistics for VRRP.

- Advertisements received—`vrrpInAdvers`
- Advertisements transmitted—`vrrpOutAdvers`
- Advertisements received, but ignored—`vrrpBadAdvers`

/stats/mp/maint

Maintenance Statistics for Management Processor

```
[MP Maintenance Statistics Menu]
      ltr     - Letter counts
      msg     - STEM message statistics
      mem     - STEM memory statistics
      amem    - All STEM memory blocks in use
      thr     - STEM thread statistics
      pkt     - Packet statistics
      tcb     - All TCP control blocks in use
      ucb     - Show All UDP control blocks in use
      uart    - UART statistics
      cpu     - Show CPU utilization
```

NOTE – The MP maintenance submenu and options contain hardware-specific information related to the management processor. The example screens shown in this section are for informational purposes only. You may be asked by an Alteon WebSystems support engineer to provide MP maintenance information for troubleshooting purposes.

Table 5-23 Maintenance Statistics (/stats/mp/maint)

Command Syntax and Usage
ltr Displays letter count statistics. To view an example, see page 130.
msg Displays Simple Thread Execution Model (STEM) message statistics. To view a sample, see page 131.
mem Displays STEM memory statistics, showing available memory. To view a sample, see page 131.
amem Displays all STEM memory blocks in use to check for leaks. To view a sample, see page 132.
thr Displays STEM thread statistics. To view a sample, see page 133.
pkt Displays packet statistics to check for leads and load. To view a sample, see page 134.
tcb Displays all TCP control blocks (TCB) in use. To view a sample, see page 134.
ucb Displays all UCP control blocks in use. To view a sample, see page 134.
uart Displays universal asynchronous receiver/transmitter (UART) statistics. To view a sample, see page 135.
cpu Displays CPU utilization for periods of up to 1, 4, and 64 seconds. There are two CPUs per port. To view a sample, see page 135.

/stats/mp/maint/ltr

Letter Statistics for Management Processor

```
Letter statistics for MP:

    ------------ Received -------------  -------------- Sent -------
SP   non-ACKs     ACKs     Dups  NonSeq   non-ACKs     ACKs    Rxmits Defers
--  ---------- ---------- ------ ------  ---------- ---------- ------ -------
 0       16809     122171      0      0      122173      16807      0    1347
 1       21889     652145      0      0      652147      21887      0   20861
 2           9      44031      0      0       44033          7      0     201
 3          11      44032      0      0       44034          9      0     204
 4        2821      43073      0      0       43081       2813      0     232
 5          16      44065      0      0       44067         14      0     320

 Letters By Type     Received        Sent        Rxmitted      Deferred
-----------------  ------------  ------------  ------------  -------
 ACK                    949517         41537             0            0
 BOOT_ME                     6             0             0            0
 DOWNLOAD                    0           264             0            0
 GO                          0             6             0            0
 INIT_DONE                   6             0             0            0
 FLAGS_CHANGE                0            20             0            8
 STATS_REQ                   0            27             0            0
 STATS_RSP                  27             0             0            0
 SLB_REQ                     0           438             0          428
 IP_ROUTE_REQ                0        753519             0        22582
 SLOT_INFO                   0            21             0            7
 LINK_DOWN_NOTE              0             6             0            0
 LINK_DOWN_DONE              6             0             0            0
 LINK_UP_OKAY                0             1             0            0
 LINK_UP_NOTE                0            18             0            0
 GDT_UPDATE                  0           714             0            0
 EVID_UPDATE                 0            42             0           42
 HEARTBEAT                   0        193579             0           20
 FDB_REQ                     0            84             0           44
 VLAN_CLASS_UPDT             0             6             0            2
 ACTIVE_LINKS                0             6             0            6
 SYSMON_REQ                  0             4             0            0
 FDB_AGE_SCAN                0           774             0           26
 FLAGS_CHANGED              24             0             0            0
 SYSMON_DATA              5593             0             0            0
```

/stats/mp/maint/msg

STEM Message Statistics for Management Processor

This menu displays the STEM message statistics of the management processor.

```
STEM message statistics for MP:
-------------------------------------------------------------------
 allocs:           6988   frees:          8399631
 callocs:       8392643   calloc_fails:         0
 sends:         8399632   returns:              0
 alloc_curr:          0   alloc_hiwat:         41
 alloc_fails:         0
```

/stats/mp/maint/mem

STEM Memory Statistics for Management Processor

This menu displays the statistics of the management processor.

```
STEM memory statistics for MP:
-------------------------------------------------------------------
 allocs:        611271   frees:          610935
 alloc_fails:        0   pool_bytes:    8388608
 bytes_curr:    371232   bytes_hiwat:   2470272
 largest:      2097152
```

/stats/mp/maint/amem

All STEM Memory Allocated Blocks for Management Processor

This menu displays the statistics of the STEM memory allocated blocks.

```
----------------------------------------------------------------------
All STEM memory allocated blocks:
Number    Caller    Blocks       Bytes
------   --------   ------  ----------
   1     004465f8       26      165888
   2     004f24b8     1183      421280
   3     004c21c4     1024       65536
   4     004c3264     1024       32768
   5     004c2260       34        2176
   6     004c3348       34        2176
   7     004c22a0       34        4352
   8     004b765c        4         256
   9     00430688        1          32
  10     00404c28        1        4096
  11     00404a2c        4       16384
  12     0054cadc      253       20864
  13     004b90b0        2        1024
  14     0042391c        6        3072
  15     004bf818        1          32
  16     004f4888        2          64
  17     0041d7b8        2        1024
  18     0041e844        1        8192
  19     004c502c        2        2048
  20     004c0a38        3         192
  21     004c0e44        3         192
  22     004b20f0        3          96
  23     004f2440       19         608
  24     004b8f68        7        3584
  25     004f45a8        1          64
  26     004b842c        3         512
  27     0044830c        1          64
  28     0044bdd8        1          64
   :
   :
```

/stats/mp/maint/thr

STEM Thread Statistics for Management Processor

This command displays the Simple Thread Execution Model (STEM) statistics of the management processor (MP). A sample follows:

```
STEM thread statistics for MP:
-------------------------------------------------------------------
  thid  name        stack              maximum pathlen
                 size   used      uSecs     cmd      word1
    1   STEM                         0  00000000  00000000
    2   STP      2028    624      2611  00000002  00000000
    4   TND                       1916  00000002  00000000
    5   CONS     8172   2048     35933  00000001  fbfbfbfb
    6   TNET     8172   2264     35862  00050001  004b38cc
    7   TNET     8172   1296      2144  00050001  004b38cc
    8   TNET     8172    144         0  00000000  00000000
    9   TNET     8172    144         0  00000000  00000000
   10   LOG      4076    840      1831  00060001  00060005
   11   TRAP     4076    248        86  00070003  00000000
   14   RMON                       64  00000002  00000010
   15   SLB      8172    544      6542  00000007  1f829f7a
   16   DSLB     8172    224        15  00000002  00000000
   17   IP       8172   1776      2233  00000005  0000ff00
   18   RIP      8172   1392     32821  00000004  009e0c8c
   19   AGR      8172    256        42  00000000  009e61ac
   20   EPI      8172   1000      2693  00000001  004b38cc
   21   SLOT     4076    992     92612  000a0005  00000001
   22   PORT     8172    848      1632  00090001  003f9120
   23   STAT     4076    616      2249  00000002  00000000
   31   VLAN     4076   1160      6738  00000002  00000000
   32   MANT     4076    144         0  00000000  00000000
   33   FDB      8172    144         0  00000000  00000000
   34   OSPF     8172    128         0  00000000  00000000
   36   VRRP     4076    280        15  00000008  0001a800
   37   SNMP     4076    128         0  00000000  00000000
   38   SMON     4076    896     80727  000e000b  fbfbfbfb
   39   CFG      4076    128         0  00000000  00000000
   53   DONE     8172   1168     32077  00000001  fbfbfbfb
    :
    :
```

/stats/mp/maint/pkt

Packet Counts for Management Processor

This menu displays the statistics of the packets.

```
Packet counts:
 allocs:          804115   frees:         804115
 mediums:              0   jumbos:             0
 smalls:               0   failures:        4723
```

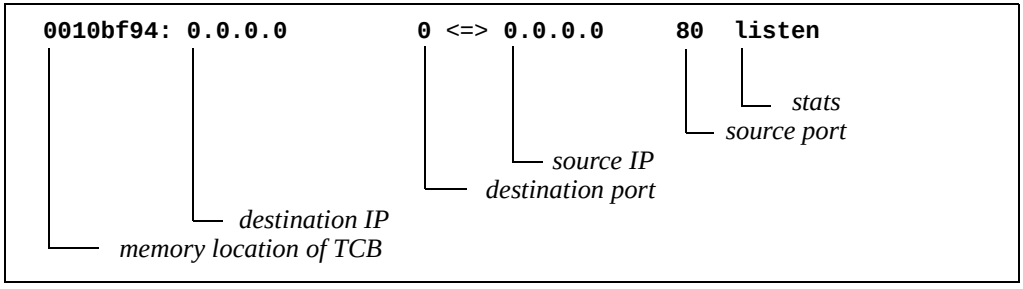
/stats/mp/maint/tcb

All TCB Allocated Control Blocks for Management Processor

This menu displays the statistics of the TCP Control Block (TCB).

```
All TCP allocated control blocks:
0090edbc:  0.0.0.0              0 <=> 0.0.0.0             80  listen
00b6ed3c:  123.255.224.3     1510 <=> 123.255.224.11      23  established
009144fc:  0.0.0.0              0 <=> 0.0.0.0             23  listen
```

The TCB statistics are described below.



/stats/mp/maint/ucb

UDP Allocated Control Blocks for Management Processor

```
All UDP allocated control blocks:
   53:  listen
  161:  listen
```

/stats/mp/maint/uart

UART Statistics for Management Processor

This command displays the statistics of the UART. A sample follows:

```
UART:
 input overflows:    0
 Rx discards:        0  Tx discards:        0
 X-OFFs seen:        0  X-ONs seen:         0
 X-OFFs sent:        0  X-ONs sent:         0
```

/stats/mp/maint/cpu

CPU Statistics for Management Processor

The CPU statistics show CPU utilization in periods of 1, 4, and 64 seconds. A sample follows:

```
CPU utilization:
cpUtil1Second:            13%  fpUtil1Second:            0%
cpUtil4Seconds:           13%  fpUtil4Seconds:           0%
cpUtil64Seconds:          13%  fpUtil64Seconds:          0%
```

/stats/sp <port number>
Switch Processor Statistics

```
[SP index 0 (port A1) Statistics Menu]
      ip      - IP statistics
      maint   - SP Maintenance Statistics Menu
      clear   - Clear statistics
```

This menu allows you to view the Switch Processor (SP) statistics of the selected port.

Table 5-24 Switch Processor Statistics Menu Options (/stats/sp)

Command Syntax and Usage	
ip	Displays IP interface statistics for the switch processor. The range of the interface number is 1 through 1024. For more information, see page 115.
maint	Displays maintenance statistics for the switch processor. This submenu contains information that may be requested by an Alteon WebSystems support engineer for troubleshooting purposes. For more information, see page 138.
clear	Clears all switch processor statistics. All values are reset to zero.

/stats/sp/ip
IP Statistics for Switch Processor

```
IP statistics for SP index 0 (port A1):
ipInReceives:          94080   ipInHdrErrors:           0
ipInDelivers:          94080   ipForwDatagrams:         0
ipInUnknownProtos:         0   ipInDiscards:            0
ipFragOKs:                 0   ipFragCreates:           0
ipDontFrags:               0   ipFragFails:             0
ipOutNoRoutes:             0
```

Table 5-25 IP Statistics (/stats/sp/ip)

Statistic	Description
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.

Table 5-25 IP Statistics (/stats/sp/ip)

Statistic	Description
ipInHdrErrors	The number of input datagrams discarded because of errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
ipInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
ipForwDatagrams	The number of input datagrams for which this device was not their final IP destination. As a result, an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter will include only those packets which were source-routed via this device, and the source- route option processing was successful.
ipInUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
ipInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting reassembly.
ipFragOKs	The number of IP datagrams that have been successfully fragmented at this device.
ipFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this device.
ipDontFrags	The number of IP datagrams that needed to be fragmented but had the “Don’t Fragment” bit set.
ipFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this device but could not. For example: because their “Don’t Fragment” flag was set.
ipOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this *no-route* criterion. Note that this includes any datagrams that a host cannot route because all of its default gateways are down.

/stats/sp/maint

Maintenance Statistics for Switch Processor

```
[SP index 0 (port A1) Maintenance Statistics Menu]
      ltr     - Letter counts
      msg     - STEM message statistics
      mem     - STEM memory statistics
      thr     - STEM thread statistics
      cpu     - Show CPU utilization
```

Table 5-26 SP Maintenance Menu Options (/stats/sp/maint)

Command Syntax and Usage
ltr Displays STEM letter count statistics for the selected port. See page 139 for sample output.
msg Displays STEM message statistics for the selected port. See page 140 for sample output.
mem Displays STEM memory statistics for the selected port. See page 140 for sample output.
thr Displays STEM thread statistics for the selected port. See page 141 for sample output.
cpu Displays CPU utilization statistics for the selected port. CPU statistics are for periods of 1, 4, and 64 seconds. Refer to page 141 for a sample output.

NOTE – The SP maintenance submenu and options contain hardware-specific information related to the switch processor. The example screens shown in this section are for informational purposes only. You may be asked by an Alteon WebSystems support engineer to provide SP maintenance information for troubleshooting purposes.

/stats/sp/maint/ltr

Letter Counts for Switch Processor

```
   Letter statistics for SP index 0 (port A1):

    ------------ Received -------------  -------------- Sent ---------------
SP   non-ACKs     ACKs     Dups  NonSeq   non-ACKs     ACKs     Rxmits Defers
--  ---------- ---------- ------ ------  ---------- ---------- ------ ----
MP    21439021     676977      0     82      676977   21439021     32    965

 Letters By Type     Received        Sent        Rxmitted      Deferred
-----------------  ------------  ------------  ------------  ------------
 ACK                   676977      21439021             0             0
 INIT_DONE                  0             1             0             0
 FLAGS_CHANGE               1             0             0             0
 STATS_REQ                 21             0             0             0
 STATS_RSP                  0            20             0             0
 SLB_REQ               292682             0             0             0
 IP_ROUTE_REQ         2531307             0             0             0
 SLOT_INFO                 12             0             0             0
 LINK_DOWN_NOTE             4             0             0             0
 LINK_DOWN_DONE             0             4             0             0
 LINK_UP_NOTE               9             0             0             0
 GDT_UPDATE               493             0             0             0
 EVID_UPDATE               29             0             0             0
 HEARTBEAT                667             0             0             0
 FDB_REQ                  222             0             0             0
 FDB_RSP                    0           164             0             0
 VLAN_CLASS_UPDT            1             0             0             0
 ACTIVE_LINKS               1             0             0             0
 SYSMON_REQ                 2             0             0             0
 FDB_AGE_SCAN        18612550             0             0             0
 FLAGS_CHANGED              0             2             0             0
 SYSMON_DATA                0          7342             0           585
 SLB_RSP                    0        550647            25            47
 IP_ROUTE_RSP               0             2             0             0
 SLB_RATE                  40             0             0             0
 SLB_RCTE                   5             0             0             0
 SLB_RSTE                   5             0             0             0
 SLB_RGTE                   3             0             0             0
 SA_LEARN                   0        117870             7           332
 QL_STATS_REQ             210             0             0             0
 QL_STATS_RSP               0           210             0             1
 CP_FLOW_CTRL             757           715             0             0
```

/stats/sp/maint/msg

STEM Message Statistics for Switch Processor

This command displays the STEM message statistics for the SP. A sample follows:

```
STEM message statistics for SP index 0 (port A1):
-----------------------------------------------------------------
 allocs:          109389   frees:           299552
 callocs:         190163   calloc_fails:         0
 sends:           299553   returns:              0
 alloc_curr:           0   alloc_hiwat:         18
 alloc_fails:          0
```

/stats/sp/maint/mem

STEM Memory Statistics for Switch Processor

This command displays the STEM memory statistics for the switch processor. A sample follows:

```
STEM memory statistics for SP index 0 (port A1):
-----------------------------------------------------------------
 allocs:          15618   frees:            15606
 alloc_fails:         0   pool_bytes:      262144
 bytes_curr:      43008   bytes_hiwat:     190720
 largest:         16384
```

/stats/sp/maint/thr

STEM Thread Statistics for Switch Processor

This command displays the STEM thread statistics for the switch processor. A sample follows:

```
STEM thread statistics for SP index 0 (port A1):
-----------------------------------------------------------------
 thid  name        stack              maximum pathlen
                 size   used      uSecs     cmd      word1
    1   STEM                         0  00000000  00000000
   42   SLB       4076    400       5180  00000002  00000000
   43   IP        4076    384      22850  00090001  003732ec
   44   SLOT      4076   1264       1315  00090001  001f1da0
   45   PORT      4076    320         99  00000001  fdfdfdfd
   46   STAT      4076    128        103  00000001  fdfdfdfd
   47   VLAN      4076    392       1119  00090001  0038bb2c
   48   DIAG      4076    128         69  00000001  fdfdfdfd
   49   FDB       4076    512       5558  00090001  00393e2c
   50   DSTA      4076    128         38  00000001  fdfdfdfd
   51   SMON      4076    360       1730  00000002  00000000
   52   DONE      1004     88         19  00000001  fdfdfdfd
   53   diag      4076     96         19  00000001  fdfdfdfd
```

/stats/sp/maint/cpu

CPU Statistics for Switch Processor

This command displays the CPU statistics for the switch processors. A sample follows:

```
>> SP index 4 (port B1) Maintenance Statistics# cpu
-----------------------------------------------------------------
CPU utilization for port SP index 4 (port B1):
cpUtil1Second:              0%  fpUtil1Second:           17%
cpUtil4Seconds:             0%  fpUtil4Seconds:          17%
cpUtil64Seconds:            0%  fpUtil64Seconds:         17%
```

/stats/dump
Statistics Dump

Use the dump command to dump all switch statistics available from the Statistics Menu (40K or more, depending on your configuration). This data can be used in tuning and debugging switch performance.

If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump commands.

CHAPTER 6

The Configuration Menu

This chapter discusses how to use the Command Line Interface (CLI) for making, viewing, and saving switch configuration changes.

NOTE – After configuring the switch, make sure to use the `apply` command to make the configurations operational, and the `save` command to retain them when the switch is next reset.

Viewing, Applying, and Saving Changes

As you use the configuration menus described in this chapter 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
- Revert, delete all pending changes

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.

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

You can view all pending configuration changes that have been applied but not saved to FLASH memory using the `diff FLASH` command. It is a global command that can be executed from any menu.

For instructions on selecting the configuration to run at the next system reset, see “Selecting a Configuration Block” on page 267.

/cfg
Configuration Menu

```
[Configuration Menu]
      sys     - System-wide configuration menu
      udef    - User-defined defaults menu
      port    - Port configuration menu
      vlan    - VLAN configuration menu
      stp     - Spanning Tree configuration menu
      snmp    - SNMP configuration menu
      trunk   - Trunk Group configuration menu
      ip      - IP configuration menu
      slb     - Layer 4 configuration menu
      vrrp    - VRRP configuration menu
      setup   - Step by step configuration set up
      dump    - Dump current configuration to script file
      putcfg  - Backup current configuration to TFTP server
      getcfg  - Restore current configuration from TFTP server
```

Each menu option is described in Table 6-1, followed by specific descriptions of each menu and submenu.

Table 6-1 Configuration Menu Options (/cfg)

Command Syntax and Usage
sys Displays the System Configuration Menu. To view menu options, see page 147.
udef Displays the User-Defined Defaults Menu. Use this menu to set user-defined default link configurations and user defined Ethertypes. To view menu options, see page 153.
port *<port number, such as b12>* Displays the Port Configuration Menu. To view menu options, see page 157.
vlan *<1-4094>* Displays the VLAN Configuration Menu for the selected VLAN. To view menu options, see page 162.
stp *<1-256>* Displays the Spanning Tree Configuration Menu for the selected group. To view menu options, see page 164.

Table 6-1 Configuration Menu Options (/cfg)

Command Syntax and Usage
snmp Displays the SNMP Configuration Menu. To view menu options, see page 168.
trunk *<1-4>* Displays the Trunk Group Configuration Menu for the selected group. Alteon 708 switches support four trunk groups and up to four wire ports per trunk group. To view menu options, see page 171.
ip Displays the IP Configuration Menu. To view menu options, see page 175.
slb Displays the Server Load Balancing Configuration Menu. To view menu options, see page 209.
vrrp Displays the Virtual Router Redundancy Protocol Configuration Menu. To view menu options, see page 241.
setup Walks you through the configuration of System Date and Time, BOOTP, IP address and subnet mask, Spanning Tree, Port and Link characteristics, and VLANs. For more information, see page 25.
dump Dumps current configuration to a script file.
putcfg *<host name or IP address> <filename on host>* Backs up current configuration to TFTP server.
getcfg *<host name or IP address> <filename on host>* Restores the backed-up configuration from TFTP server. Reads the configuration file into the *new* (pending) configuration block, not the current block. Also, you must apply the changes. See "Viewing, Applying, and Saving Changes" on page 143.

/cfg/sys
System Configuration

```
[System Menu]
      syslog  - Syslog Menu
      radius  - RADIUS Authentication Menu
      user    - User Access Control Menu (passwords)
      sshd    - SSH Server Menu
      date    - Set system date
      time    - Set system time
      idle    - Set timeout for idle CLI sessions
      bootp   - Enables or disables use of BOOTP
      snmp    - Enables or disables SNMP management access
      web     - Enables or disables Web management access
      wport   - Set Web server port number
      mnet    - Set management network
      mmask   - Set management netmask
      telnet  - Enable/Disable telnet access
      banner  - Set login banner
      extbuf  - Enables or disables Extended Buffer Mode
      current - Display current system-wide configuration
```

This menu provides configuration of switch management parameters, such as user and administrator privilege mode passwords, Web-based management settings, and management access list.

Table 6-2 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
syslog Displays the Syslog Configuration menu. This menu controls the level and destination of system log messages. For menu options, see page 149.
radius Displays the RADIUS Authentication menu. For menu options, see page 150.
user Displays the User Access Control menu. For menu options, see page 151.
sshd Displays the Secure Shell (SSH) Server menu. For menu options, see page 151.
date Prompts the user for the system date.

Table 6-2 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
time Prompts the user for the system time using a 24-hour clock format.
idle *<time-out in minutes>* Sets the idle time allowed on console and Telnet sessions to the switch CLI. The range is 1 to 60 minutes. The default is 30 minutes.
bootp disable|enable (**d|e**) Enables or disables the use of BOOTP; if you enable BOOTP, the switch will query its BOOTP server for switch IP parameters, such as address, default gateway, etc.
snmp disable|enable (**d|e**) Enables or disables SNMP-based network management. If disabled, you cannot configure or manage the switch using an SNMP network management station.
web disable|enable (**d|e**) Enables or disables HTTP access to the Web OS Browser-Based Interface (BBI)
wport *<TCP port number (1-65535)>* Sets the switch port to be used for serving switch Web content. The default is HTTP port 80. If Global Server Load Balancing is to be used, change this parameter to use a different port (such as 8080).
mnet *<IP subnet, such as 192.4.17.0>* Sets the base source IP subnet allowed to access switch management through Telnet, SNMP, RIP, or the Web OS Web interface. A range of IP addresses is produced when used with mmask (below). Specify an IP address in dotted decimal notation.
mmask *<IP subnet mask, such as 255.255.0.0>* This IP address mask is used with mnet to set a range of source IP addresses allowed to access switch management functions. Specify the mask in dotted decimal notation.
telnet disable|enable (**d|e**) Enables or disables telnet access to command line interface sessions. This command is only available through a console connection.
banner *<string, maximum 80 characters>* Configures a login banner of up to 80 characters. When a user or administrator logs into the switch via Telnet, the login banner is displayed; it is also displayed as part of the output from the /info/sys command.

Table 6-2 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
extbuf disable|enable (**d|e**) Enables or disables extended buffer mode to handle 64KB burst traffic. Enabling `extbuf` reduces the SLB table entries from 256K to 128K so that the buffer can be increased. To activate enabling or disabling `extbuf`, the switch must be reset. This configuration is saved in NVRAM, not the configuration file.
current Displays the values for most system parameters (excluding passwords).

/cfg/sys/syslog

System Host Log Configuration

```
[Syslog Menu]
      host    - Set IP address of first syslog host
      host2   - Set IP address of second syslog host
      facil   - Set facility of first syslog host
      facil2  - Set facility of second syslog host
      console - Enables or disables console output of syslog messages
      current - Display current syslog settings
```

This menu provides configuration of the system host log, such as the IP address of the hosts, and the level of the syslog messages.

Table 6-3 System Host Log Configuration Menu Options (/cfg/sys/syslog)

Command Syntax and Usage
host *<IP address, such as 192.4.17.0>* Sets the IP address, in dotted decimal notation, of the first host where system log (syslog) messages are sent.
host2 *<IP address, such as 192.4.17.0>* Sets the IP address, in dotted decimal notation, of the second host where syslog messages are sent.
facil *<level (0-7)>* Sets the facility level of syslog messages sent to the first syslog host. The security level identifies the switch to the syslog host.
facil2 *<level (0-7)>* Sets the facility level of syslog messages sent to the second syslog host. The security level identifies the switch to the syslog host.

Table 6-3 System Host Log Configuration Menu Options (/cfg/sys/syslog)

Command Syntax and Usage
console enable|disable (d|e) Enables or disables sending syslog messages to the console port.
current Displays the current syslog configuration parameters.

/cfg/sys/radius

Radius Server Configuration

```
[RADIUS Server Menu]
      prisrv  - Set primary RADIUS server address
      secsrv  - Set secondary RADIUS server address
      secret  - Set RADIUS secret
      port    - Set RADIUS port
      retries - Set RADIUS server retries
      timeout - Set RADIUS server timeout
      telnet  - Enables or disables RADIUS backdoor for telnet
      on      - Turn RADIUS authentication ON
      off     - Turn RADIUS authentication OFF
      current - Display current RADIUS configuration
```

Remote Authentication Dial-In User Interface (RADIUS) is an access server authentication, authorization, and accounting protocol that can be configured to secure remote access to networks and network services against unauthorized access.

Table 6-4 Radius Server Menu Options (/cfg/sys/radius)

Command Syntax and Usage
prisrv *<RADIUS server address, such as 192.4.17.101>* Sets the primary RADIUS server address.
secsrv *<RADIUS server address, such as 192.4.17.101>* Sets the secondary RADIUS server address.
secret *<1-32 character string>* This is the shared secret between the switch and the RADIUS server(s). The string length can be from 1 to 32 characters.
port *<RADIUS port (1500-3000)>* Sets the number of the UDP port. The default value is 1645.

Table 6-4 Radius Server Menu Options (/cfg/sys/radius)

Command Syntax and Usage
retries *<1-3>* Sets the number of failed authentication requests before switching to a different RADIUS server. The default is 3 retries.
timeout *<RADIUS server time-out seconds (1-10)>* Sets the amount of time, in seconds, before a RADIUS server authentication attempt is considered to have failed. The default is 3 seconds.
telnet disable|enable (**d**|**e**) Enables or disables the RADIUS backdoor for telnet. `telnet` also applies to SSH/SC connections. The default setting is enabled.
on Enables the RADIUS server.
off Disables the RADIUS server.
current Displays the current Radius configuration.

/cfg/sys/user

User Access Control Configuration

```
[User Access Control Menu]
      usrpw   - Set user password (user)
      sopw    - Set SLB operator password (slboper)
      l4opw   - Set L4 operator password (l4oper)
      opw     - Set operator password (oper)
      sapw    - Set Slb administrator password (slbadmin)
      l4apw   - Set L4 administrator password (l4admin)
      admpw   - Set administrator password (admin)
      current - Display current user statis
```

Changing any password requires validation by entering the current administrator password, and verification by entering the new password twice. The maximum length of a password is 15 characters. To disable a user, set the password to NULL.

Table 6-5 User Access Menu Options (/cfg/sys/user)

Command Syntax and Usage
usrpw Sets the user password. The default password is `user`. This user level is always enabled.
sopw Sets the SLB operator password. The default password is NULL—disabled. This user level will be disabled until a password is assigned.
l4opw Sets the L4 operator password. The default password is NULL—disabled. This user level will be disabled until a password is assigned.
opw Sets the operator password. The default password is NULL—disabled. This user level will be disabled until a password is assigned.
sapw Sets the SLB administrator password. The default password is NULL—disabled. This user level will be disabled until a password is assigned.
l4apw Sets the Layer 4 administrator password. The default password is NULL—disabled. This user level will be disabled until a password is assigned.
admpw Sets the administrator password. The default password is `admin`. The administrator user level is always enabled
current Displays the current user status—enabled, always enabled, or disabled. The administrator and user levels are always enabled.

/cfg/sys/sshd

SSH Server Configuration

```
[SSHD Menu]
      current - Display current SSH server configuration.
```

Table 6-6 SSH Server Configuration Menu Options (/cfg/sys/sshd)

Command Syntax and Usage
current Displays the current SSH configuration.

A sample display of current SSH configuration follows.

```
>> SSHD# current
RSA server key autogen disabled
RSA host key currently not ready to service
RSA server key currently not ready to service
SCP apply and save currently disabled
SSH server currently OFF
Note: There are pending config changes; use "diff" to see them.
```

/cfg/udef
User-Defined Defaults Configuration

```
[User-defined defaults Menu]
      link    - Link defaults menu
      etype   - User-defined Ethertype menu
```

Use this menu to set up to two user-defined default link configurations.

Table 6-7 User-Defined Defaults Configuration Menu Options (/cfg/udef)

Command Syntax and Usage
link FE|GE-SX|1000Base-T Displays the Link Defaults Menu. This menu enables you to create user-defined default link configurations, for Fast Ethernet, 1000Base-SX Gigabit Ethernet (GE-SX), or 1000Base-T Gigabit Ethernet over Copper. To view menu options, see "User-Defined Link Defaults Configuration" on page 155.
etype *<Ethertype index (1-10)>* Displays the user-defined Ethertype Menu. This menu enables you to configure up to 10 user-defined Ethertypes per switch for PVID configuration. To view menu options, see page 156.

/cfg/udef/link <FE|GE-SX|1000Base-T>

User-Defined Link Defaults Configuration

```
[FE|GE-SX|1000Base-T Default Menu]
      speed   - Set link speed
      mode    - Set full or half-duplex mode
      fctl    - Set flow control
      auto    - Control autonegotiation
      current - Display current link configuration
```

User-defined link types can be used to globally configure links for a particular type of Line Card Module. For example, instead of configuring each port, you can choose to set all Fast Ethernet port links to run at 100 Mbps. Then, if any Fast Ethernet LCM is inserted, it will always run at this speed.

However, you may need to make an exception for one or more ports. For example, if connecting port b12 to a laptop that requires a 10 Mbps link, you would configure the link speed for port b12 using the menu for "Port Link Configuration" on page 158.

Table 6-8 User-Defined Link Defaults Menu Options (/cfg/udef/link)

Command Syntax and Usage
speed 10|100|1000|any (not all options are valid on all ports) Sets the link speed; the choices include: ■ "Any" for automatic detection (default) ■ 10 Mbps ■ 100 Mbps ■ 1000 Mbps
mode The operating mode is set for full duplex.(This parameter is fixed).
fctl Flow control is set for none. (This parameter is fixed).
auto on|off Enables or disables autonegotiation for the port.
current Displays current port parameters.

/cfg/udef/etype <index>

Ethertype Defaults Configuration

NOTE – The *index* (1-10) represents the number of the Ethertype you wish to configure.

```
[User-defined Ethertype[1] Menu]
      ether   - Specify Ethertype value
      current - Display current Ethertype value
      dump    - Display all user-defined Ethertypes
```

For Port VLAN ID (PVID) configuration, this menu enables you to configure up to 10 user-defined Ethertypes per switch.

In the following example, to set an Ethertype for an AppleTalk network, enter the Ethertype value for AppleTalk.

```
>> User-defined Ethertype[1]# ether 809b
Current user-defined Ethertype[1]:      undefined
New pending user-defined Ethertype[1]:  0x809b
```

Table 6-9 Ethertype Defaults Menu Options (/cfg/udef/etype)

Command Syntax and Usage
ether 0800-ffff|0 Specifies the Ethertype value in the frame. The values 0800-ffff identify the range of Ethertype values. The value 0 deletes the value.
current Displays the current Ethertype value.
dump Displays all user-defined Ethertypes.

/cfg/port <slot-number port-number >
Port Configuration

```
[Port A1 Menu]
      link    - Link menu ("link none" to reset link type)
      pvid    - Port VLAN ID menu
      tag     - Forward/Discard all VLAN tagged frames
      utag    - Forward/Discard all untagged frames
      ptag    - Forward/Discard all priority tagged frames
      rmon    - Enable/Disable RMON for port
      enable  - Enable port
      disable - Disable port
      current - Display current port configuration
```

The Port Menu enables you to configure settings for individual switch ports.

Table 6-10 Port Configuration Menu Options (/cfg/port)

Command Syntax and Usage
link none|FE|GE-SX|1000Base-T Displays the Link Menu for the specified port. Enter the new link type `none|FE|GE-SX|1000Base-T` to display the link menu for that type of link. `none` resets the existing link. To view menu options, see page 158.
pvid Displays the Port VLAN ID Menu, which is used to configure default PVIDs for each port. The PVID is the default VLAN number that will be used to forward frames that are not VLAN-tagged. The factory default value is 1. To view menu options, see page 160.
tag forward|discard (**f|d**). Forwards or discards all VLAN tagged frames.
utag forward|discard (**f|d**) Forwards or discards all untagged frames.
ptag forward|discard (**f|d**) Forwards or discards all priority tagged frames.
rmon disable|enable (**d|e**) Enables or disables RMON support on the port.
enable Enables the port.

Table 6-10 Port Configuration Menu Options (/cfg/port)

Command Syntax and Usage
disable Disables the port.
current Displays current port configuration.

/cfg/port <slot-number port-number>/link <link type>

Port Link Configuration

```
[Port A1 Link Menu]
      speed   - Set link speed
      mode    - Set full or half duplex mode
      fctl    - Set flow control
      auto    - Control autonegotiation
      current - Display current link configuration
```

The port link configuration menu enables you to configure the link speed and autonegotiation for a port.

After you have configured a port, it will operate at the configured settings, even if you swap the Line Card Module (LCM) type: for example, if you have already configured port a1 of a Fast Ethernet LCM to operate at 100 Mbps and want to swap it with a Gigabit-SX LCM.

After the Gigabit LCM has come online, a syslog message will indicate that the default link speed of the current hardware is different from the configured link speed. However, the port will continue to operate at 100 Mbps until you specify otherwise. Example of a syslog error message:

```
Mar 210:10:42 ALERT slot mgr-unexpected LCM type in slot C
```

NOTE – Both ends of a port link must be configured at the same speed, mode, and autonegotiation settings. For example, if port a1 is connected to another switch, router, Web server or hub, the port link settings for that device must match the settings for port a1.

Use menu options below to set port parameters for the port link.

Table 6-11 Port Link Configuration Menu Options (/cfg/port/link)

Command Syntax and Usage
speed 10|100|1000|any (not all options are valid on all ports) Sets the link speed; the choices include: ■ Any for automatic detection (default) ■ 10 Mbps ■ 100 Mbps ■ 1000 Mbps
mode The operating mode is set to full duplex.
fctl The flow control is set to none.
auto on|off (not all options are valid on all ports) Enable or disable autonegotiation for the port.
current Displays current port parameters.

/cfg/port <slot-number port-number>/pvid
Port VLAN ID (PVID) Configuration

```
[PVID Menu]
      ip      - Set PVID for IP/ARP/RARP
      ipx     - Set PVID for IPX
      user    - Set PVID for user-specified Ethertypes menu
      other   - Set PVID for all other Ethertypes
      current - Display current PVID configuration
```

Use this menu to configure default Port VLAN IDs (PVIDs) for ports. There can be up to eight PVIDs per port: IP, IPX, and five user-defined Ethernet types as well as the PVID that can be configured for each port to classify untagged frames that do not match the other Ethertypes. The factory default values for the PVIDs is 1.

The Alteon 708 switches allow you to configure PVIDs for each different Ethernet type.

Table 6-12 Port VLAN ID Configuration Menu Options (/cfg/port/pvid)

Command Syntax and Usage
ip *<VLAN number (1-4094)>* Sets the PVID for Internet Protocol (IP), Address Resolution Protocol (ARP), Reverse Address Resolution Protocol (RARP).
ipx *<VLAN number (1-4094)>* Sets the PVID for IPX.
user *<user PVID index number (1-5)>* This menu is for setting the PVID for one user specified Ethertype. To view menu options, see page 161.
other *<VLAN number (1-4094)>* Sets the PVID for all other Ethertypes.
current Displays the current PVIDs for the port.

/cfg/port <slot-number port-number>/pvid/user <pvid number>

User PVID Port Configuration

NOTE – The *PVID number* (1-5) represents the index you wish to configure.

```
[User PVID 1 Menu]
      etype   - Specify Ethertype index
      vlan    - Specify VLAN ID
      undef   - Undefine this user PVID
      current - Display current user PVID configuration
```

This menu allows you to configure the Port VLAN identification.

Table 6-13 User PVID Port Configuration Menu Options(/cfg/port/pvid/user)

Command Syntax and Usage
etype *<(1-10)>* Specifies the Ethertype index. The default value is *undefined*.
vlan *<(1-4094)>* Specifies the VLAN ID.
undef Undefines the user PVID—`etype` is reset to *undefined*.
current Displays the current Ethertype and VLAN for this user PVID.

/cfg/vlan <vlan number>
VLAN Configuration

NOTE – The *group number* (1-1024) represents the number of the VLAN you wish to configure. A maximum of 1024 VLANs can be configured.

```
[VLAN 1 Menu]
      name    - Assign VLAN name
      stg     - Assign VLAN to a Spanning Tree Group
      add     - Add port(s) or trunk group(s) to VLAN
      remove  - Remove port(s) or trunk group(s) from VLAN
      clear   - Remove all ports and trunk groups from VLAN
      enable  - Enable VLAN
      disable - Disable VLAN
      delete  - Delete VLAN
      current - Display current VLAN configuration
```

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 more information on configuring VLANs, see "Setup Part 3: VLANs" on page 32.

NOTE – You cannot add a port to more than one VLAN unless the port has VLAN tagging turned on (see "Port Configuration" on page 157).

NOTE – All ports must belong to at least one VLAN. Any port that is removed from a VLAN and that 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.

Table 6-14 VLAN Configuration Menu Options(/cfg/vlan)

Command Syntax and Usage
name *<text name>* Assigns the name to the VLAN, maximum 32 characters. The default name is VLAN.
stg *<(1–256)>* Assigns VLAN to a Spanning Tree Group. The default value is 1.
add *<slot and port, such as b12>* Adds port(s) or trunk group(s) to the VLAN. One port or trunk group is added per line. End the string with NULL.
remove *<slot and port, such as b12>* Removes selected port(s) or trunk group(s) from the VLAN. One port or trunk group is removed per line. End the string with NULL.
clear Removes all ports and trunk groups from the VLAN.
enable Enables this VLAN.
disable Disables this VLAN.
delete Deletes this VLAN.
current Displays the current configuration values that can be set from this menu.

/cfg/stp <Spanning Tree Group index>
Spanning Tree Configuration

NOTE – The *STP index* (1-256) represents the number of the Spanning Tree Group you wish to configure.

```
[Spanning Tree Group 1 Menu]
      brg     - Bridge parameter menu
      port    - Port parameter menu
      add     - Add VLAN(s) to Spanning Tree Group
      remove  - Remove VLAN(s) from Spanning Tree Group
      clear   - Remove all VLANs from Spanning Tree Group
      on      - Globally turn Spanning Tree ON
      off     - Globally turn Spanning Tree OFF
      default - Default Spanning Tree and Member parameters
      current - Display current Spanning Tree configuration
```

Web OS 7.2 supports the IEEE 802.1d Spanning Tree Protocol (STP). STP is used to prevent loops in the network topology. Up to 256 Spanning Trees can be configured per switch.

NOTE – When Virtual Router Redundancy Protocol (VRRP) is used for active/active redundancy, STP must be enabled.

Table 6-15 Spanning Tree Configuration Menu Options(/cfg/stp)

Command Syntax and Usage
brg Displays the Bridge Spanning Tree Menu. To view menu options, see page 165.
port *<port as slot and port, such as b12>* Displays the Spanning Tree Port Menu. To view menu options, see page 167.
add *<VLAN number (1-4094)>* Adds a VLAN to this Spanning Tree Group.
remove *< VLAN number (1-4094)>* Removes a VLAN from a Spanning Tree Group.
clear Removes all VLANs from a Spanning Tree Group.

Table 6-15 Spanning Tree Configuration Menu Options(/cfg/stp)

Command Syntax and Usage
on Globally enables the Spanning Tree Group.
off Globally disables the Spanning Tree Group.
default Displays the default Spanning Tree and Member parameters.
cur Displays the current Spanning Tree parameters set by this menu.

/cfg/stp <*index*>/brg
Bridge Spanning Tree Configuration

NOTE – The *index* (1-256) represents the number of the Bridge STP you wish to configure.

```
[Bridge Spanning Tree Menu]
      prior   - Set bridge Priority [0-65535]
      hello   - Set bridge Hello Time [1-10 secs]
      maxage  - Set bridge Max Age [6-40 secs]
      frwd    - Set bridge Forward Delay [4-30 secs]
      aging   - Set bridge Aging Time [10-65535 secs]
      current - Display current bridge parameters
```

This menu allows you to configure the STP bridge. STP bridge parameters affect the global STP operation of the switch. The parameters include the following:

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

Table 6-16 Bridge Spanning Tree Menu Options (/cfg/stp/brg)

Command Syntax and Usage
prior *<bridge priority (0-65535)>* 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 *<seconds (1-10)>* Configures the bridge hello time in seconds.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.
maxage *<seconds (6-40) >* Configures the bridge maximum age in seconds. 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.
frwd *<seconds (4-30) >* 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 *<seconds (1-65535|0)>* 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 10 to 65535 seconds, and the default is 300 seconds. To disable aging, set this parameter to 0.
current Displays the current bridge STP parameters set by this menu.

When configuring STP bridge parameters, use these formulas:

- 2*(*frwd*-1) $\geq$ *maxage*
- 2*(*hello*+1) $\leq$ *maxage*

/cfg/stp <index>/port < port or trunk group>
Spanning Tree Port Configuration

NOTE – The *index* (1-256) and the *port/trunk group* represent the STP you wish to configure.

```
[Spanning Tree Port A1 Menu]
      prior   - Set port Priority (0-255)
      cost    - Set port Path Cost (1-65535, 0 for default)
      on      - Turn port's Spanning Tree ON
      off     - Turn port's Spanning Tree OFF
      current - Display current port Spanning Tree parameters
```

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

- Port priority
- Port path cost

Table 6-17 Spanning Tree Port Menu Options(/cfg/stp/port)

Command Syntax and Usage
prior *<new port priority (0-255)>* 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 *<new port path cost (1-65535|0 for default)>* 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 100 Mbps ports, and 1 for gigabit ports. A value of 0 indicates that the cost will be computed for an autonegotiated link speed—this value is default.
on Enables STP on the port.
off Disables STP on the port.
current Displays the current STP port parameters set by this menu.

/cfg/snmp
SNMP Configuration

```
[SNMP Menu]
      name    - Set SNMP "sysName"
      locn    - Set SNMP "sysLocation"
      contact - Set SNMP "sysContact"
      rcomm   - Set SNMP read community string
      wcomm   - Set SNMP write community string
      trap    - SNMP trap host configuration menu
      auth    - Disable/enable SNMP "sysAuthenTrap"
      linkt   - Disable/enable SNMP link up/down trap
      current - Display current SNMP information
```

The Web OS 7.2 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)

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

Table 6-18 SNMP Configuration Menu Options (/cfg/snmp)

Command Syntax and Usage
name *<new string, maximum 64 characters>* Configures the name for the system. The name can have a maximum of 64 characters.
locn *<new string, maximum 64 characters>* Configures the system location information (for example, “Bldg. 1"). The system location can have a maximum of 64 characters.
contact *<new string, maximum 64 characters>* Configures the system contact information (for example, a phone number). The system contact can have a maximum of 64 characters.
rcomm *<new SNMP read community string, maximum 32 characters>* 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 *<new SNMP write community string, maximum 32 characters>* 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.
trap *<trap host index* ***1***|***2****>* Displays the SNMP Trap Host Configuration Menu. The SNMP trap host is the device that receives SNMP trap messages from the switch. To view menu options, see page 170.
auth disable|enable (**d**|**e**) Enables or disables the use of the system authentication trap facility. The default setting is disabled.
linkt *<port>* **disable|enable** (**d**|**e**) Enables or disables the sending of SNMP link up and link down traps for the given port.
current Displays the current STP port parameters set by this menu.

/cfg/snmp/trap <trap host number>

Trap Host SNMP Configuration

NOTE – The *trap host number (*1-2) represents the trap host you wish to configure.

```
[Trap host 1 Menu]
      addr    - Set SNMP trap host IP address
      comm    - Set community string for SNMP trap hosts
```

This menu allows you to configure the IP address and the community string of the SNMP trap host.

Table 6-19 Trap Host SNMP Configuration Menu Options (/cfg/snmp/trap)

Command Syntax and Usage
addr *<new SNMP trap host IP address, such as 192.4.17.101>* Configures the IP address of this SNMP trap host using dotted decimal notation. The SNMP trap host is the device that receives SNMP trap messages from the switch.
comm *<new trap host community string>* Configures the community string for this trap host. The maximum allowed string is 32 characters. The default value is `public`. To set `comm` to NULL, enter the string to `none`.

/cfg/trunk <group number>
Trunk Configuration

NOTE – The *group number (*1-4) represents the trunk group you wish to configure.

```
[Trunk Group 1 Menu]
      add     - Add port(s) to trunk group
      remove  - Remove port(s) from trunk group
      pvid    - Default port VLAN id configuration menu
      tag     - Forward/Discard all VLAN tagged frames
      utag    - Forward/Discard all untagged frames
      ptag    - Forward/Discard all priority tagged frames
      enable  - Enable trunk group
      disable - Disable trunk group
      current - Display current trunk group configuration
```

A *trunk* is a group of ports that act together, combining their bandwidth to create a single, larger port. The following restrictions apply.

- Any physical switch port can belong to no more than one trunk group.
- Up to four trunk groups can be configured per switch.
- Up to four ports can belong to the same trunk group, with the bandwidth at 100 Mbps or 1000 Mbps.
- Best performance is achieved when all ports in any given trunk group are configured for the same speed.
- Trunking from non-Alteon WebSystems devices must comply with Cisco® EtherChannel® technology.

Table 6-20 Trunk Configuration Menu Options (/cfg/trunk)

Command Syntax and Usage
add *<port as slot and port, such as b12>* Adds a physical port to the current trunk group.
remove *<port as slot and port, such as b12>* Removes a physical port from the current trunk group.
pvid Displays the Default Port VLAN ID configuration menu. To view menu options, see page 173.

Table 6-20 Trunk Configuration Menu Options (/cfg/trunk)

Command Syntax and Usage
`tag forward|discard` (**f**|**d**) Forwards or discards all VLAN-tagged frames.
`utag forward|discard` (**f**|**d**) Forwards or discards all untagged frames.
`ptag forward|discard` (**f**|**d**) Forwards or discards all priority tagged frames.
`enable` Enables the current trunk group.
`disable` Disables the current trunk group.
`current` Displays current trunk group parameters set by this menu.

/cfg/trunk <trunk group number>/pvid

Port VLAN ID Trunk Configuration

```
[PVID Menu]
      ip      - Set PVID for IP/ARP/RARP
      ipx     - Set PVID for IPX
      user    - Set PVID for user-specified Ethertypes menu
      other   - Set PVID for all other Ethertypes
      current - Display current PVID configuration
```

Table 6-21 PVID Trunk Configuration Menu Options (/cfg/trunk/pvid)

Command Syntax and Usage
ip *<VLAN number (1-4094)>* Sets the PVID for IP/ARP/RARP. The default value is 1.
ipx *<VLAN number (1-4094)>* Sets the PVID for IPX. The default value is 1.
user *<user PVID index number (1-5)>* Displays the PVID Configuration Menu for user defined Ethertypes. To view menu options, see page 174.
other *<VLAN number (1-4094)>* Sets the PVID for all other Ethertypes.
current Displays current PVID configuration set by this menu.

/cfg/trunk <trunk group number>/pvid/user <user PVID index number>

User PVID Port Configuration Menu

NOTE – The *trunk group number* (1-5) and the PVID index number (1-5) represent the PVID index you wish to configure.

```
[User PVID 1 Menu]
      etype   - Specify Ethertype index
      vlan    - Specify VLAN ID
      undef   - Undefine this user PVID
      current - Display current user PVID configuration
```

Table 6-22 PVID Trunk Configuration Menu Options (/cfg/trunk/pvid)

Command Syntax and Usage
etype *<Ethertype index (1-10)>* Specifies the Ethertype index for this PVID. This menu is used with etype commands to configure Ethertypes. For information about Ethertype commands, see page 156.
vlan *<VLAN number (1-4094)>* Specifies the VLAN ID for this Ethertype.
undef Sets the user PVID to disabled.
current Displays current configuration for the PVID set by this menu.

/cfg/ip
IP Configuration

```
[IP Menu]
      if       - Interface menu
      gw       - Default gateway menu
      route    - Static route menu
      frwd     - Forwarding menu
      nwf      - Network Filters menu
      rmap     - Route Map menu
      rip      - Routing Information Protocol (RIP) menu
      ospf     - Open Shortest Path First (OSPF) menu
      bgp      - Border Gateway Protocol (BGP) menu
      port     - Port menu
      dns      - Domain Name System menu
      bootp    - Bootstrap Protocol Relay Menu
      as       - Set Autonomous System (AS) number
      rtrid    - Set router ID
      metric   - Set default gateway metric
      rearp    - Set re-ARP period in minutes
      current  - Display current IP configuration
```

Table 6-23 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage
if *<interface number (1-1024)>* Displays the IP Interface Menu. To view menu options, see page 177.
gw *<default gateway number (1-4)>* Displays the IP Default Gateway Menu. To view menu options, see page 178.
route Displays the IP Static Route Menu. To view menu options, see page 179.
frwd Displays the IP Forwarding Menu. To view menu options, see page 180.
nwf *<1-256>* Displays the Network Filters Menu. To view menu options, see on page 181
rmap *<1-32>* Displays the Route Map Menu. To view menu options, see page 182.
rip Displays the Routing Interface Protocol Menu. To view menu options, see page 186.

Table 6-23 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage
ospf Displays the Open Shortest Path First (OSPF) Configuration menu. For menu options, see page 189.
bgp Displays the BGP configuration menu. For menu options, see page 200.
port *<port as slot and port, such as b12>* Displays the IP Port Menu. To view menu options, see page 205.
dns Displays the IP Domain Name System Menu. To view menu options, see page 206.
bootp Displays the BOOTSTRAP Protocol menu. The view menu options, see page 207.
as *<1-65535>* Sets the Autonomous System (AS) number.
rtrid *<router ID, such as 10.10.10.1>* Sets a static router ID for this switch. The router ID is expressed in dotted decimal IP address format. Each BGP peer requires a unique router ID to identify itself from the peers. OSPF, when enabled, uses the router ID to identify the routing device. If no router ID is specified, or if the router ID is set to 0.0.0.0 and the switch is rebooted, the switch dynamically selects one of the active IP interfaces on the switch as the router ID. When using OSPF virtual links, the router ID must be set. See the *Web OS 7.2 Application Guide* for more information about OSPF.
metric strict|round robin (s|r) Sets the default gateway metric to `strict` or `roundrobin`. For more information, see page 208.
rearp *<minutes (2-120)>* Sets the re-ARP period. The range is 2 to 120 minutes. The default value is 10. The switch periodically sends ARP (Address Resolution Protocol) requests to refresh its address cache. This command is used for setting the interval between ARP refresh requests.
current Displays the current IP configuration set by this menu.

/cfg/ip/if <interface number>

IP Interface Configuration

NOTE – The *interface number (*1-1024) represents the IP interface you wish to configure.

```
[IP Interface 1 Menu]
      addr    - Set IP address
      mask    - Set subnet mask
      bcast   - Set broadcast address
      vlan    - Set VLAN number
      enable  - Enable interface
      disable - Disable interface
      delete  - Delete interface
      current - Display current interface configuration
```

The switch can be configured with up to 1024 Internet Protocol (IP) interfaces. Each IP interface represents the switch on an IP subnet on your network.

Table 6-24 IP Interface Menu Options (/cfg/ip/if)

Command Syntax and Usage
addr *<IP address, such as 192.4.17.101>* Configures the IP address of the switch interface using dotted decimal notation.
mask *<IP subnet mask, such as 255.255.0.0>* Configures the IP subnet address mask for the interface using dotted decimal notation.
bcast *<broadcast address, such as 192.4.17.255>* Configures the IP broadcast address for the interface using dotted decimal notation.
vlan *<vlan number (1-4094)>* Configures the VLAN number for this interface. Each interface can belong to only one VLAN; however, any VLAN can have multiple IP interfaces in it.
enable Enables this interface.
disable Disables this interface.
delete Deletes this interface.
current Displays the current interface settings.

/cfg/ip/gw <gateway number>

Default IP Gateway Configuration

NOTE – The *gateway number* (1-4) represents the gateway you wish to configure.

```
[Default gateway 1 Menu]
      addr    - Set IP address
      inter   - Set interval between ping attempts
      retry   - Set number of failed attempts to declare gateway DOWN
      arp     - Enables or disables ARP only health checks
      enable  - Enable default gateway
      disable - Disable default gateway
      delete  - Delete default gateway
      current - Display current default gateway configuration
```

Table 6-25 Default IP Gateway Menu Options (/cfg/ip/gw)

Command Syntax and Usage
addr *<default IP gateway address, such as 192.4.1.1>* Configures the IP address of the default IP gateway using dotted decimal notation.
inter *<0-60>* The value range is from 0 to 60 seconds. The switch pings the default gateway to verify that the gateway is up. The **inter** option lets you choose the time between health checks. The default interval is 2 seconds. Setting the interval to 0 disables gateway pinging, and the interface is always assumed to be up.
retry *<1-120>* 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 eight attempts.
arp disable|enable (**d**|**e**) Enables or disables ARP-only health checks.
enable Enables the gateway for use.
disable Disables the gateway.
delete Delete this gateway from the configuration.
current Displays the current gateway settings.

/cfg/ip/route

IP Static Route Configuration

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

Table 6-26 IP Static Route Menu Options (/cfg/ip/route)

Command Syntax and Usage
add *<destination> <mask> <gateway> <interface number (1-1024)>* Adds a static route. You will be prompted to enter a destination IP address, destination subnet mask, gateway address, and interface number. Enter all addresses using dotted decimal notation.
remove *<destination> <mask>* Removes a static route. Specify the destination address and mask of the route to remove, using dotted decimal notation.
current Displays the current IP static routes.

/cfg/ip/frwd

IP Forwarding Configuration

```
[IP Forwarding Menu]
      dirbr   - Enables or disables forwarding directed broadcasts
      on      - Globally turn Forwarding ON
      off     - Globally turn Forwarding OFF
      current - Display current IP Forwarding configuration
```

The IP Forwarding Menu is used for setting IP forwarding (routing) on or off.

Table 6-27 IP Forwarding Menu Options (/cfg/ip/frwd)

Command Syntax and Usage
dirbr disable|enable (**d|e**) Enables or disables forwarding directed broadcasts.
on Enables IP forwarding (routing).
off Disables IP forwarding (routing).
current Displays the current IP forwarding settings.

/cfg/ip/nwf <filter number>

IP Network Filters Configuration

NOTE – The *filter number* (1-256) represents the network filter you wish to configure.

```
[IP Network Filter  1 Menu]
      addr    - IP Address
      mask    - IP Subnet mask
      enable  - Enable Network Filter
      disable - Disable Network Filter
      delete  - Delete Network Filter
      current - Display current Network Filter configuration
```

Network filters work with route mapping and Open Shortest Path First (OSPF) redistribution.

Table 6-28 IP Network Filter Menu Options (/cfg/ip/nwf)

Command Syntax and Usage
addr *<IP address, such as 192.4.17.44>* Sets the starting IP address for this filter. The default address is 0.0.0.0.
mask *<subnet mask, such as 255.255.255.0>* Sets the IP subnet mask that is used with `/cfg/ip/nwf/addr` to define the range of IP addresses that will be accepted by the peer when the filter is enabled. The default value is 0.0.0.0. ■ For Border Gateway Protocol (BGP), assign the network filter to a route map, then assign the route map to the peer.
enable Enables the network filter.
disable Disables the network filter.
delete Deletes the network filter.
current Displays the current network filter configuration.

/cfg/ip/rmap <map number>

Routing Map Configuration

NOTE – The *map number* (1-32) represents the routing map you wish to configure.

```
[IP Route Map 1 Menu]
      alist   - Access List number
      aspath  - AS Filter Menu
      ap      - Set as-path prepend of the matched route
      lp      - Set local-preference of the matched route
      metric  - Set metric of the matched route
      type    - Set OSPF metric-type of the matched route
      prec    - Set the precedence of this route map
      weight  - Set weight of the matched route
      enable  - Enable route map
      disable - Disable route map
      delete  - Delete route map
      current - Display current route map configuration
```

Routing maps control and modify routing information.

Table 6-29 Routing Map Menu Options (/cfg/ip/rmap)

Command Syntax and Usage
alist *<1-8>* Displays the Access List menu. For more information, see page 184.
aspath *<1-8>* Displays the Autonomous System (AS) Filter menu. For more information, see page 185.
ap *<AS number> [<AS number>] [<AS number>]* Sets the AS path preference of the matched route. One to three path preferences can be configured.
lp *<(0-4294967294)*|**none**>* Sets the local preference of the matched route, which affects both inbound and outbound directions. The path with the higher preference is preferred.
metric *<(0-4294967294)*|**none**>* Sets the metric of the matched route.

Table 6-29 Routing Map Menu Options (/cfg/ip/rmap)

Command Syntax and Usage
type <**1**|**2**|**none**>
Assigns the type of OSPF metric. The default is type 1. ■ Type 1—External routes are calculated using both internal and external metrics. ■ Type 2—External routes are calculated using only the external metrics. Type 2 routes have more cost than Type 2. ■ none—Removes the OSPF metric.
prec *<1-255>*
Sets the precedence of the route map. The smaller the value, the higher the precedence.
weight *<(0-65534)*|**none**>
Sets the weight of the route map.
enable
Enables the route map.
disable
Disables the route map.
delete
Deletes the route map.
current
Displays the current route configuration.

/cfg/ip/rmap <rmap number>/alist <number>

IP Access List

NOTE – The *rmap number (*1-32) and the *alist number* (1-8) represent the IP access list you wish to configure.

```
[IP Access List 1 Menu]
      nwf     - Network Filter number
      metric  - Metric
      action  - Set Network Filter action
      enable  - Enable Access List
      disable - Disable Access List
      delete  - Delete Access List
      current - Display current Access List configuration
```

Table 6-30 IP Access List Menu Options (/cfg/ip/rmap/alist)

Command Syntax and Usage

nwf *<filter (1-256)>*

Sets the network filter number.

metric *<(1-4294967294)*|none>

Sets the metric value in the AS-External (ASE) LSA.

action <permit|deny (p|d)>

Permits or denies action for the access list.

enable

Enables the access list.

disable

Disables the access list.

delete

Deletes the access list.

current

Displays the current Access List configuration.

/cfg/ip/rmap <rmap>/aspath <path>

Autonomous System Filter Path

NOTE – The *rmap number (*1-32) and the *path number* (1-8) represent the AS path you wish to configure.

```
[AS Filter 1 Menu]
      as      - AS number
      action  - Set AS Filter action
      enable  - Enable AS Filter
      disable - Disable AS Filter
      delete  - Delete AS Filter
      current - Display current AS Filter configuration
```

Table 6-31 AS Filter Menu Options (/cfg/ip/rmap/aspath)

Command Syntax and Usage

as *<1-8>*

Sets the AS path number.

action *<permit|deny (p|d>*

Permits or denies AS filter action.

enable

Enables the AS filter.

disable

Disables the AS filter.

delete

Deletes the AS filter.

current

Displays the current AS configuration.

/cfg/ip/rip

Routing Information Protocol Configuration

```
[Routing Information Protocol Menu]
      if     - RIP Interface menu
      update  - Set update period in seconds
      on      - Globally turn RIP ON
      off     - Globally turn RIP OFF
      current - Display current RIP configuration
```

Table 6-32 Routing Information Protocol Menu Options (/cfg/ip/rip)

Command Syntax and Usage
if *<interface number (1-1024)>* Displays the RIP Interface Menu. To view menu options, see page 187.
update *<update period (1-120)>* Sets the RIP update period in seconds. The default is 30 seconds.
on Globally turns RIP on.
off Globally turns RIP off.
current Displays the current Routing Information Protocol (RIP) configuration.

/cfg/ip/rip/if <interface number>

Routing Information Protocol Interface Configuration

NOTE – The *map number* (1-1024) represents the RIP interface you wish to configure.

```
[RIP Interface 1 Menu]
      version - Set RIP version
      supply  - Enables or disables supplying route updates
      listen  - Enables or disables listening to route updates
      default - Set default route action
      poison  - Enables or disables poisoned reverse
      trigg   - Enables or disables triggered updates
      mcast   - Enables or disables multicast updates
      metric  - Set metric
      auth    - Set authentication type
      key     - Set authentication key
      current - Display current RIP interface configuration
```

The RIP Interface Menu is used for configuring Routing Information Protocol parameters.

NOTE – Do not configure RIPv2 parameters if your routing equipment uses RIPv1.

Table 6-33 RIP Interface Menu Options (/cfg/ip/rip/if)

Command Syntax and Usage
version *<1|2>* Sets the version of the Routing Information Protocol, RIPv1 or RIPv2.
supply disable|enable (d|e) When enabled, the switch supplies routes to other routers.
listen disable|enable (d|e) When enabled, the switch learns routes from other routers.
default none|listen|supply|both ■ When `listen` is enabled, the switch accepts default routes from other routers and gives them priority over configured default gateways. ■ When `supply` is enabled, the switch advertises RIP default routes to other routers. ■ When `both` is enabled, the switch accepts and advertises RIP default routes. ■ When `none` is enabled, the switch rejects RIP default routes.

Table 6-33 RIP Interface Menu Options (/cfg/ip/rip/if)

Command Syntax and Usage
poison disable|enable (d|e) When enabled, the switch uses split horizon with poisoned reverse. When disabled, the switch uses only split horizon.
trigg disable|enable (d|e) Enables or disables triggered updates.
mcast disable|enable (d|e) Enables or disables multicast updates. This parameter can only be configured in RIPv2.
metric *<1-15>* Sets the metric.
auth none|password Sets the authentication type. This parameter can only be configured in RIPv2.
key *<key>* Sets the authentication key. This parameter can only be configured in RIPv2.
current Displays the current RIP settings.

/cfg/ip/ospf

Open Shortest Path First Configuration

```
[Open Shortest Path First Menu]
      aindex  - OSPF Area (index) menu
      range   - OSPF Summary Range menu
      if      - OSPF Interface menu
      virt    - OSPF Virtual Links menu
      host    - OSPF Host Entry menu
      redist  - OSPF Route Redistribute menu
      lsdb    - Set the LSDB limit
      traps   - Enable/Disable the OSPF traps
      default - Originate default route information
      on      - Globally turn OSPF ON
      off     - Globally turn OSPF OFF
      current - Display current OSPF configuration
```

OSPF uses flooding to exchange link state updates between routers. Any change in routing information is flooded to all routers in the network.

Table 6-34 OSPF IP Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
aindex *<area index (0-7)>* Displays the OSPF Area menu. This menu is used for defining OSPF area numbers and parameters. For menu options, see page 191. **Note—**The area index specified in this menu option does not represent the actual OSPF area number. It is an arbitrary index used only on the switch. The actual value is defined with `/cfg/ip/opsf/aindex/areaid` (see page 191). For more information, refer to the *Web OS 7.2 Application Guide*.
range *<range number (1-64)>* Displays the OSPF Summary Range menu. This menu is used for defining OSPF summary routes and condensing OSPF routing information. For menu options, see page 192.
if *<IP interface number (1-1024)>* Displays the OSPF Interface menu. This menu is used for attaching IP interface networks to OSPF areas. For menu options, see page 194.
virt *<virtual link number (1-8)>* Displays the OSPF Virtual Links menu. This menu is used for connecting partitioned areas together. For menu options, see page 196. For more information, see the *Web OS 7.2 Application Guide*.

Table 6-34 OSPF IP Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
host *<host number (1-1024)>* Displays the Host Entry menu. This menu is used for configuring host routes for server load balancing or ABR load sharing and failover. For menu options, see page 198.
redist fixed|static|rip|ebgp|ibgp Displays the OSPF Route Redistribute menu. For menu options, see page 199.
lsdb *<interface number (0-4000)>* This command sets the maximum size limit for the Link State Database (LSDB). This number depends on the size of the network. In most cases, the default size of 2000 entries is sufficient. The value 0 indicates no limit.
traps *<trap number (1-16)*`all`> <**disable|enable** (**d|e**)> Enable or disable OSPF traps. You can specify a single trap number or specify `all` to affect all traps.
default *(<metric value (1-16777215)> <metric type* (**1|2**)*>)* | **none** The metric value sets the priority for choosing this switch for default routes where multiple Area Boundary Routers (ABR) or Autonomous System Boundary Routers (ASBR) exist in an area. Metric type determines the method for influencing routing decisions for external routes. Type 1 external routes use the OSPF internal link state metrics. Type 2 external routes use the cost of the interface. The default is type 1. Selecting `none` sets *no default routes*. For more information, see the *Web OS 7.2 Application Guide*.
on Globally turns on OSPF.
off Globally turns off OSPF.
current Displays the current OSPF configuration.

/cfg/ip/ospf/aindex <area index (0-7)>

OSPF Area Configuration

NOTE – The *area index* (0-7) represents the OSPF area you wish to configure.

```
[OSPF Area (index) 1 Menu]
      areaid  - Set area ID
      type    - Set area type
      metric  - Set stub area metric
      auth    - Set authentication type
      spf     - Set time interval between two SPF calculations
      enable  - Enable area
      disable - Disable area
      delete  - Delete area
      current - Display current OSPF area configuration
```

Table 6-35 OSPF Area Menu Options (/cfg/ip/ospf/aindex)

Command Syntax and Usage

areaid *<area ID, such as 0.0.0.0>*

Sets the OSPF area number in dotted decimal notation. The area number can be set using the last octet format (0.0.0.1 for area 1) or using multi-octet format (1.1.1.1), though the same format should be used throughout an area. For more information, see the *Web OS 7.2 Application Guide*.

type transit|stub|nssa

Sets the area type:

- `transit` for the backbone or any area that contains a virtual link.
- `stub` for any area that contains no external routes.
- `nssa` for any area that can process external routes but does not advertise external routes originating from outside its area.

The default type is `transit`. For more information, see the *Web OS 7.2 Application Guide*.

metric *<1-65535>*

Sets the stub area metric. Other routing devices add this value to the cost of routing to this stub area when building their SPF tree.

auth none|password

Sets the authentication type for this area:

- `none` turns OSPF authentication off.
- `password` turns type 1 (plain text) password authentication on. The passwords are set using the `key` options for OSPF interfaces (see page 195) and OSPF virtual links (see page 197).

Table 6-35 OSPF Area Menu Options (/cfg/ip/ospf/aindex)

Command Syntax and Usage
spf *<time interval (0-255)>* Sets the time interval, in seconds, between each calculation of the shortest path tree. The default is 10.
enable Enables this area.
disable Disables this area.
delete Deletes this area.
current Displays the current OSPF area configuration.

/cfg/ip/ospf/range *<range number>*

OSPF Summary Range Configuration

NOTE – The *range number* (1-64) represents the OSPF summary route you wish to configure.

```
[OSPF Summary Range 1 Menu]
      addr    - Set IP address
      mask    - Set IP mask
      aindex  - Set area index
      hide    - Enables or disables hide range
      enable  - Enable range
      disable - Disable range
      delete  - Delete range
      current - Display current OSPF summary range configuration
```

This menu is used for defining OSPF summary routes. Without summarization, each routing device in an OSPF network would retain a route to every subnet in the network. With summarization, routing devices can reduce some sets of routes to a single advertisement, reducing both the load on the routing device and the perceived complexity of the network. The importance of route summarization increases with network size.

Table 6-36 OSPF Summary Range Menu Options (/cfg/ip/ospf/range)

Command Syntax and Usage
addr *<IP address, such as 192.4.17.1>* Sets the base IP address for the summary range, using dotted decimal notation.
mask *<IP mask, such as 255.255.255.0>* Sets the IP mask for the summary range, using dotted decimal notation.
aindex *<area index (0-7)>* Sets the area index number into which the summary range is to be injected.
hide enable|disable When enabled, this forces the address range to be removed from any other summary ranges being injected into the defined area by this switch. This is useful for removing sections from large summary ranges that are not fully contiguous or contain gaps. This option is disabled by default.
enable Enables this range.
disable Disables this range.
delete Deletes this range.
current Displays the current OSPF summary range configuration.

/cfg/ip/ospf/if <interface number>

OSPF Interface Configuration

NOTE – The *interface number* (1-1024) represents the OSPF interface you wish to configure.

```
[OSPF Interface 1 Menu]
      aindex  - Set area index
      prio    - Set interface router priority
      cost    - Set interface cost
      hello   - Set hello interval in seconds
      dead    - Set dead interval in seconds
      trans   - Set transit delay in seconds
      retra   - Set retransmit interval in seconds
      key     - Set authentication key
      enable  - Enable interface
      disable - Disable interface
      delete  - Delete interface
      current - Display current OSPF interface configuration
```

NOTE – The hello interval (`hello`), dead interval (`dead`), transit interval (`trans`) and retransmit interval (`retra`) must be the same on *all* OSPF routing devices within an area. Using incompatible values could keep adjacencies from forming and could stop or loop routing updates.

Table 6-37 OSPF Interface Menu Options (/cfg/ip/ospf/if)

Command Syntax and Usage
aindex *<area index (0-7)>* Sets the OSPF area index to attach to the network for the current IP interface.
prio *<0-127>* Sets the IP interface (IF) priority that is used when electing a Designated Router (DR) and Backup Designated Router (BDR) for the area. The default is 1 (lowest priority). A value of 0 specifies that the elected interface is `DROTHER` and cannot be used as a DR or BDR. For more information, see the *Web OS 7.2 Application Guide*.
cost *<output cost (1-65535)>* Sets the cost of output routes on this interface. Cost is used in calculating the shortest path tree throughout the AS. Cost is based on bandwidth. Low cost indicates high bandwidth. The default is 1.

Table 6-37 OSPF Interface Menu Options (/cfg/ip/ospf/if)

Command Syntax and Usage
hello *<hello interval(1-65535)>* Sets the hello interval in seconds. The switch sends hello messages to inform neighbors that the link is up. The default is 10 seconds. This value must be the same on all routing devices within the area.
dead *<dead interval (1-65535)>* Sets the router dead interval, in seconds. If the switch does not receive `hello` on the IP interface within the dead interval, the switch will declare the interface to be down. Typically, the dead value is four times the value of `hello`. The default is 40 seconds. This value must be the same on all routing devices within the area.
trans *<transit delay (1-65535)>* Sets the transit delay, in seconds. This is the estimated time required to transmit an LSA to adjacencies on this interface, taking into account transmission and propagation delays. The default is 1 second. This value must be the same on all routing devices within the area.
retra *<time interval (1-65535)>* Sets the time interval, in seconds, between each transmission of LSAs to adjacencies on this interface. The default value is five seconds. This value must be the same on all routing devices within the area.
key *<type 1 password>* When `auth` is set to `password`, the key option sets the password to be used for OSPF authentication. Specify a type 1 (plain text) password of up to eight characters. This option is used with the `auth` option (see page 191). When `auth` is set to `none`, the key option is ignored. To clear the key, specify `none` as the value.
enable Enables this interface.
disable Disables this interface.
delete Deletes this interface.
current Displays the current OSPF interface configuration.

/cfg/ip/ospf/virt <link number>

OSPF Virtual Link Configuration

NOTE – The *link number* (1-8) represents the OSPF virtual link you wish to configure.

```
[OSPF Virtual Links Menu]
      aindex  - Set area index
      hello   - Set hello interval in seconds
      dead    - Set dead interval in seconds
      trans   - Set transit delay in seconds
      retra   - Set retransmit interval in seconds
      nbr     - Set router ID of virtual neighbor
      key     - Set authentication key
      enable  - Enable interface
      disable - Disable interface
      delete  - Delete interface
      current - Display current OSPF interface configuration
```

Virtual links are typically created to connect one area to the backbone through another non-backbone area. The virtual link must be configured at each endpoint of the virtual link, though they may traverse multiple routing devices.

The minimum requirements for configuring a virtual link are the `aindex` and `nbr` options in this menu and the `rtrid` option in the IP Configuration menu (see page 175). For more information, see the *Web OS 7.2 Application Guide*.

NOTE – The hello interval (`hello`), dead interval (`dead`), transit interval (`trans`) and retransmit interval (`retra`) must be the same on *all* OSPF routing devices within an area. Using incompatible values could keep adjacencies from forming and may stop or loop routing updates.

Table 6-38 OSPF Virtual Link Menu Options (/cfg/ip/ospf/virt)

Command Syntax and Usage
aindex *<area number (0-7)>* Sets the OSPF area index through which the virtual link passes.
hello *<value (1-65535)>* Sets the hello interval, in seconds. The switch sends hello messages to inform other network devices that the virtual link is up. The default is 10 seconds. This value must be the same on all routing devices within the area.

Table 6-38 OSPF Virtual Link Menu Options (/cfg/ip/ospf/virt)

Command Syntax and Usage
dead *<dead interval (1-65535)>* Sets the dead interval, in seconds. If the switch does not receive a hello on the IP interface within the deal interval, the switch will declare the virtual link to be down. Typically, the dead value is four times the hello value. The default is 40 seconds. This value must be the same on all routing devices within the area.
trans *<transit delay (1-65535)>* Sets the transit delay, in seconds. This is the estimated time required to transmit an LSA to adjacencies, taking into account transmission and propagation delays. The default is one second. This value must be the same on all routing devices within the area.
retra *<time interval (1-65535)>* Sets the time interval, in seconds, between each transmission of LSAs to adjacencies. The default is five seconds. This value must be the same on all routing devices within the area.
nbr *<router ID, such as 192.4.17.101>* Set router ID of the recipient neighbor (endpoint of the virtual link). The neighbor router ID is specified in dotted decimal format.
key *<type 1 password>* This option is used in conjunction with the `auth` option (see page 191). When `auth` is set to *password*, the key option sets the password to be used for OSPF authentication. Specify a type 1 (plain text) password of up to eight characters. When `auth` is set to `none`, the key option is ignored. To clear the key, specify `none` as the value.
enable Enables this virtual link.
disable Disables this virtual link.
delete Deletes this virtual link.
current Displays the current OSPF virtual link configuration.

/cfg/ip/ospf/host <range number>

OSPF Host Entry Configuration

NOTE – The *range number* (1-1024) represents the OSPF host entry you wish to configure.

```
[OSPF Host Entry 1 Menu]
      addr    - Set host entry IP address
      aindex  - Set area index
      cost    - Set cost of this host entry
      enable  - Enable host entry
      disable - Disable host entry
      delete  - Delete host entry
      current - Display current OSPF host entry configuration
```

This menu is used for configuring host routes. Host routes are used for advertising network device IP addresses to external networks, allowing Server Load Balancing (SLB) with OSPF. It also allows ABR load sharing and failover with OSPF. For more information, see the *Web OS 7.2 Application Guide*.

Table 6-39 OSPF Host Entry Menu Options (/cfg/ip/ospf/host)

Command Syntax and Usage
addr *<IP address, such as 192.4.17.1>* Specify the host IP address and mask that you want to advertise to the network.
aindex *<0-7>* Sets the area index into which the host entry is to be announced.
cost *<1-65535>* Sets the cost of output routes on this host. Cost is used in calculating the shortest path tree throughout the AS. Cost is usually based on bandwidth (low cost indicates high bandwidth) but can be used to set route preferences. The default is 1.
enable Enables this host entry.
disable Disables this host entry.
delete Deletes this host entry.
current Displays the current OSPF host entry configuration.

/cfg/ip/ospf/redist <route>

OSPF Host Entry Configuration

NOTE – The *route* represents the route you wish to configure—fixed|static|rip|ebgp|ibgp

```
[OSPF Redistribute fixed|static|Rip|eBGP|iBGP  Menu]
      addo    - Add rmap into out-rmap list
      remo    - Remove rmap from out-rmap list
      export  - Export all routes of this protocol
      cur     - Display current route-maps added
```

With this menu, you can specify which routes of the protocol to export into OSPF.

Table 6-40 OSPF Redistribution Menu Options (/cfg/ip/ospf/redist)

Command Syntax and Usage
addo *(<route map (1-32)> <route map (1-32)> ...)*|**all** Adds selected routing maps to the rmap list.To add all 32 route maps, enter `all`. To add specific route maps, enter routing map numbers one per line, NULL at end.
remo *(<route map (1-32)> <route map (1-32)> ...)*|**all** Removes routing maps from the `rmap` list. To remove all 32 route maps, enter all. To remove specific route maps, enter routing map numbers one per line, NULL at end.
export *<metric [1-16777215]>*|**none***><metric type [***1**|**2***]>* Exports the routes of this protocol as external OSPF AS-external LSAs in which the metric and metric type are specified. To remove a previous configuration and stop exporting the routes of the protocol, enter `none`.
cur Displays the current route map settings.

/cfg/ip/bgp

Border Gateway Protocol Configuration

```
[Border Gateway Protocol Menu]
      peer    - Peer menu
      aggr    - Aggregation menu
      pref    - Set Local Preference
      on      - Globally turn BGP ON
      off     - Globally turn BGP OFF
      current - Display current BGP configuration
```

This menu allows you to configure routers to share internal routing information, advertise external information, and select the best route.

Table 6-41 BGP Option Menu Options (/cfg/ip/bgp)

Command Syntax and Usage
peer *<peer number (1-16)>* Displays the Peer menu. To view the menu options, see page 201.
aggr *<aggregate number (1-16)>* Displays the Aggregation menu. To view the menu options, see page 204.
pref *<local preference (0-4294967294)>* Sets the local preference. The path with the higher value is preferred.
on Globally turns on BGP.
off Globally turns off BGP.
current Displays the current BGP configuration.

/cfg/ip/bgp/peer <peer number>

BGP Peer Configuration

NOTE – The *peer number* (1-16) represents the BGP peer you wish to configure.

```
[BGP Peer 1 Menu]
      redist  - Redistribution menu
      addr    - Set remote IP address
      ras     - Set remote autonomous system number
      hold    - Set hold time
      alive   - Set keep alive time
      advert  - Set min time between advertisement
      ttl     - Set time-to-live of IP datagrams
      addi    - Add rmap into in-rmap list
      addo    - Add rmap into out-rmap list
      remi    - Remove rmap from in-rmap list
      remo    - Remove rmap from out-rmap list
      enable  - Enable peer
      disable - Disable peer
      delete  - Delete peer
      current - Display current peer configuration
```

The value of *peer number* can be from 1 to 16. This menu allows you to configure upstream BGP providers (peers).

Table 6-42 BGP Peer Menu Options (/cfg/ip/bgp/peer)

Command Syntax and Usage
redist Displays the Redistribution Menu. To view the menu options, see page 203.
addr *<IP address, such as 192.4.17.101>* Sets the remote IP address.
ras *<RAS number (1-65535)>* Sets the Remote Autonomous System (RAS) number.
hold *<hold time (0|3-65535)>* Sets hold time. The unit of time is seconds.
alive *<keep alive time (0 |1-21845)>* Sets keep alive time. The unit of time is seconds.

Table 6-42 BGP Peer Menu Options (/cfg/ip/bgp/peer)

Command Syntax and Usage
advert *<advert time (1-65535)>* Sets the minimum time between advertisements. The unit of time is seconds.
ttl *<time to live (1-255)>* Sets time to live—the maximum number of hops—for the IP datagrams.
addi *<route map (1-32)>* Adds the selected route map to an *in-rmap* list.
addo *<route map (1-32)>* Adds the selected route map to the *out-rmap* list.
remi *<route map (1-32)>* Removes the selected route map from the *in-rmap* list.
remo *<route map (1-32)>* Removes the selected route map from the *out-rmap* list.
enable Enables the peer.
disable Disables the peer.
delete Deletes the peer.
current Displays the current peer configuration.

/cfg/ip/bgp/peer/redist

BGP Redistribution Configuration

```
[Redistribution Menu]
      metric  - Set default-metric of advertised routes
      default - Set default route action
      rip     - Enable/disable advertising RIP routes
      ospf    - Enable/disable advertising OSPF routes
      fixed   - Enable/disable advertising fixed routes
      static  - Enable/disable advertising static routes
      vip     - Enable/disable advertising VIP routes
      current - Display current redistribution configuration
```

With this menu, you can specify which routes of the protocol to export into Border Gateway Protocol.

Table 6-43 BGP Redistribution Menu Options (/cfg/ip/bgp/peer/redist)

Command Syntax and Usage
metric <*(1-4294967294)*>|**none**> Sets the default metric of the advertised route.
default none|import|originate|redistribute Sets the default route action.
rip disable|enable (**d**|**e**) Enables or disables advertising RIP routes.
ospf disable|enable (**d**|**e**) Enables or disables advertising OSPF routes.
fixed disable|enable (**d**|**e**) Enables or disables advertising fixed routes.
static disable|enable (**d**|**e**) Enables or disables advertising static routes.
vip disable|enable (**d**|**e**) Enables or disables advertising VIP routes.
current Displays the current redistribution configuration.

/cfg/ip/bgp/aggr <aggregate number>

Aggregate Routing Configuration

NOTE – The *aggregate number* (1-16) represents the aggregation route you wish to configure.

```
[BGP Aggr 1 Menu]
      addr    - Set aggregation IP address
      mask    - Set aggregation network mask
      enable  - Enable aggregation
      disable - Disable aggregation
      delete  - Delete aggregation
      current - Display current aggregation configuration
```

This menu allows you to configure aggregate routing to condense the number of routes between internal and external peer routers.

Table 6-44 Aggregate Menu Options (/cfg/ip/bgp/aggr)

Command Syntax and Usage
addr *<IP address, such as 192.4.17.101>* Adds the IP address to the selected aggregate.
mask *<IP subnet mask, such as 255.255.255.0>* Sets the IP mask for the selected aggregate.
enable Enables the selected aggregate.
disable Disables the selected aggregate.
delete Deletes the selected aggregate.
current Displays the current aggregate configuration.

/cfg/ip/port <slot-number port-number>

IP Forwarding Port Configuration

NOTE – The *slot-port number* represents the IP port you wish to configure.

```
[IP Forwarding Port A1 Menu]
      on      - Turn Forwarding ON
      off     - Turn Forwarding OFF
      current - Display current port configuration
```

The IP Port Menu allows you to turn IP forwarding on or off on a port-by-port basis.

Table 6-45 IP Forwarding Port Menu Options (/cfg/ip/port)

Command Syntax and Usage
on Enables IP forwarding for this port.
off Disables IP forwarding for this port.
current Displays the current IP forwarding settings for this port.

/cfg/ip/dns

Domain Name System Configuration

```
[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
      current - Display current DNS configuration
```

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.

Table 6-46 Domain Name Service Menu Options (/cfg/ip/dns)

Command Syntax and Usage
prima *<IP address, such as 192.4.17.101>* Sets the IP address for your primary DNS server. Use dotted decimal notation.
secon *<IP address, such as 192.4.17.101>* Sets the IP address for your secondary DNS server. If the primary DNS server fails, the configured secondary will be used instead. Use dotted decimal notation.
dname *<dotted DNS notation>*|**none** Sets the default domain name used by the switch. For example: `mycompany.com`
current Displays the current DNS settings.

/cfg/ip/bootp

Bootstrap Protocol Relay Configuration

```
Bootstrap Protocol Relay Menu]
      addr    - Set IP address of BOOTP server
      addr2   - Set IP address of second BOOTP server
      on      - Globally turn BOOTP relay ON
      off     - Globally turn BOOTP relay OFF
      current - Display current BOOTP relay configuration
```

BOOTP configuration can set the switch to forward Dynamic Host Configuration Protocol (DHCP) messages. BOOTP relay is useful when BOOTP clients are on different VLANs, and the switch is the primary router between those VLANs.

Table 6-47 BOOTP Menu Options (/cfg/ip/bootp)

Command Syntax and Usage
addr *<IP address, such as 192.4.17.101>* Sets the IP address of the primary BOOTP server.
addr2 *<IP address, such as 192.4.17.1>* Sets the IP address of the secondary BOOTP server.
on Globally sets BOOTP relay on.
off Globally sets BOOTP relay off.
current Displays the current BOOTP relay configuration.

/cfg/ip/metric <metric name>

Default Gateway Metrics Configuration

If multiple default gateways are configured and enabled, a metric can be set to determine which primary gateway is selected. There are two metrics; each is described in the table below.

Table 6-48 Default Gateway Metrics (/cfg/ip/metric).

Option	Description
`s (strict)`	The gateway number determines its level of preference. Gateway #1 acts as the preferred default IP gateway until it fails or is disabled. At that point, the next in line will take over as the default IP gateway.
`r (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.

/cfg/slb
SLB Configuration

```
[Layer 4 Menu]
      real    - Real server menu
      group   - Real server group menu
      virt    - Virtual server menu
      filt    - Filtering menu
      port    - Layer 4 port menu
      gslb    - Global SLB menu
      sync    - Config synch menu
      adv     - Layer 4 advanced menu
      on      - Globally turn Layer 4 processing ON
      off     - Globally turn Layer 4 processing OFF
      current - Display current Layer 4 configuration
```

The switch can be configured to balance user session traffic among the available servers in a server group and to set up filters.

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.

Table 6-49 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage
`real` *<real server number (1-1024)>* Displays the menu for configuring real servers. To view menu options, see page 211.
`group` *<real server group number (1-256)>* Displays the menu for placing real servers into real server groups. To view menu options, see page 214.
`virt` *<virtual server number (1-1024)>* Displays the menu for defining virtual servers. To view menu options, see page 218.
`filt` *<filter ID (1-1024)>* Displays the menu for Filtering and Application Redirection. To view menu options, see page 222.

Table 6-49 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage
port *<port as slot and port, such as b12>* Displays the menu for setting physical switch port states for Layer 4 activity. To view menu options, see page 228.
gslb Displays the menu for configuring Global Server Load Balancing. To view menu options, see page 230.
sync Displays the configuration synchronization menu. To view menu options, see page 236.
adv Displays the Layer 4 Advanced Menu. To view menu options, see page 238.
on Globally turns 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 260). Enabling Layer 4 services is not necessary for using filters only to `allow`, `deny`, or NAT traffic (see "Defining IP Address Ranges for Filters" on page 295).
off Globally disables 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.
current Displays the current Server Load Balancing configuration.

/cfg/slb/real <server number>

Real Server SLB Configuration

NOTE – The *real server number* (1-1024) represents the real server that you wish to configure.

```
[Real server 1 Menu]
      rip     - Set IP addr of real server
      name    - Set server name
      weight  - Set server weight
      maxcon  - Set maximum number of connections
      backup  - Set backup real server
      inter   - Set interval between health checks
      retry   - Set number of failed attempts to declare server DOWN
      restr   - Set number of successful attempts to declare server
                UP
      remote  - Enables or disables remote site operation
      proxy   - Enables or disables client proxy operation
      enable  - Enable real server
      disable - Disable real server
      delete  - Delete real server
      current - Display current real server configuration
```

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

- Real server IP address
- Enabling the real server

Table 6-50 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
rip *<server IP address>* Sets the IP address of the real server in dotted decimal format. When this command is used, the address entered is PINGed to determine if the server is up, and the administrator will be warned if the server does not respond.
name *<string, maximum 15 characters>* Defines a 15-character alias for each Real Server. This will enable the network administrator to quickly identify the server by a natural language keyword value.

Table 6-50 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
weight *<server weight (1-48)>* Sets the weight value (1-48) 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. Setting the weight to 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 216).
maxcon *<maximum connections (0-65535)>* Sets 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 set maximum connection (`maxcon`) limit is reached. New connections will be issued again to this server once the number of current connections has decreased below the `maxcon` setting. If all the servers of a real server group representing a virtual server reach their `maxcon` limit at the same time, client requests will be dropped by the virtual server.
backup *<real server number (1-1024)|***none***>* Sets the real server used as the backup/overflow server for this real server. To prevent loss of service when 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 `maxcon` 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.
inter *<number of seconds between health checks (0-60)>* Sets 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 0 to 60 seconds. The default interval is two seconds. Setting the interval to 0 disables health checking. The real server is assumed to be always available.
retry *<number of consecutive health checks (1-63)>* Sets 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.

Table 6-50 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
restr *<number of consecutive health checks (1-63)>* Sets 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.
remote disable|enable (**d|e**) Enables or disables remote site operation for this server. This should be enabled when the real IP address supplied above represents a remote server (real or virtual) that the switch will access as part of its Global Server Load Balancing network.
proxy disable|enable (**d|e**) Enables or disables proxy IP address translation. With this option enabled (default), a client request from any application can be proxied using a load-balancing Proxy IP address (PIP).
enable 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/enable` in Table 7-5 on page 258 for an operations-level command.
disable Disables 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` and `save` commands are used, disables this real server until it is explicitly re-enabled. This option *does not* perform a graceful server shutdown. See `/oper/slb/disable` on Table 7-5 on page 258 for an operations-level command.
delete Deletes 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.
current Displays the current configuration information for this real server.

/cfg/slb/group <group number>

Real Server Group SLB Configuration

NOTE – The *group number* (1-256) represents the number of the real server group you wish to configure.

```
[Real server group 1 Menu]
      add     - Add real server
      remove  - Remove real server
      metric  - Set metric used to select next server in group
      content - Set health check content
      health  - Set health check type
      backup  - Set backup real server or group
      name    - Set real server group name
      realthr - Set real server failure threshold
      delete  - Delete real server group
      current - Display current group configuration
```

This menu is used for combining real servers into real server groups. Each real server can belong to more than one group. Real server groups are used for both Server Load Balancing and Application Redirection.

Each real server group should consist of all the real servers that provide a specific service for load balancing. Each group must consist of at least one real server.

Table 6-51 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
add *<real server number (1-1024)>* Adds a real server (1-1024) to this real server group.
rem *<real server number (1-1024)>* Removes a real server from this real server group. You will be prompted for the ID number of the real server to remove from this group.
metric leastconns|roundrobin|minmisses|hash Sets the load balancing metric used for determining which real server in the group will be the target of the next client request. See "Server Load Balancing Metrics" on page 216.
content *<filename>|<host>|<filename>*|**none** Defines the specific content that is examined during health checks. The content depends on the type of health check specified in the **health** option (see page 215).

Table 6-51 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
health **icmp**|**tcp**|**http**|**dns**|**pop3**|**smtp**|**nntp**|**ftp**|**imap**|**radius**|**script<n>** Sets 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, uses HTTP 1.1 `GETS` when a `HOST:` header is required to check that the URL content specified in `content` is accessible on the server. Otherwise, an `HTTP/1.0 GET` occurs. ■ `dns` For Domain Name Service, check that the domain name specified in `content` can be resolved by the server. ■ `pop3` For user mail service, check that the *user:password* account specified in `content` exists on the server. ■ `smtp` For mail-server to mail-server services, check that the user specified in `content` is accessible on the server. ■ `nntp` For newsgroup services, check that the newsgroup name specified in `content` is accessible on the server. ■ `ftp` For FTP services, check that the filename specified in `content` is accessible on the server through anonymous login. ■ `imap` For user mail service, check that the *user:password* value specified in `content` exists on the server. ■ `radius` For Remote Access Dial-up Service (RADIUS) authentication, check that the *user:password* value specified in `content` exists on the switch and the server. To perform application health checking to a RADIUS server, the network administrator must also configure the `/cfg/slb/adv/secret` parameter. The `secret` value is a field of 16 alphanumeric characters that is used by the switch to encrypt a password during the RSA Message Digest Algorithm (MD5) and by the RADIUS server to decrypt the password during verification. ■ `script` Enables the use of script-based health checks in send/expect format to check for application and content availability.
backup < **r**<*real server number*>|**g**<*group number*>**|none** Sets 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 should fail, 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 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 `maxcon` (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.
name <*string, maximum 15 characters*> Defines a 15-character alias for each Real Server Group. This will enable the network administrator to quickly identify the server group by a natural language keyword value.

Table 6-51 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
realthr *<real server failure threshold (1255)|0 for disabled>* Sets real server failure threshold. Once the configured number of servers have failed, a syslog message is generated.
delete Deletes this real server group from the Layer 4 software configuration. This removes the group from operation under all virtual servers to which it is assigned. *Use this command with caution*; if you remove the only group assigned to a virtual server, the virtual server will become inoperative.
current Displays the current configuration parameters for this real server group.

Server Load Balancing Metrics

Using the *metric* command, you can set a number of metrics for selecting which real server in a group gets the next client request. These metrics are described in the following table:

Table 6-52 Real Server Group Metrics

Option	Description
minmisses	This minimum misses (minmisses) metric is optimized for Application Redirection. 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. The metric 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. It is recommended that minmisses be used only when persistence is a must. The metric minmisses cannot be used for Firewall Load Balancing (FWLB) because the real server IP addresses used in calculation are different on each side of the firewall.

Table 6-52 Real Server Group Metrics

Option	Description
`hash`	The `hash` metric, like `minmisses`, 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 `min-misses` 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.
`leastconns`	With the least connections (`leastconns`) 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.
`roundrobin`	With the round robin (`roundrobin`) 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 `weight` option on page 211), a higher proportion of connections is 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`.

/cfg/slb/virt <server number>

Virtual Server SLB Configuration

NOTE – The *virtual server number* (1- 1024) represents the number of the virtual server you wish to configure.

```
[Virtual server 1 Menu]
      vip     - Set IP addr of virtual server
      dname   - Set domain name of virtual server
      layer3  - Enables or disables layer 3 only balancing
      service - Virtual service menu
      enable  - Enable virtual server
      disable - Disable virtual server
      delete  - Delete virtual server
      current - Display current virtual configuration
```

This menu is used for configuring the virtual servers for client requests for server load balancing. The required parameters to configure are as follows.

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

Table 6-53 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
vip *<server IP address>* Sets the IP address of the virtual server using dotted decimal notation. The virtual server created within the switch will respond to ARPs and PINGs 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 *<domain name>*|**none** Sets 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, and so on). An example would be "foocorp.com." It does not include the hostname portion (www, www2, ftp, and so on). To define the hostname, see hname in Table 6-54 on page 220. To clear the dname, specify the name as none.

Table 6-53 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
layer3 disable|enable (**d|e**) Normally, the client IP address is used with the client Layer 4 port number to produce a session identifier. When the `layer3` 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 exists between them. This is necessary for some server applications where state information about the client system is divided across different simultaneous connections and in applications where TCP fragments are generated. If the real server to which the client is assigned becomes unavailable, the Layer 4 software will allow the client to connect to a different server.
service *< name or virtual port (2 - 65535)>* Displays the Virtual Service Menu. The virtual port name can be a well-known port name, such as http, ftp, and so on. To view a list of well-known ports, see page 224. To view menu options, see page 220.
enable Enables this virtual server. This option activates the virtual server within the switch so that it can service client requests sent to its defined IP address.
disable This option disables the virtual server so that it no longer services client requests.
delete 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.
current Displays the current configuration for the specified virtual server.

/cfg/slb/virt <server number>/service <virtual port or name>

Virtual Server Service Configuration

```
[Virtual Server 1 2 Service Menu]
      group   - Set real server group number
      rport   - Set real port
      hname   - Set hostname
      udp     - Enables or disables UDP balancing
      frag    - Enables or disables remapping UDP server fragments
      nonat   - Enables or disables only substituting MAC addresses
      delete  - Delete virtual service
      cur     - Display current virtual service configuration
```

This menu is used for configuring services assigned to a virtual server. For information about access to real filters, see "Direct Access to Real Servers" on page 287.

Table 6-54 Virtual Server Service Configuration Menu Options (/cfg/slb/virt/service)

Command Syntax and Usage
group *<real server group number (1-256)>* Sets a real server group for this service. You will be prompted to enter the number (1-256) of the real server group to add to this service.
rport *<real port (2-65535)>* Defines the real server TCP or UDP port assigned to this service. By default, this port is the same as the virtual port (service virtual port). If `rport` is configured to be different than the virtual port defined in **/cfg/slb/virt/service** *<virtual port>*, the switch will map the virtual port to this real port.
hname *<hostname>*|**none** Sets the hostname for a service added. This is used in conjunction with the command `dname` (page 218) to create a full host/domain name for individual services. The format for this command is as follows: # `hname` *<hostname>* For example, to add a hostname for Web services, you could specify "www" as the hostname. If a `dname` of "foocorp.com" was defined (page 218), "www.foocorp.com" would be the full host/domain name for the service. To clear the hostname for a service, use the command `hname none`.

Table 6-54 Virtual Server Service Configuration Menu Options (/cfg/slb/virt/service)

Command Syntax and Usage
udp disable|enable|stateless (**d**|**e**|**s**) Enables, disables, or makes stateless UDP balancing for a virtual port. You can configure this option if the service(s) to be load balanced include UDP and TCP. (For example, DNS uses UDP and TCP.) In those environments, you must activate UDP balancing for the particular virtual servers that clients will communicate with using UDP.
frag disable|enable (**d**|**e**) Enables or disables substitution of IP addresses in server response fragments with virtual addresses when a virtual server is load-balancing a real server.
nonat disable|enable (**d**|**e**) Enables or disables substituting the MAC addresses of the real server. This option does not substitute IP addresses.
delete This command removes this virtual service 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 service.
current Displays the current configuration of services on the specified virtual server.

/cfg/slb/filt <filter number>

SLB Filter Configuration

NOTE – The *filter number* (1-1024) represents the number of the virtual server you wish to configure.

```
[Filter 1 Menu]
      smac    - Set source MAC address
      dmac    - Set destination MAC address
      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
      ack     - Enables or disables TCP ack matching
      invert  - Enables or disables filter inversion
      action  - Set action
      group   - Set real server group for redirection
      rport   - Set real server port for redirection
      nat     - Set which addresses are network address translated
      name    - Set filter name
      adv     - Filter advanced menu
      enable  - Enable filter
      disable - Disable filter
      delete  - Delete filter
      current - Display current filter configuration
```

The switch supports up to 1024 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 parameters to configure are as follows:

- Set the address, masks, and/or protocol that 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.

For information about defining IP address for filters, see "Defining IP Address Ranges for Filters" on page 295.

Table 6-55 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
smac *<MAC address|***any***>* Sets the source MAC address.
dmac *<MACaddress|***any***>* Sets the destination MAC address
sip any|*<IP address, such as 192.4.17.0>* 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. Before configuring `sip`, `smac` must be set to `any`. For more information about defining an IP address range for filters, see Appendix A, "Defining IP Address Ranges for Filters
smask This IP address mask is used with the source IP address to select traffic that this filter will affect. Before configuring `smask`, `smac` must be set to `any`. See details below for more information on producing address ranges.
dip any|*<IP address, such as 192.4.17.0>* 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 command `dmask` below.
dmask *<IP subnet mask, such as 255.255.255.0>* This IP address mask is used with the command (`dip`)to select traffic that this filter will affect. See details below for more information on producing address ranges.
proto any|*<number>*|*<name>* If defined, traffic from the specified protocol is affected by this filter. The protocol number, name, or `any` can be specified. Following is a list of well-known protocols. ***Number*** ***Name*** 1 icmp 2 igmp 6 tcp 17 udp 89 ospf 112 vrrp

Table 6-55 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
sport <**any**|<*name*>|<*port*>|<*port*>-<*port*>> Sets the source TCP/UDP port or range. The IP protocol **proto** must be configured before configuring **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. Following is a list of well-known ports:

Number	*Name*	*Number*	*Name*
7	echo	179	bgp
9	discard	194	irc
11	sysdat	220	imap3
13	daytime	389	ldap
15	netstat	443	imap
19	chargen	520	rip
20	ftp-data	554	rtsp
21	ftp	1812	radius
22	ssh	1985	hsrp
23	telnet		
25	smtp		
37	time		
42	name		
43	whois		
53	domain		
67	bootps		
68	bootpc		
69	tftp		
70	gopher		
79	finger		
80	http		
109	pop2		
110	pop3		
111	sunrpc		
119	nntp		
123	ntp		
143	imap		
144	news		
161	snmp		
162	snmptrap		

Table 6-55 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
dport <**any**|<*name*>|<*port*>|<*port*>-<*port*>> Sets the destination TCP/UDP port or range. The IP protocol `proto` must be configured before configuring `sport`. 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.
ack disable|**enable** (**d**|**e**) Enables or disables `TCP ack` matching. Filters with this option enabled match only those frames that have the `TCP ACK` or `RST` flag set. This prevents servers from beginning a TCP connection (with a `TCP SYN`) from a source, such as TCP port 25. The server will drop any frames that have the `ACK` flag "spoofed" in them and will not allocate space for a new connection.
invert disable|**enable** (**d**|**e**) Inverts the filter logic. ■ If the conditions of the filter are met, no action is taken. ■ If the conditions for the filter are not met, the assigned action is performed.
action allow|**deny**|**redir**|**nat** Specifies the action this filter takes: **allow** Allows frames to pass. **deny** Discards frames that fit this filter's profile. This can be used for building basic security profiles. **redir** Redirects 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 209). **nat** Performs 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 (see page 226) and can also be combined with proxies.
group *<real server group number (1-256)>* This option applies only when redir is specified at the filter action. Define a real server group (1-256) to which redirected traffic will be sent.
rport *<real server port (0-65535)>* This option applies only when `redir` is specified at the filter action. The command `rport` defines the real server TCP or UDP port to which redirected traffic will be sent. For valid Layer 4 health checks, `rport` must be configured whenever TCP protocol traffic is redirected. Also, if transparent proxies are used for Network Address Translation (NAT) on the switch (see `pip` on page 229), `rport` must be configured for all Application Redirection filters.

Table 6-55 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
`nat source|dest` When NAT is set as the filter action (see "action allow|deny|redir|nat" on page 225), this command specifies whether the source or the destination information is remapped. If you specify source, the frame's source IP address (`sip`) and source port number (`sport`) are replaced with the destination ip address (`dip`) and destination port address (dport) values. If you specify `dest`, the frame's destination ip (`dip`) address and destination port number (`dport`) are replaced with the `sip` and `sport` values.
`name` *<string, maximum 15 characters>* Sets the filter name.
`adv` Displays the Filter Advanced Menu. To view menu options, see page 227.
`enable` Enables this filter.
`disable` Disables this filter.
`delete` Deletes this filter.
`current` Displays the current configuration for this filter.

/cfg/slb/filt <filter number>/adv

Advanced Filter Configuration

NOTE – The *filter number* (1-1024) represents the number of the filter you wish to configure.

```
[Filter 1 Advanced   Menu]
      proxy   - Enables or disables client proxy
      cache   - Enables or disables caching sessions that match
                filter
      log     - Enables or disables logging
      fwlb    - Enables or disables firewall redirect hash method
      cur     - Display current advanced filter configuration
```

Table 6-56 Advanced Filter Menu Options (/cfg/slb/filt/adv)

Command Syntax and Usage

proxy disable|enable (**d**|**e**)

Enables or disables client proxy. This option applies only when `redir` or `nat` is specified as 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 proxy IP address (see page 229) is not performed for traffic that meets the filter criteria. This is useful when some types of traffic must retain original IP address information or when other forms of translation (such as Application Redirection or NAT) are preferred.

cache proxy disable|enable (**d**|**e**)

Enables or disables caching sessions that match the filter.

log disable|enable (**d**|**e**)

Enables or disables logging filter messages.

fwlb disable|enable (**d**|**e**)

Enables or disables firewall redirect hash method.

cur

Displays the current advanced filter configuration.

/cfg/slb/port <slot-number port-number >

Port SLB Configuration

```
[SLB Port A1 Menu]
      client  - Enables or disables client processing for port
      server  - Enables or disables server processing for port
      hotstan - Enables or disables hot-standby processing
      intersw - Enables or disables inter-switch processing
      pip     - Set Proxy IP address for port
      submac  - Enables or disables source MAC address substitution
      filter  - Enables or disables filtering for port
      add     - Add filter to port
      remove  - Remove filter from port
      current - Display current port configuration
```

The Web OS switch software allows you to enable or disable processing independently for each type of Layer 4 traffic (client and server), expanding your topology options.

Table 6-57 Port Configuration Menu Options (/cfg/slb/port)

Command Syntax and Usage
client disable|enable (**d|e**) For server load balancing, the port can be enabled or disabled to process client Layer 4 traffic. Ports configured to process client request traffic bind servers to clients and provide address translation from the virtual server IP address to the real server IP address, remapping virtual server IP addresses and port values to real server IP addresses and ports. Traffic not associated with virtual servers is switched normally. Maximizing the number of these ports on the Layer 4 switch will improve the switch's potential for effective server load balancing.
server disable|enable (**d|e**) Ports configured to provide real server responses to client requests require real servers to be connected to the Layer 4 switch directly or through a hub, router, or another switch. When server processing is enabled, the switch port remaps real server IP addresses and Layer 4 port values to virtual server IP addresses and Layer 4 ports. Traffic not associated with virtual servers is switched normally.
hotstan disable|enable (**d|e**) Enables or disables hot standby processing. The standby switch blocks all ports configured as standby ports, whereas the master switch enables these same ports.
intersw disable|enable (**d|e**) Enables or disables interswitch processing. Use this option in conjunction with VRRP hot standby failover.

Table 6-57 Port Configuration Menu Options (/cfg/slb/port)

Command Syntax and Usage
pip *<proxy IP address (such as 192.4.17.0)>* 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, `pip` forces response traffic to return through the switch as required, rather than around it, as is possible in complex routing environments. Proxies are also useful for Application Redirection and Network Address Translation (NAT). When `pip` is used with Application Redirection filters, each filter's `rport` parameter must also be defined (see `rport` on page 220).
submac disable|enable (**d|e**) Enables or disables substituting the source MAC address of server load balancing (SLB) and Web Cache Redirection (WCR) client frames with one of the switch's IP interface MAC addresses.
server disable|enable (**d|e**) Ports configured to provide real server responses to client requests require real servers to be connected to the Layer 4 switch, directly or through a hub, router, or another switch. When server processing is enabled, the switch port remaps real server IP addresses and Layer 4 port values to virtual server IP addresses and Layer 4 ports. Traffic not associated with virtual servers is switched normally.
filter disable|enable (**d|e**) Enables or disables filtering on this port.
add *<filter ID (1-1024)>* Adds a filter for use on this port.
remove *<filter ID (1-1024)>* Removes a filter from use on this port.
current Displays the current configuration of this 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 in Table 7-5 on page 258).

/cfg/slb/gslb

Global SLB Configuration

```
[Global SLB Menu]
      site    - Remote Site menu
      lookup  - Network Preference Lookup Menu
      dns     - Enables or disables DNS handoffs
      ttl     - Set Time To Live of DNS resource records
      local   - Enables or disables DNS responses with only local
                addresses
      one     - Enables or disables DNS responses with only one address
      always  - Enables or disables DNS responses at least one address
      geo     - Enables or disables geographic awareness
      http    - Enables or disables HTTP redirects
      usern   - Enables or disables HTTP redirect to real server name
      mincon  - Set minimum number of site connections
      inter   - Set interval between remote site updates
      weight  - Set local weight
      on      - Globally turn Global SLB ON
      off     - Globally turn Global SLB OFF
      current - Display current Global SLB configuration
```

You can balance server traffic load across multiple physical sites. Up to 128 Global Server Load Balance (GSLB) sites with a total aggregate of 8192 service/site combinations can be configured.

Table 6-58 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage
site *<remote site (1-128)>* Displays the Remote Site Menu for up to 128 remote sites. To view menu options, see page 232.
lookup Displays the Network Preference Lookup menu. To view menu options, see page 234.
dns disable|enable (**d**|**e**) Enables or disables DNS hand-offs to peer sites by this switch. This should be enabled for proper GSLB operation. If disabled, whenever the switch receives a DNS request for a configured service, it will respond only with its own virtual server IP address, regardless of performance or load considerations.

Table 6-58 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage
ttl *<time to live in seconds (0-65535)>* Specifies the duration (from 0 to 65535 seconds) that the DNS response from the switch (indicating site of best service) will remain in the cache of DNS servers. A lower value may increase the ability of the GSLB 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 GSLB's response to sudden traffic changes.
local disable|enable (**d**|**e**) Enables or disables switch responses to DNS queries with local virtual server IP addresses. When enabled, the switch will always respond to DNS queries by providing a local virtual server IP address. This response will always occur as long as the virtual server IP address has healthy real servers with an aggregate of at least 1024 available connections (the total from each server's configured `maxcons` value, minus the server's current number of connections). When the real servers for the local virtual server IP addresses are unavailable or saturated, the switch will respond to DNS requests using normal GSLB rules.
one disable|enable (**d**|**e**) Enables or disables DNS responses with only one address. At most, one IP address is included in each DNS response.
always disable|enable (**d**|**e**) Enables or disables DNS responses (with) at least one address. At least one IP address is included in each DNS response. Even if all remote sites cannot handle another request, the local virtual server IP address is returned in a DNS response to eliminate long DNS time-outs caused by an empty response.
geo disable|enable (**d**|**e**) Enables or disables geographic awareness, such as the IANA table. If this option is disabled, all clients and sites will be assumed to exist in the same geographic region, allowing all sites to be eligible for each client.
http disable|enable (**d**|**e**) Enables or disables 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.
usern disable|enable (**d**|**e**) Enables or disables an HTTP redirect to a real server name. When a site redirects a client to another site using an HTTP redirect, the client is redirected to the new site's IP address. If `usern` is enabled, the client will be redirected to the domain name specified by the remote real server name and the virtual server domain name: `<remote real server name><virtual server domain name>`

Table 6-58 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage
mincon *<0-65535>* Sets the minimum number of available site connections. If the site's available sessions fall below this value, traffic won't be redirected to the site. A site is not eligible for more requests (such as DNS or HTTP redirects) once the number of available connections at a site drops below this threshold.
inter *<1-120>* Sets the time between Distributed Site State Protocol (DSSP) updates between this switch and its peers. The range is from 1 to 120 minutes.
weight *<1-48>* Sets the local weight. The range is from 1 to 48. The higher the weight value, the more connections that will be directed to the local site. The default is 1. The response time of this site is divided by *this weight* before the best site is assigned to a client. *Remote site* response times are divided by the *real server weight* before selection occurs.
on Activates Global Server Load Balancing (GSLB) for this switch. This option can be performed only once the optional GSLB software is activated (refer to "Activating Optional Software" on page 260).
off Turns GSLB off for this switch. Any active remote sites will still perform GSLB services with each other but will not hand off requests to this switch.
current Displays current Global SLB configuration.

/cfg/slb/gslb/site *<site number>*

GSLB Remote Site Configuration

```
[Remote site 1 Menu]
      name    - Set remote site name
      prima   - Set primary switch IP address of remote site
      secon   - Set secondary switch IP address of remote site
      update  - Enables or disables remote site updates
      enable  - Enable remote site
      disable - Disable remote site
      delete  - Delete remote site
      current - Display current remote site configuration
```

Up to 128 remote sites can be configured.

Table 6-59 GSLB Remote Site Menu Options (/cfg/slb/gslb/site)

Command Syntax and Usage
name *<string, maximum 15 characters>* Sets the name of the remote site. A maximum of 15 characters is allowed.
prima *<primary server IP address, such as 192.4.17.0>* Defines the IP interface IP address of the primary switch at the remote site used for Global Server Load Balancing. Use dotted decimal notation.
secon *<secondary server IP address, such as 192.4.17.>* 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.
update disable|enable (**d|e**) Enables or disables 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.
enable Enables this remote site for use with Global Server Load Balancing.
disable Disables this remote site. The switch will no longer use this remote site for Global Server Load Balancing.
delete Removes this remote site from operation and deletes its configuration.
current Displays the current remote site configuration.

NOTE – When `update` is enabled, Global Server Load Balancing uses service port 80 on the IP interface for DSSP updates. By default, the Web OS Web-based interface also uses port 80. Both services cannot use the same port. If both are enabled, configure the Web OS interface to use a different service port (see the `/cfg/sys` options in Table 6-2 on page 147).

/cfg/slb/gslb/lookup

Global Server Load Balance Lookup Configuration

```
[Global SLB Lookup Menu]
      network - Internet Network Preference Menu
      dname   - Set domain name for internal lookup table
      lookups - Enables or disables network preference lookups
      current - Display current lookup configuration
```

Table 6-60 GSLB Lookup Menu Options (/cfg/slb/gslb/lookup)

Command Syntax and Usage
network *<1-128>* Displays the Internet Network Preference Menu. To view menu options, see page 234.
dname *<domain name*|**none>** Sets the domain name for the internal lookup table.
lookups disable|**enable** (**d**|**e**) Enables or disables network preference lookups.
current Displays the current lookup configuration.

/cfg/slb/gslb/lookup/network *<number>*

GSLB Internet Lookup Network Preference Configuration

NOTE – The *number* (1-128) represents the network preference you wish to configure.

```
[Network 1 Menu]
      sip     - Set Source IP address
      mask    - Set Net mask
      vip1    - Set VIP address
      vip2    - Set VIP address
      delete  - Delete internet network entry
     current - Display current internet network entry configuration
```

The options in this menu allow you to define client networks to overwrite the IANA table. You should use regular GSLB to respond to a DNS request under the following conditions:

- Queried domain is not matched.
- Client IP address does not match address in the Network Preference Menu and no default entry is configured.
- There is an entry match in the Network Preference Menu. However, VIP1 and VIP2 are not healthy—they are down or over the minimum connections (`mincon`).

The *default entry* is one where the source IP address and mask are not configured (both are 0.0.0.0) and only the VIP1 and VIP2 are configured. All client networks not in the Network Preference Menu will use this entry to respond to a DNS request.

Table 6-61 GSLB Internet Network Preference Menu Options (/cfg/slb/gslb/lookup/network)

Command Syntax and Usage
`sip` *<IP address, such as 192.4.17.0>* Sets the source IP address. Specify an IP address in dotted decimal notation, or `any`. A range of IP addresses is produced when used with the `mask` option.
`mask` *<IP address, such as 255.255.0.0>* This IP address mask is used with the source IP address to find a correct virtual server IP address to respond to a DNS request.
`vip1` *<IP address, such as 192.4.17.100>* Sets the first virtual server IP address. The VIP address can be either a local or remote virtual server. The switch returns the VIP address with the least response time that is over the (`mincon`).
`vip2` *<IP address, such as 192.4.17.0>* Sets the second virtual server IP address.
`del` Deletes the specified network entry.
`cur` Displays the current GSLB lookup's Internet network entry configuration.

/cfg/slb/sync

Synchronizing Peer Switch Configurations

```
[Config Synchronization Menu]
   peer    - Synch peer switch menu
   filt    - Enables or disables syncing filter configuration
   ports   - Enables or disables syncing port configuration
   prios   - Enables or disables syncing VRRP priorities
   pips    - Enables or disables syncing proxy IP addresses
   current - Display current Layer 4 sync configuration
```

Synchronizing certain configurations—including proxy IP addresses Layer 4 port configuration, and filter configuration—can be enabled or disabled.

To synchronize the configuration between two switches, a peer must be configured and enabled on each switch. Switches being synchronized must use the same administrator password. Peers are sent SLB, FILT, and VRRP configuration updates using `/oper/slb/synch`.

NOTE – When using both VRRP and GSLB, the `/cfg/sys/wport` (Browser-Based Interface port) value of the target switch (the switch that is being synchronized to) must be changed to a port other than port 80 before VRRP synchronization begins.

Table 6-62 Synch Peer Switch Menu Options (/cfg/slb/synch)

Command Syntax and Usage
peer *<peer switch number (1-4)>* Displays the Peer switch menu. To view a list of menu options, see page 237.
filt disable|enable (**d**|**e**) Enables or disables synchronizing filter configuration.
ports disable|enable (**d**|**e**) Enables or disables synchronizing Layer 4 port configuration.
prios disable|enable (**d**|**e**) Enables or disables synchronizing VRRP priorities.
pips disable|enable (**d**|**e**) Enables or disables the current synchronizing proxy IP address.
current Displays the current configuration for this peer switch.

/cfg/slb/sync/peer <peer switch>

Peer Switch Configuration

NOTE – The *peer switch number* (1-4) represents the peer switch you wish to configure.

```
[Peer switch 1 Menu]
      addr    - Set peer switch IP address
      enable  - Enable peer switch
      disable - Disable peer switch
      delete  - Delete peer switch
      current - Display current peer switch configuration
```

The peer switch menu for each peer switch (1-4) allows you to set the IP address of the peer switch, enable, disable, or delete the current peer switch, or display the current configuration.

To synchronize the configuration between two switches, a peer must be configured and enabled on each switch. Switches being synchronized must use the same administrator password.

Table 6-63 Peer Switch Menu Options (/cfg/slb/sync/peer)

Command Syntax and Usage
addr *<IP address, such as 192.4.17.101>* Sets the peer switch IP address in dotted decimal notation.
enable Enables the peer for this switch.
delete Deletes the peer for this switch.
current Displays the current peer switch configuration.

/cfg/slb/adv

Advanced Layer 4 Configuration

```
[Layer 4 Advanced Menu]
      script  - Scriptable health check menu
      direct  - Enables or disables Direct Access Mode
      grace   - Enables or disables graceful real server failure
      octets  - Enables or disables octet counters
      bclose  - Enables or disables bidirectional close for DAM
      tpcp    - Enables or disables Transparent Proxy Cache Protocol
      imask   - Set virtual and real IP address mask
      mnet    - Set management network
      mmask   - Set management subnet mask
      pmask   - Set persistent mask
      timeout - Set minutes inactive session remains open
      secret  - Set RADIUS secret
      current - Display current Layer 4 advanced configuration
```

Table 6-64 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage
script *<health script number (1-8)>* Displays the Scriptable Health Check Menu. For more information, see page 240.
direct disable|enable (**d|e**) Enables or disables Direct Access Mode (DAM) to real servers/services. This option also allows any virtual server to load balance any real server.
grace disable|enable (**d|e**) Enables or disables graceful real server failure. When enabled, state information about existing sessions is maintained and traffic associated with existing sessions continues to be sent to the server. Connection requests to, and traffic associated with, other load-balanced services continue to be processed by the server.
octets disable|enable (**d|e**) Enables or disables octet counters.
bclose disable|enable (**d|e**) Enables or disables bidirectional close for DAM.
tpcp disable|enable (**d|e**) Enables or disables Transparent Proxy Cache Protocol (TPCP).
imask *<IP subnet mask, such as 255.255.255.0>* Configures the real and virtual server IP address mask using dotted decimal notation. For more information, see "Configuring the imask" on page 240.

Table 6-64 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage
mnet *<server IP address, such as 192.4.17.0>* 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` option.
mmask *<IP subnet mask, such as 255.255.255.0>* Sets management subnet mask. This IP address mask is used with the `mnet` to select management traffic that is allowed access to direct real servers.
pmask *<IP subnet mask, such as 255.255.255.0>* Sets persistent mask, which is used to increase granularity of persistence to a subnet level.
timeout *<even number of minutes [4-60]>* Sets 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. If a client application is abnormally terminated by the client's system, TCP 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 `time-out` 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. *Note:* A session entry can be removed before the timeout value *only* for filtered or TCP sessions in which a FIN/RST is seen. All other connections such as UDP, ICMP, and normal TCP sessions, must time out through inactivity, using the time-out value.
secret *<1-32 character secret>* To perform application health checking to a RADIUS server, the network administrator must configure two parameters in the switch: the `/cfg/slb/adv/secret` value and the `content` parameter with a *username:password* value. The `secret` value is a field of up to 32 alphanumeric characters that is used by the switch to encrypt a password during the RSA Message Digest Algorithm (MD5). It is also used by the RADIUS server to decrypt the password during verification.
current Displays the current Layer 4 advanced configuration.

Configuring the imask

The imask determines how many different IP addresses each real and virtual server will represent and to which it will respond. 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 server IP address 172.16.10.45, the unmasked portion of the virtual server 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.

/cfg/slb/adv/script <*script number*>

Scriptable Health Checks Configuration

NOTE – The *script number* (1-8) represents the script you wish to modify.

```
[Health Script 1 Menu]
      open    - Add open command to end of script
      send    - Add send command to end of script
      expect  - Add expect command to end of script
      close   - Add close command to end of script
      remove  - Remove last command from script
      delete  - Delete script
      current - Display current script configuration
```

Scriptable health checks provide a robust and extensible way to health check a group of servers. Users can use scripts to define their own health checks rather than rely on hard-coded health checks that may not meet their particular needs.

The approach introduces ASCII-based scripts to control how a group of servers are health checked. Only TCP services can be health checked, since UDP protocols are usually not ASCII-based.

A very rudimentary interface controls the addition and deletion of ASCII commands to each script. New commands are added and removed from the end of the script. Commands exist to open a connection to a specific TCP port, send an ASCII request to the server, expect an ASCII string, and close a connection. The string configured with an `expect` command is searched for in each response packet. If it is not seen in any response packet before the real server health check interval expires, the server does not pass the `expect` step and fails the health check.

A script can contain any number of these commands, up to the allowable number of characters that a script supports.

Example:

```
Current health script 1:
    open 80
    send GET / HTTP/1.0\r\n\r\n
    expect HTTP/1.0 200
    close

The ability to use a script to health check a real server group was
added:

>> Real server group 1# health ?
Usage: health
icmp|tcp|http|dns|pop3|smtp|nntp|ftp|imap|radius|sslh|script<n>
```

/cfg/vrrp
VRRP Configuration

```
[Virtual Router Redundancy Protocol Menu]
      vr      - VRRP Virtual Router menu
      group   - VRRP Virtual Router Group menu
      if      - VRRP Interface menu
      track   - VRRP Priority Tracking menu
      hotstan - Enable/disable hot-standby processing
      on      - Globally turn VRRP ON
      off     - Globally turn VRRP OFF
      current - Display current VRRP configuration
```

Virtual Router Redundancy Protocol (VRRP) on Alteon Web switches provides redundancy between routers in a LAN. If the master fails, a backup virtual routers assumes routing authority and takes control of the shared virtual router IP address. This is set up by configuring the same virtual router IP address and ID number on each participating VRRP-capable routing device.

Table 6-65 Virtual Router Redundancy Protocol Menu Options (/cfg/vrrp)

Command Syntax and Usage
vr *<virtual router number (1-1024)>* Displays the VRRP Virtual Router Menu. This menu is used for configuring up to 1024 virtual routers on the switch. To view menu options, see page 243.
group Displays the VRRP Virtual Router Group Menu. To view menu options, see page 247.
if *<interface number (1-1024)>* Displays the VRRP Virtual Router Interface Menu. This menu is used for setting VRRP authentication parameters for up to 1024 IP interfaces. To view menu options, see page 250.
track Displays the VRRP Tracking Menu. This menu is used for weighting the criteria used when modifying priority levels in the master router election process. To view the menu options, see page 249.
on Globally enables VRRP on this switch.
off Globally disables VRRP on this switch.
current Displays the current VRRP parameters.

/cfg/vrrp/vr <virtual router number>

Virtual Router Configuration

NOTE – The *virtual router number* (1-1024) represents the virtual router you wish to configure.

```
[VRRP Virtual Router 1 Menu]
      vrid    - Set virtual router ID
      addr    - Set IP address
      if      - Set interface number
      prio    - Set renter priority
      adver   - Set advertisement interval
      preem   - Enables or disables preemption
      share   - Enables or disables sharing
      track   - Priority tracking menu
      enable - Enable virtual router
      disable - Disable virtual router
      delete - Delete virtual router
      current - Display current VRRP virtual router configuration
```

This menu is used for configuring up to 1024 possible virtual routers for this switch. A virtual router is defined by its virtual router ID and IP address. On each VRRP-capable routing device participating in redundancy for this virtual router, a virtual router will be configured to share the same virtual router ID and IP address.

Table 6-66 VRRP Virtual Router Menu Options (/cfg/vrrp/vr)

Command Syntax and Usage
vrid *<virtual router ID (1-255)>* Defines the virtual router ID. This is used in conjunction with the IP address to define a virtual router on this switch. To create a pool of VRRP-enabled routing devices that can provide redundancy to each other, each participating VRRP device must be configured with the same virtual router—one that shares the same Virtual Router ID (VRID) and IP address combination. The VRID for standard virtual routers (where the virtual router IP address is not the same as any virtual server) can be any integer between 1 and 255. All VRID values must be unique within the VLAN to which the virtual router's IP interface belongs.
addr *<IP address, such as 192.114.17.101>* Defines the IP address for this virtual router using dotted decimal notation. This IP address is used in conjunction with the VRID to configure the same virtual router on each participating VRRP device. The default address is 0.0.0.0.

Table 6-66 VRRP Virtual Router Menu Options (/cfg/vrrp/vr) (Continued)

Command Syntax and Usage
`if` *<interface number (1-1024)>* Selects a switch IP interface (from 1 to 1024). If the IP interface has the same Virtual Router IP address as the `addr` option above, this switch is considered the "owner" of the defined virtual router. An owner has a special priority of 255 (highest) and will always assume the role of master router, even if it must preempt another virtual router that has assumed master routing authority. This preemption occurs even if the `preem` option is disabled.
`prio` *<priority (1-254)>* Defines the election priority bias for this virtual server. This can be any integer from 1 to 254. The default value is 100. During the master router election process, the routing device with the highest virtual router priority number wins. If there is a tie, the device with the highest IP interface address wins. If this virtual router's IP address (`addr`) is the same as the one used by the IP interface, the priority for this virtual router will automatically be set to 255. When priority tracking is used (`/cfg/vrrp/track` or `/cfg/vrrp/vr #/track`), this base priority value can be modified according to a number of performance and operational criteria.
`adver` *<seconds (1-255)>* Defines the time interval between VRRP master advertisements. The range is from 1 to 255 seconds. The default is 1.
`preem disable|enable` (**`d|e`**) Enables or disables master preemption. When enabled, if this virtual router is in backup mode but has a higher priority than the current master, this virtual router will preempt the lower priority master and assume control. Even when `preem` is disabled, this virtual router will always preempt any other master if this switch is the owner (the IP interface address and virtual router `addr` are the same). By default, this option is enabled.
`share disable|enable` (**`d|e`**) Enables or disables virtual router sharing, an Alteon WebSystems proprietary extension to VRRP. When enabled, this switch will process any traffic addressed to this virtual router, even when in backup mode. By default, this option is enabled.
`track` Displays the VRRP Priority Tracking Menu for this virtual router. Tracking is an Alteon WebSystems proprietary extension to VRRP, used for modifying the standard priority system and for electing the master router. Tracking is not needed if sharing is enabled. For more information, see "Virtual Router Priority Tracking Configuration" on page 245.
`enable` Enables this virtual router.

Table 6-66 VRRP Virtual Router Menu Options (/cfg/vrrp/vr) (Continued)

Command Syntax and Usage
disable Disables this virtual router.
delete Deletes this virtual router from the switch configuration.
current Displays the current configuration information for this virtual router.

/cfg/vrrp/vr <*virtual router number*>/track

Virtual Router Priority Tracking Configuration

NOTE – The *virtual router number* (1-1024) represents the virtual router of the priority tracking you wish to configure.

```
[VRRP Virtual Router 1 Priority Tracking Menu]
      vrs     - Enables or disables tracking other virtual routers
      ifs     - Enables or disables tracking other interfaces
      ports   - Enables or disables tracking VLAN switch ports
      l4pts   - Enables or disables tracking L4 switch ports
      reals   - Enables or disables tracking L4 real servers
      hsrp    - Enables or disables tracking HSRP
      current - Display current VRRP virtual router configuration
```

This menu is used for modifying the priority system used when electing the master router from a pool of virtual routers. Various tracking criteria can be used to bias the election results. Each time one of the tracking criteria is met, the priority level for the virtual router is increased by an amount defined through the VRRP Tracking Menu (see page 251).

Criteria are tracked dynamically, continuously updating virtual router priority levels when enabled. If the virtual router preemption option (see `preem` in Table 6-66 on page 243) is enabled, this virtual router can assume master routing authority when its priority level rises above that of the current master.

Some tracking criteria (virtual routers, interfaces and ports) apply to standard virtual routers, also called *virtual interface routers*. Other tracking criteria (Layer 4 switch ports, real servers and HSRP) apply to extended virtual routers, or *virtual server routers*, which perform Layer 4 server load balancing functions in addition to their standard VRRP operation. A virtual *server* router is defined as any virtual router whose IP address is the same as any configured virtual server IP address.

Table 6-67 VRRP Priority Tracking Menu Options (/cfg/vrrp/vr #/track)

Command Syntax and Usage
vrs disable|enable (**d**|**e**) When enabled, the priority for this virtual router will be increased for each virtual router in master mode on this switch. This is useful for ensuring that traffic for any particular client/server pairing are handled by the same switch, increasing routing and load-balancing efficiency.
ifs disable|enable (**d**|**e**) When enabled, the priority for this virtual router will be increased for each IP interface active on this switch. An IP interface is considered active when there is at least one active port on the same VLAN. This helps elect the virtual routers with the most available routes as the master.This command is disabled by default.
ports disable|enable (**d**|**e**) When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. This helps elect the virtual routers with the most available ports as the master. This command is disabled by default.
l4pts disable|enable (**d**|**e**) When enabled for virtual server routers, the priority for this virtual router will be increased for each physical switch port which has active Layer 4 processing on this switch. This helps elect the main Layer 4 switch as the master. This command is disabled by default.
reals disable|enable (**d**|**e**) When enabled for virtual server routers, the priority for this virtual router will be increased for each healthy real server behind the virtual server IP address of the same IP address as the virtual router on this switch. This helps elect the switch with the largest server pool as the master, increasing Layer 4 efficiency. This command is disabled by default.
hsrp disable|enable (**d**|**e**) Hot Standby Router Protocol (HSRP) is used with some types of routers for establishing router failover. In networks were HSRP is used, enable this switch option to increase the priority of this virtual router for each Layer 4 client-only port that receives HSRP advertisements. Enabling `hsrp` helps elect the switch closest to the master HSRP router as the master, optimizing routing efficiency. This command is disabled by default.
current Displays the current configuration for priority tracking for this virtual router.

/cfg/vrrp/group

Virtual Router Group Configuration

```
[VRRP Virtual Router Group Menu]
      track   - Priority tracking menu
      vrid    - Set virtual router ID
      if      - Set interface number
      prio    - Set renter priority
      adver   - Set advertisement interval
      preem   - Enable/disable preemption
      share   - Enable/disable sharing
      enable  - Enable virtual router
      disable - Disable virtual router
      delete  - Delete virtual router
      current - Display current VRRP virtual router configuration
```

The menu allows you to configure the virtual router group.

Table 6-68 Virtual Router Group Menu Options (/cfg/vrrp/group)

Command Syntax and Usage
track Displays the Priority Tracking Menu. To view the menu options, see page 249.
vrid *<virtual router ID (1-255)>* Defines the virtual router ID.
if *<interface number (1-1024)>* Selects a switch IP interface number.
prio *<priority (1-254)>* Sets the renter priority. A renter is a VRRP router that does not own the IP address it uses.
adver *<1-255>* Defines the time interval between VRRP master advertisements in seconds. The value range is from 1 to 255. The default value is 1 second.
preem disable|enable (**d|e**) Enables or disables master preemption. When enabled, if this virtual router is in backup mode but has a higher priority than the current master, this virtual router will preempt the lower priority master and assume control. Even when **preem** is disabled, this virtual router will always preempt any other master if this switch is the owner (the IP interface address and virtual router address are the same). By default, this option is enabled.

Table 6-68 Virtual Router Group Menu Options (/cfg/vrrp/group) (Continued)

Command Syntax and Usage
share disable|enable (**d|e**) Enables or disables virtual router sharing. When enabled, a master/slave pair of switches share the workload. If the slave receives a request, it will process that request instead of forwarding it to the master switch.
enable Enables the virtual router group.
disable Disables the virtual router group.
delete Deletes the virtual router group.
current Displays the current configuration information for the virtual router group.

/cfg/vrrp/group/track

Virtual Router Group Priority Tracking Configuration

```
[VRRP Virtual Router  Priority Tracking Menu]
      vrs     - Enable/disable tracking master virtual routers
      ifs     - Enable/disable tracking other interfaces
      ports   - Enable/disable tracking VLAN switch ports
      l4pts   - Enable/disable tracking L4 switch ports
      reals   - Enable/disable tracking L4 real servers
      hsrp    - Enable/disable tracking HSRP
      current - Display current VRRP virtual router configuration
```

Tracking only affects hot-standby and active-standby configurations. It has no effect when sharing (active-active) configurations are used.

For tracking to have an effect on virtual router operation, preemption must be enabled (`/cfg/vrrp/vr`).

Table 6-69 Virtual Priority Tracking Menu Options (/cfg/vrrp/group/track)

Command Syntax and Usage
vrs disable|enable (**d|e**) When enabled, the priority for this virtual router will be increased for each virtual router in master mode on this switch. This is useful for ensuring that traffic for any particular client/server pairing is handled by the same switch, increasing routing and load balancing efficiency. This command is disabled by default.
ifs disable|enable (**d|e**) When enabled, the priority for this virtual router will be increased for each other IP interface active on this switch. An IP interface is considered active when there is at least one active port on the same VLAN. This helps elect the virtual routers with the most available routes as the master. This command is disabled by default.
ports disable|enable (**d|e**) When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. This helps elect the virtual routers with the most available ports as the master. This command is disabled by default.
l4pts disable|enable (**d|e**) When enabled for virtual server routers, the priority for this virtual router will be increased for each physical switch port which has active Layer 4 processing on this switch. This helps elect the main Layer 4 switch as the master. This command is disabled by default.

Table 6-69 Virtual Priority Tracking Menu Options (/cfg/vrrp/group/track)

Command Syntax and Usage
reals disable|enable (**d**|**e**) When enabled for virtual server routers, the priority for this virtual router will be increased for each healthy real server behind the virtual server IP address of the same IP address as the virtual router on this switch. This helps elect the switch with the largest server pool as the master, increasing Layer 4 efficiency. This command is disabled by default.
hsrp disable|enable (**d**|**e**) Hot Standby Router Protocol (HSRP) is used with some types of routers for establishing router failover. In networks were HSRP is used, enable this switch option to increase the priority of this virtual router for each Layer 4 client-only port that receives HSRP advertisements. Enabling HSRP helps elect the switch closest to the master HSRP router as the master, optimizing routing efficiency. This command is disabled by default.
current Displays the current configuration for priority tracking for this virtual router.

/cfg/vrrp/if <interface number>

VRRP Interface Configuration

NOTE – The *interface number* (1-1024) represents the IP interface on which authentication parameters must be configured.

```
[VRRP Interface 1 Menu]
     auth    - Set authentication types
     passw   - Set plain-text password
     delete  - Delete interface
     current - Display current VRRP interface configuration
```

This menu is used for configuring VRRP authentication parameters for the IP interfaces used with the virtual routers.

Table 6-70 VRRP Interface Menu Options (/cfg/vrrp/if).

Command Syntax and Usage
auth none|password Defines the type of authentication: ■ `none` No authentication used. ■ `password` Password authentication will be used.

Table 6-70 VRRP Interface Menu Options (/cfg/vrrp/if).

Command Syntax and Usage
passw *<password>* Defines a plain-text password up to eight characters long. This password will be added to each VRRP packet transmitted by this interface when password authentication is chosen (see `auth` above).
delete Clears the authentication configuration parameters for this IP interface. The IP interface itself is not deleted.
current Displays the current configuration for this IP interface's authentication parameters.

/cfg/vrrp/track

VRRP Tracking Configuration

```
[VRRP Tracking Menu]
      vrs     - Set priority increment for virtual router tracking
      ifs     - Set priority increment for IP interface tracking
      ports   - Set priority increment for VLAN switch port tracking
      l4pts   - Set priority increment for L4 switch port tracking
      reals   - Set priority increment for L4 real server tracking
      hsrp    - Set priority increment for HSRP tracking
      current - Display current VRRP Priority Tracking configuration
```

This menu is used for setting weights for the various criteria used to modify priority levels during the master router election process. Each time one of the tracking criteria is met ("Virtual Router Priority Tracking Configuration" on page 245), the priority level for the virtual router is increased by an amount defined through this menu.

Table 6-71 VRRP Tracking Menu Options (/cfg/vrrp/track)

Command Syntax and Usage
vrs *<0-254>* Defines the priority increment value (0 through 254) for virtual routers in master mode detected on this switch. The default value is 2.
ifs *<0-254>* Defines the priority increment value (0 through 254) for active IP interfaces detected on this switch. The default value is 2.

Table 6-71 VRRP Tracking Menu Options (/cfg/vrrp/track) (Continued)

Command Syntax and Usage
ports *<0-254>* Defines the priority increment value (0 through 254) for active ports on the virtual router's VLAN. The default value is 2.
l4pts *<0-254>* Defines the priority increment value (0 through 254) for physical switch ports with active Layer 4 processing. The default value is 2.
reals *<0-254>* Defines the priority increment value (0 through 254) for healthy real servers behind the virtual server router. The default value is 2.
hsrp *<0-254>* Defines the priority increment value (0 through 254) for switch ports with Layer 4 client-only processing that receive Hot Standby Router Protocol (HSRP) broadcasts. The default value is 10.
current Displays the current configuration of priority tracking increment values.

NOTE – These priority tracking options only define increment values. These options do not affect the VRRP master router election process until options under the VRRP Virtual Router Priority Tracking Menu (see page 245) are enabled.

CHAPTER 7

The Operations Menu

The Operations Menu is generally used for commands that immediately affect operation of the switch 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), with the understanding that when the switch is reset, the port returns to its normally configured operation.

To make permanent changes to switch configurations, use commands under Chapter 6, "The Configuration Menu."

/oper
Operations Menu

```
[Operations Menu]
      port    - Operational Port menu
      sp      - Operational SP menu
      lcm     - Operational LCM menu
      slb     - Operational Layer 4 menu
      vrrp    - Operational Virtual Router Redundancy menu
      ip      - Operational IP Menu
      swkey   - Enter key to enable software feature
      rmkey   - Enter software feature to be removed
```

The commands of the Operations Menu enable you to alter switch operational characteristics without affecting switch configuration.

Table 7-1 Operations Menu Options (/oper)

Usage
`port` *<port and slot number, such as b12>* Displays the Operational Port menu. To view menu options, see page 255.
`sp` Displays the Operational Switch Processor menu. To view menu options, see page 256.
`lcm` Displays the Operational Line-Card Module menu. To view menu options, see page 257.
`slb` Displays the Operational Layer 4 menu. To view menu options, see page 258.
`vrrp` Displays the Operational Virtual Router Redundancy Protocol Menu. To view menu options, see page 259.
`ip` Displays the Operational IP Menu. To view menu options, see page 259.
`swkey` Enter the 16 digit key to activate an optional software feature, such as GSLB. For more information, see page 260.
`rmkey` This option is used for deactivating any optional software, such as GSLB. For more information, see page 261.

/oper/port <slot-number port-number> Operations-Level Port Options

```
[Operations Port 1 Menu]
      up      - Bring port up
      down    - Shut port down
      rmon    - Up/down RMON for port
      current - Current port state
```

NOTE – The `ena` and `dis` commands, previously in Release 6.0, have been renamed `up` and `down`. The `ena` and `dis` commands always apply to *configuration*, whereas `up` and `down` always apply to *operation.*

Operations-level port options are used for temporarily disabling or enabling a port and for changing RMON status on a port.

Table 7-2 Operations-Level Port Menu Options (/oper/port)

Command Syntax and Usage
up Temporarily brings up the port. The port will be returned to its configured operation mode when the switch is reset.
down Temporarily brings down the port. The port will be returned to its configured operation mode when the switch is reset.
rmon Temporarily brings up/down RMON on the port. The port will be returned to its configured operation mode when the switch is reset.
current Displays the current operational status of the port.

/oper/sp
Operations-Level Switch Processor Options

```
[SP Operations Menu]
      up      - Bring up one or all SPs
      down    - Shut down one or all SPs
      post    - Run POST on one or all SPs
```

NOTE – The following commands will bring up or down all ports on the switch processor.

Table 7-3 Operations-Level Switch Processor Menu Options (/oper/sp)

Command Syntax and Usage
up *<SP as port (such as b12) or index (such as x5)>* Temporarily brings up the specified switch processor.
down *<SP as port (such as b12) or index (such as x5)>* Temporarily brings down the specified switch processor.
post *<SP as port (such as b12) or index (such as x5)>* Runs Power On Self Test (POST) on the specified switch processor.

/oper/lcm <module number>
Operations-Level Line Card Module Options

```
[LCM Operations Menu]
      up      - Bring an LCM up
      down    - Shut an LCM down
```

Table 7-4 Operations-Level Switch Processor Menu Options (/oper/lcm)

Command Syntax and Usage
up *<LCM index (A-D)>* Temporarily brings up the specified line card module (LCM).
down *<LCM index (A-D)>* Temporarily brings down the specified line card module.

/oper/slb
Operations-Level SLB Options

```
[Server Load Balancing Operations Menu]
      synch   - Synchronize SLB, FILT, and VRRP configuration on peers
      enable  - Enable real server
      disable - Disable real server
      clear   - Clear session table on port
      current - Current layer 4 operational state
```

When the optional Layer 4 software is enabled, the operations-level server load balancing options are used for temporarily disabling or enabling real servers and synchronizing the configuration between the switches.

The options are described in the following table:

Table 7-5 Server Load Balancing Operations Menu Options (/oper/slb)

Command Syntax and Usage
synch Synchronizes the SLB, filter, and VRRP configuration on a peer switch (a switch that owns the IP address). To take effect, VRRP must be globally enabled on the peer switch.
enable *<real server number (1-1024)>* Temporarily enables a real server. The real server will be returned to its configured operation mode when the switch is reset.
disable *<real server number (1-1024)>* Temporarily disables a real server, removing it from operation within its real server group and virtual server. Sessions are not sent to the real server and old sessions are retained. The real server will be returned to its configured operation mode when the switch is reset.
clear *<port as slot and port, such as b12>* Clears the session table for a specific port and allows port filter changes to take effect immediately. Note: This disrupts current Server Load Balancing and Application Redirection sessions.
current Displays the current SLB operational state.

/oper/vrrp
Operations-Level VRRP Options

```
[VRRP Operations Menu]
        back    - Set virtual router to backup
```

This menu is used to force a master virtual router to become backup router.

Table 7-6 Virtual Router Redundancy Operations Menu Options (/oper/vrrp)

Command Syntax and Usage
back *<virtual router number (1-1024)>* Forces the specified master virtual router on this switch into backup mode. This is generally used for passing master control back to a preferred switch once the preferred switch has been returned to service after a failure. When this command is executed, the current master gives up control and initiates a new election, by temporarily advertising its own priority level as 0 (lowest). After the new election, the virtual router that had been forced into backup mode by this command, will resume master control in the following cases: ■ This switch owns the virtual router (the IP addresses of the virtual router and its IP interface are the same) ■ This switch's virtual router has a higher priority and preemption is enabled. ■ There are no other virtual routers available to take master control.

/oper/ip
Operations-Level IP Options

```
[IP Operations Menu]
      bgp   - Operational Border Gateway Protocol Menu
```

Table 7-7 Operations Level Ip Options (/oper/ip)

Command Syntax and Usage
bgp Displays the Border Gateway Protocol menu. To view the menu options, see page 260.

/oper/ip/bgp

Operations-Level BGP Options

```
[Border Gateway Protocol Operations Menu]
      start   - Start peer session
      stop    - Stop peer session
      current - Current BGP operational state
```

Table 7-8 Operations Level BGP Options (/oper/ip/bgp)

Command Syntax and Usage
start *<peer number (1-16)>* Temporarily starts the peer session. The peer will be returned to its configured operation mode when the switch is reset.
stop *<peer number (1-16)>* Temporarily stops the peer session. The peer will be returned to its configured operation mode when the switch is reset.
current Displays the current state of the Border Gateway Protocol.

/oper/swkey

Activating Optional Software

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 WebSystems 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 WebSystems, and provide the MAC address of the Web OS switch that will run the optional software. Alteon WebSystems 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 received your License Password, perform the following actions:

1. **After the switch reboots, connect to the switch's command line interface and log in as the administrator (see Chapter 1, "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.
```

/oper/rmkey
Removing Optional Software

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 [L4]: l4
```

CHAPTER 8

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 the following options:

- 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
      config  - 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
```

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 from the CD-ROM 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`. The switch also stores 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 the 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 Configuration" on page 206).

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.**

   To minimize mistakes, you must type the full name (“image1”, “image2”, or “boot”).

```
Enter 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 as to which image is currently set to be loaded at the next reset and prompts you to enter a new choice:

```
Currently set to use software "image2" on next boot.
Specify new image to use [image1|image2]:
```

   To minimize mistakes, you must type the full name (“image1”, “image2”, or “boot”).

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# config
```

2. **Enter the name of the configuration block you want the switch to use:**

The system informs you 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 9

The Maintenance Menu

The Maintenance Menu is used to manage dump information and forward database information. It also includes a debugging menu to help with troubleshooting.

NOTE – To use the Maintenance Menu, you must be logged in to the switch as the administrator.

/maint
Maintenance Menu

```
[Maintenance Menu]
      dump    - FLASH Dump Manipulation Menu
      sys     - System Maintenance Menu
      sp      - SP Maintenance Menu
      fdb     - Forwarding Database Manipulation Menu
      debug   - Debugging Menu
      arp     - ARP Cache Manipulation Menu
      route   - IP Route Manipulation Menu
      panic   - Dump MP information to FLASH and reboot
      tsdump  - Tech support dump
```

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.

Table 9-1 Maintenance Menu Options (/maint)

Command Syntax and Usage
dump Displays the FLASH Dump Manipulation Menu. To view menu options, see page 271.
sys Displays the System Maintenance Menu. To view menu options, see page 274.
sp *<port number, such as b12>* Displays the SP Maintenance Menu. To view menu options, see page 275.
fdb Displays the Forwarding Database Manipulation Menu. To view menu options, see page 275.
debug Displays the Debugging Menu. To view menu options, see page 277.
arp Displays the ARP Cache Manipulation Menu. To view menu options, see page 279.
route Displays the IP Route Manipulation Menu. To view menu options, see page 281.
panic Dumps MP information to FLASH and reboots. For more information, see page 282.
tsdump Dumps all switch information, statistics, and configurations.

/maint/dump

FLASH Dump Manipulation Options

```
[Dump Menu]
      current - Display currently-available FLASH dumps
      uuout   - Uuencode a FLASH dump to standard out
      tftp    - TFTP a FLASH dump to tftpserver
      delete  - Delete a FLASH dump
      mksave  - Mark a FLASH dump as "saved"
```

Table 9-2 FLASH Dump Maintenance Menu Options (/maint/dump)

Command Syntax and Usage

current
Displays currently available FLASH dumps. Information similar to the following is shown:
`index  MP/SP  ver     date and time taken        cause     state`
`-----  -----  ---  -------------------------  ---------  --------`
`  0     MP     1   11:08:56 Thu May 25, 2000     W/DOG     unread`
`----------------------------------------------------------------------`

uuout Uuencodes a FLASH dump to standard out. For more information, see page 272.
tftp Sends a FLASH dump to a TFTP server. For more information, see page 273.
delete Deletes a FLASH dump. For more information, see page 273.
mksave Marks a FLASH dump as “saved.” A new dump is called UNREAD, and it will not be overlaid with a later dump until it has been read (with `uuout` or TFTP), or deleted. Marking a dump as SAVED will prevent it from being overlaid, even after it is read.

/maint/dump/uuout

Uuencode FLASH Dump

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 WebSystems Customer Support for help in 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 `uuout` command. This will ensure that you do not lose any information. Once entered, the `uuout` command will cause data to be displayed on your screen and copied into the file.

Using the `uuout` 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 273.

To access dump information, at the `Maintenance#` prompt, enter:

```
Maintenance# dump/uuout
```

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.
```

/maint/dump/tftp *server/filename*

TFTP System Dump Put

Use this command to put (save) the system dump to a server via TFTP.

NOTE – If the TFTP server is running SunOS or the Solaris operating system, the specified target dump file must exist *prior* to executing the `tftp` command and must be writable (set with proper permission, and not locked by any application). The contents of the specified file will be replaced with the current dump data.

To save dump information via TFTP, at the `Dump#` prompt, enter:

```
Maintenance# tftp server filename
```

where *server* is the TFTP server IP address or hostname, and *filename* is the target dump file.

/maint/dump/delete

Clearing Dump Information

To clear dump information from FLASH memory, at the `Maintenance#` prompt, enter:

```
Maintenance# delete
```

/maint/sys

System Maintenance Options

This menu is reserved for use by Alteon WebSystems Customer Support. The options are used to perform system debugging.

```
[System Maintenance Menu]
      sw      - Display software build information
      flags   - Set NVRAM flag word; "bits" to list bits
      tmask   - Set current MP trace mask
      sptmask - Set default SP trace mask
```

Table 9-3 System Maintenance Menu Options (/maint/sys)

Command Syntax and Usage
sw Displays software build information. To view a sample, see page 275.
flags Sets an NVRAM flag word that can be used by Customer Support to turn on special messages or additional information used for debugging purposes. Enter the flag in hexadecimal or standard format.
tmask Sets the current Management Processor trace mask that can be used by Customer Support to turn on special messages or additional information used for debugging purposes. Enter the mask in hexadecimal or standard format.
sptmsk Sets the default Switch Processor trace mask that can be used by Customer Support to turn on special messages or additional information used for debugging purposes. Enter the mask in hexidecimal or standard format.

/maint/sys/sw

Software Build Information

```
Software build information:
  Software build information:
  FW_VERSION: #2 Thu Jun 22 14:05:57 PDT 2000
  FW_COMPILE_TIME: 14:05:57
  FW_COMPILE_BY: devrcs
  FW_SANDBOX: /projects/sandbox/gne_7_0_24/src
  FW_COMPILER: gcc version 2.7.2
```

/maint/sp *<switch processor index>*

Switch Processor Maintenance Options

```
[Maint SP index 2 (port A3) Menu]
      tmask   - Temporarily set SP's trace mask
```

Table 9-4 Switch Processor Maintenance Menu Options (/maint/sp)

Command Syntax and Usage
tmask Temporarily sets the switch processor's trace mask, which can be used by Customer Support to turn on special messages or additional information used for debugging purposes.

/maint/fdb

Forwarding Database Options

```
[FDB Manipulation 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
      dump    - Show all FDB entries
      count   - Show count of FDB entries
      delete  - Delete an FDB entry
      clear   - Clear entire FDB
```

The Forwarding Database Manipulation Menu can be used to view forwarding database entries, find entries based on MAC address, port, VLAN, and switch processor. It can also be used 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.

Table 9-5 FDB Manipulation Menu Options (/maint/fdb)

Command Syntax and Usage
`find` *<MAC address> <VLAN>* Displays a single database entry by its MAC address. 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`.
`port` *<port number and slot number, (such as b12)>* Displays all FDB entries for a particular port.
`vlan` *<VLAN number (1-4094)>* Displays all FDB entries on a single VLAN.
`dump` Displays all entries in the FDB. For more information, see page 64.
`count` Displays the count of all FDB entries.
`delete` *<MAC address> <VLAN>* Removes a single FDB entry.
`clear` Clears the entire FDB from switch memory.

/maint/debug

Debugging Options

```
[Miscellaneous Debug Menu]
      tbuf    - Display MP trace buffer
      xtbuf   - Display MP XTRACE buffer
      sptbuf  - Display SP trace buffer
      spxtbuf - Display SP XTRACE buffer
```

The Miscellaneous Debug Menu displays trace buffer information about events that can be helpful in understanding switch operation. You can view the following information using the debug menu:

- Events traced by the Management Processor (MP)
- Events traced by a Switch Processor (SP)
- Events traced to a buffer area when a reset occurs

If the switch resets because of a crash, the MP trace buffer is saved into the MP dump, and the SP trace buffer is saved into the SP dump, both of which are written to FLASH. No trace buffers are saved if the switch is powered off.

The output from these commands can be interpreted by the Alteon WebSystems Customer Support organization.

Table 2 Miscellaneous Debug Menu Options (/maint/debug)

Command Syntax and Usage
tbuf Displays the Management Processor trace buffer. The MP trace buffer displays normal traces containing detailed operational events that may give details about what events were occurring at the moment of a crash. Header information similar to the following is shown: `MP trace buffer at 18:46:10 Thu Jun 29, 2000; mask: 0x01dff6df` The buffer information is displayed after the header.
xtbuf Displays the Management Processor xtrace buffer. The MP xtrace buffer displays more significant operational events going back further in time. Header information similar to the following is shown: `MP XTRACE buffer at 18:48:16 Thu Jun 29, 2000` The buffer information is displayed after the header
sptb *<slot-number port-number (such as b12)>* Displays the Switch Processor trace buffer. The SP trace buffer displays normal traces containing detailed operational events that may give details about what events were occuring at the moment of a switch processor crash. Header information similar to the following is shown: `SP index 0 (A1) at 18:48:26 Thu Jun 29, 2000; mask: 0x01dff6df` The buffer information is displayed after the header.
spxtbuf *<slot-number port-number (such as b12)>* Displays the Switch Processor xtrace buffer. To view a sample, see "/maint/debug/spxtbuf" on page 279.

/maint/debug/spxtbuf

SP Xtrace Buffer

Following is a sample of the output of the SP trace buffer. The first column is the label, column three is the internal time stamp, and columns two and four through seven are debug information.

```
>> Miscellaneous Debug# spx a1
SP index 0 (A1) XTRACE buffer at 19:18:00 Mon Mar 12, 2001
ST_PATH!  17000000 0018475a  00000000 00000000 000083d4 00000036
ST_PATH!  17000000 001655bc  00000000 00000000 0002776e 00000026
:
:
```

/maint/arp

ARP Cache Options

```
[Address Resolution Protocol Menu]
      find    - Show a single ARP entry by IP address
      port    - Show ARP entries on a single port
      refsp   - Show ARP entries referenced by a single port
      vlan    - Show ARP entries on a single VLAN
      add     - Add a permanent ARP entry
      delete  - Delete an ARP entry
      clear   - Clear ARP cache
      dump    - Show all ARP entries
      addr    - Show ARP address list
```

Table 9-6 Address Resolution Protocol Menu Options (/maint/arp)

Command Syntax and Usage
find *<IP address, such as 192.4.17.101>* Shows a single ARP entry by IP address.
port *<slot-number port-number>* Shows ARP entries on a single port.
refsp *<slot-number port-number>* Shows all ARP entries referenced by a single port.
vlan *<1-4094>* Shows ARP entries on a single VLAN.

Table 9-6 Address Resolution Protocol Menu Options (/maint/arp)

Command Syntax and Usage
add *<IP address> <MAC address> <VLAN number> <port> <interface>* Adds a permanent ARP entry from switch memory.
delete *<IP address, such as 192.4.17.101>* Removes a single ARP entry from switch memory.
clear Clears the entire ARP list from switch memory.
dump Shows all ARP entries.
addr Shows the list of IP addresses to which the switch will respond for ARP requests.

NOTE – To display all ARP entries currently held in the switch, or a portion according to one of the options listed on the menu above (`find`, `port`, `vlan`, `refpt`, `dump`), you can also refer to “ARP Information” on page 62.

/maint/route

IP Route Manipulation

```
[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
      clear   - Clear route table
      dump    - Show all routes
```

Table 9-7 IP Route Manipulation Menu Options (/maint/route)

Command Syntax and Usage
find *<IP address (such as 192.4.17.101)>* Shows a single route by destination IP address.
gw *<default gateway address, such as 192.4.17.44>* Shows routes to a single gateway.
type indirect|direct|local|broadcast|martian|multicast Shows routes of the selected type.
tag fixed|static|snmp|addr|rip|ospf|bgp|icmp|broadcast|martian|multicast|dynamic Shows routes of the selected tag.
if *<1-1024>* Shows routes on the selected interface.
clear Clears the route table from switch memory.
dump Shows all routes.

NOTE – To display all routes, you can also refer to “IP Routing Information” on page 71.

/maint/panic

Panic Command

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 January 31, 2001...

Boot version 1.0.1

Alteon 708

Rebooted because of console PANIC command.

Booting complete 11:55:01 Thursday January 31, 2001:
```

/maint/tsdump

Tech Support Dump

The tsdump command is used to provide Alteon WebSystems technical support with all available system information. It causes the switch to dump the following information to FLASH memory:

- /info/dump
- /stats/dump
- /cfg/dump

To select `tsdump`, at the `Maintenance#` prompt, enter:

```
Maintenance# tsdump
```

Enter **y** to confirm the command:

```
Confirm dumping all information, statistics, and configuration
[y|n]: y
```

All information about the switch will be displayed. Save the dumped information from your console into a text file, and then email it to Alteon WebSystems Technical Support for evaluation.

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 23, 2000. Use /maint/dump/uuout to
      extract the dump for analysis and /maint/dump/delete to
      clear the FLASH region. The region must be cleared
      before another dump can be taken.
```

CHAPTER 10

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)**

  MP is the processor that handles management of the switch. It processes the CLI, Telnet, SNMP operation, and Spanning Tree.

- **Switch Processor (SP)**

  SP is 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)**

  STP is the IEEE 802.1d-specified loop prevention protocol widely used in Ethernet bridge networks.

- **Bridge Protocol Data Unit (BPDU)**

  BPDU are 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 to which the `ping`ing station is connected 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 used by some other device in the network.

Actions

- Check `/cfg/ip/current` 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/current` to be sure SNMP community strings are correct.
- Check `/cfg/snmp/current` 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, 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 in “auto,” the other device must have the same configuration. Flow control can be configured through CLI for Transmit and Receive, separately or together.
- 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 crossover cable should be used unless there is an uplink enable/disable switch used instead on the switch or hub.

  Table 10-1 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 that in turn, will generate a snap trace, core dump, and switch reset.

Actions

- Messages: Any message(s) on the console should be recorded and sent to Alteon WebSystems Customer Support.
- Core dump: Retrieve the core dump (if available) by accessing the Maintenance menu and invoking the `uudmp` option. Alternately, you can enter **`/maint/uudmp`** to retrieve the coredump. Any core dump should be sent to Alteon WebSystems Customer Support.

Switch Lockup Problems

No response to console input. Cannot `ping` or Telnet. LEDs are lit solid yellow, indicating non connected ports.

Potential Causes

- SP crash
- Switch reset

Things to Check

- Check for any port-to-port traffic.
- Check the status of the LEDs (both yellow and green).
- Check for any response to ping or Telnet to the switch.

Actions

- Verify that you can connect a console to the switch and get to the CLI.
- Verify that the messages stopped printing to the console in mid-stream.

Switch Boot Failure

The switch will not boot.

Possible Causes

- Corrupted firmware
- Firmware and configuration was corrupted when rebooting with an older firmware image

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 `tftp` 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 10-2 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-X> 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 115200 baud
Power cycle to end xmodem.
```

5. **Reconfigure your terminal emulation software for the following parameters:**

Parameter	Value
Baud Rate	115,200
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 because of 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      VLAN   Port   State
00:00:00:00:00:01     1     A3     FWD
```

Spanning Tree Protocol Problems

The topology in the following figure is used to illustrate the STP problems in this section.

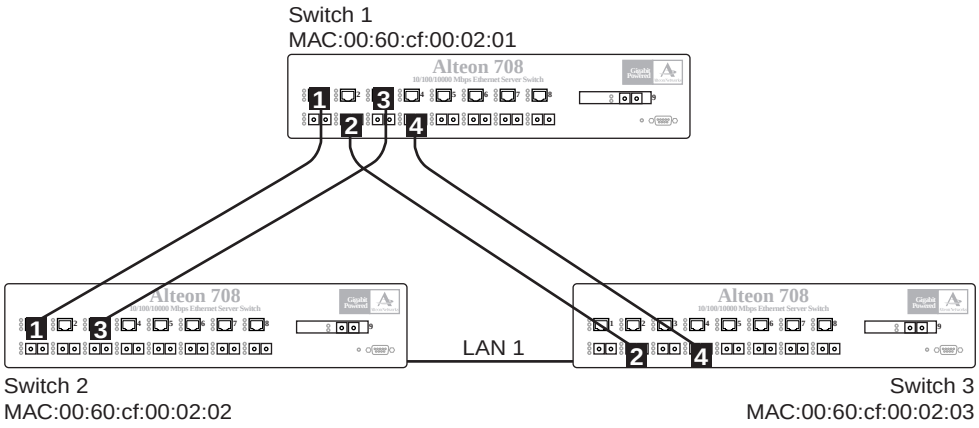


Figure 10-1 Spanning Tree Topology

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 111) and the Server Load Balancing information (page 80) for anything unexpected.
- On the switch, check that the real servers, real server groups, virtual servers, and so forth, have been *enabled*.
- Check that the real servers are physically functioning.
- Check that all the services that 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 143).
- 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 connecting clients and servers to the same switch port.
- Make sure that the port state for each switch port is properly configured as `client`, `server`, or `none` (see “SLB Port Configuration” on page 228).
- Make sure that the switch is configured to accept the TCP/UDP port numbers on which each particular service is expected to run.

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

- The topology or port state is invalid: the real server is connected to the switch through a port configured in the “client” or “application redirection” 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 three-way TCP handshake for opening a TCP connection, followed by a four-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.

APPENDIX A

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 10-3 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

Index

Symbols

/ 43
? (help) 43
[ ] 12

Numerics

802.1d Spanning-Tree Protocol 293

A

abbreviating commands (CLI) 44
ack (SLB filtering option) 225
action (SLB filtering option) 225
activating optional software 260
active configuration block 144, 267
add
    ARP entry 280
    SLB port option 229
    SLB virtual server option 220
addr
    IP route info 73
address list
    ARP entries 280
Address Resolution Protocol (ARP)
    add, delete entries 280
    address list 280
    interval 176
    statistics 113
administrator account 22, 25
administrator privileges 23
aging
    STP bridge option 166
    STP information 60
application health checking 215
application redirection 211, 225
    filter states 80
    filters 209
    within real server groups 214
apply (global command) 144
applying configuration changes 144
ARP. *See* Address Resolution Protocol.
ASCII terminal 16
authentication, application health checking 239
autoconfiguration
    link 31
    link speed 30
auto-negotiation
    enable/disable on port 155, 159
    setup 31

B

backup
    SLB real server group option 215
    SLB real server option 212
backup configuration block 144, 267
banner (system option) 148
baud rate
    console connection 16, 289
    serial download 290
BBI 15
binary firmware image 289
binding failure 106
BLOCKING (port state) 60
Boot Options Menu 263
BOOTP 17
    setup (enable/disable) 29
    system option 148
BPDU. *See* Bridge Protocol Data Unit.
bridge parameter menu, for STP 164
bridge priority 60

Bridge Protocol Data Unit (BPDU) 60, 285
STP transmission frequency 166
Bridge Spanning-Tree parameters 166
broadcast
IP route tag 73
IP route type 72
broadcast IP address 34

C

capture dump information to a file 272
Cisco EtherChannel 171
clear 113
ARP entries 280
dump information 273
FDB entry 276
routing table 281
CLI 15
client slowpath frames 107
client traffic processing 228
Command-Line Interface (CLI) ..... 15 to 23, 25, 41
commands
abbreviations 44
conventions used in this manual 12
shortcuts 44
stacking 44
tab completion 45
configuration
apply changes 144
default gateway interval, for health checks 178
default gateway IP address 178
flow control 155, 159
imask 240
IP broadcast address 177
IP static route 179
operating mode 155, 159
port link speed 155, 159
save changes 144
scriptable health checks 240
switch IP address 177
view changes 143
VLAN IP interface 177
configuration block
active 267
backup 267
factory 267
selection 267
connecting
via console 16
via Telnet 17
console port
communication settings 16, 289
connecting 16
serial download settings 290
contacting Alteon WebSystems 13
coredump 288
cost
STP information 60
STP port option 167
counters
No Server Available (dropped frames) 107
crossover cable 287
cur (system option) 149
current bindings 106
customer support 13

D

DAM sessions created 107
DAM slowpath frames 107
date
setup 28
system option 147
debugging 269
default gateway
information 49, 70
interval, for health checks 178
metrics 208
round robin, load balancing for 208
default password 22
delete
ARP entry 280
deny (filtering) 107
diff (global) command, viewing changes 143
dip (destination IP address for filtering) 295
direct (IP route type) 72
DISABLED (port state) 60
disconnect idle timeout 23
Distributed Site State Protocol (DSSP)
setting update interval 232
dmask
destination mask for filtering 295
Domain Name System (DNS)
health checks 215
peer site handoffs 230
downloading software 264

dropped frames (No Server Available) counter 107
dump
maintenance 269
state information 282
duplex mode
link status 48, 54
dynamic routes 281

E

EtherChannel
as used with port trunking 171

F

factory configuration block 267
factory default configuration 23, 25, 26
failure threshold 216
filter
protocols 223
required parameters 222
filtered (denied) frames 107
filters
IP address ranges 295
firmware image 289
first-time configuration ..... 23, 25 to ??
fixed (IP route tag) 73
flow control 48, 54
configuring 155, 159
forwarding database (FDB) 269
description 285
Forwarding Database Information Menu 64
forwarding state (FWD) 60, 61, 65
FTP server health checks 215
fwd (STP bridge option) 166
FwdDel (forward delay), bridge port 60

G

GSLB
server port 233

H

hash metric 217
health check
script 240
health checking (SLB real server group option) 215
health checks 212
default gateway interval, retries 178
layer information 80
parameters for most protocols 215
redirection (rport) 225
hello
STP information 60
help 43
Hot Standby Router Protocol (HSRP)
priority increment value for L4 client ports 252
use with VRRP 246, 250
VRRP priority increment value 252
HP-OpenView 15
HSRP. *See* Hot Standby Router Protocol.
HTTP
application health checks 215
redirects (Global SLB option) 231

I

ICMP 112
IP route tag 73
Layer 3 health checks 215
ICMP slowpath frames 107
idle timeout
overview 23
IEEE standards
802.1d Spanning-Tree Protocol 59, 164, 293
if 112
ifInDiscards 291
ifInErrors 291
ifOutDiscards 291
image
downloading 264
software, selecting 266
IMAP server health checks 215
imask (IP address mask) 238
incorrect VIPs (statistic) 107
incorrect Vports (dropped frames counter) 107
indirect (IP route type) 72
Information Menu 47
intr
SLB real server option 212

IP address 34
ARP information 62
BOOTP 17
configuring default gateway 178
filter ranges 295
IP interface 34
Telnet 17
IP address mask for SLB 238
IP configuration via setup 34
IP forwarding 205
IP forwarding information 49
IP Forwarding Menu 180
IP Information Menu 49, 70
IP interface
broadcast address (broad) 177
configuring address 177
configuring VLANs 177
subnet address maskconfiguration
IP subnet address 177
IP interfaces 34, 72
information 49, 70
IP route tag 73
priority increment value (ifs) for VRRP 251
IP options 107
IP Port Menu 205
IP Route Manipulation Menu 281
IP routing 34
tag parameters 73
IP Static Route Menu 179
IP subnet mask 34
IP, switch processor statistics for 112

L

L4 disabled 106
Layer 4
administrator account 22
layer 4 user account 22
LEARNING (port state) 60
least connections (SLB Real Server metric) 217
LED patterns 294
licence certificate 260
license password 260
lines (display option) 44
link
speed, configuring 155, 159
troubleshooting 287
link speed 48
auto-sense 30
link status
command 54
duplex mode 48, 54
port speed 54
speed 48
linkt (SNMP option) 169
LISTENING (port state) 60
lmask (routing option) 49
lnet (routing option) 49
local (IP route type) 72
Lookup failures 106

M

MAC (media access control) address .. 48, 51, 62, 64, 260, 276
switch location 17
Main Menu 41
Command-Line Interface (CLI) 23
summary 42
Maintenance Menu 269
Management Processor (MP) 277, 285
display MAC address 48, 51
manual style conventions 12
martian
IP route tag (filtered out) 73
IP route type (filtered out) 72
mask
IP interface subnet address 177
master router 245
MaxAge (STP information) 60
maxcon (maximum connections) 215
mcon (maximum connections) 107
SLB real server option 212
media access control. *See* MAC address.
metrc (SLB real server group option) 214
metrics, SLB 216
minimum misses (SLB real server metric) 216
Miscellaneous Debug Menu 277
mmask
IP address mask for SLB 239
mnet
management traffic IP address for SLB 239
system option 148
MP. *See* Management Processor.
multicast
IP route tag 73
IP route type 72

N

Network Address Translation (NAT)
  filter action 225
network management 15
non TPC/IP frames 107
Not ready 106

O

online help 43
operating mode, configuring 155, 159
Operations Menu 253
Operations-Level Port Options 255
operations-level SLB options 258
operations-level VRRP options 259
optional software 50, 84
  activating 260
  removing 261
overflow servers 212
Overflows 106

P

panic
  switch (and Maintenance Menu option) 269
parameters
  tag 73
  type 72
password
  administrator account 22
  default 22
  L4 administrator account 22
  l4 user account 22
  user account 21
  VRRP authentication 251
passwords 16, 22
ping 43, 211
  troubleshooting 286
poisoned reverse, as used with split horizon 188
POP3
  server health checks 215
port speed 54
port states
  UNK (unknown) 65
port trunking
  description 171
ports
  configuration 30
  information 55
  IP status 49, 70
  membership of the VLAN 49, 58
  priority 60
  SLB state information 80
  STP port priority 167
  VLAN ID 48
preemption
  assuming VRRP master routing authority 245
  virtual router 244
priority (STP port option) 167
proxies
  IP address translation 213
proxy IP address (PIP) 80
PVID (port VLAN ID) 48
pwd 43

Q

quiet (screen display option) 44

R

RADIUS
  server authentication 215, 239
read community string (SNMP option) 169
real server
  menu options 211
  statistics 101
real server groups
  combining servers into 214
  statistics 102
real servers
  backup 215
  priority increment value (reals) for VRRP 252
  SLB state information 80
reboot 269, 282
redir (SLB filtering option) 225
reference ports 65
remote site servers 213
removing optional software 261
reset 277
reset key combination 269
restarting switch setup 27
restr (SLB real server UDP option) 213

retry
health checks for default gateway 178
SLB real server option.....212
rip (IP routing tag).....73
RIP. *See* Routing Information Protocol.
RIP1 information 70
rmkey 261
round robin
as used in gateway load balancing.....208
roundrobin
SLB Real Server metric.....217
route
switch statistics for Route protocol.....113
router redundancy 242
Routing Information Protocol (RIP) 73
options 187
parameters 187
poisoned reverse.....188
split horizon.....188
Routing Information Protocol Menu.....186, 187

S

save (global command) 144
noback option 144
save command 267
serial cable 16
serial download.....289
Server Load Balancing
client traffic processing 228
information.....80
menu options 210
metrics 216
operations-level options.....258
port options.....229
real server group options.....214
real server weights 212
server traffic processing.....228, 229
troubleshooting 293
Server Load Balancing Maintenance Statistics Menu... 111
Server Load Balancing Metrics.....216
server pool.....211
server port mapping.....80
server slowpath frames 107
server traffic processing.....228, 229
Session Binding Table 239
session identifier 219
setup facility 23, 25
BOOTP 29
IP configuration 34
IP subnet mask 34
port auto-negotiation mode..... 31
port configuration 30
restarting..... 27
Spanning-Tree Protocol 29
starting 26
stopping..... 27
system date 28
system time 28
VLAN name 32
VLAN port numbers 33
VLANs..... 32
shortcuts (CLI)..... 44
SIP (source IP address for filtering)..... 295
slowpath sessions created 107
smask
source mask for filtering 295
SMTP server health checks..... 215
snap traces 288
SNMP 15, 112
HP-OpenView..... 15
IP route tag 73
menu options..... 169
set and get access 169
troubleshooting..... 286
software
image 264
image file and version..... 48, 51
license 260
Spanning-Tree Protocol..... 61, 285, 293
bridge aging option 166
bridge parameters..... 166
bridge priority 60
forwarding state 293
port cost option..... 167
port priority option 167
recalculation..... 293
root bridge 60, 166
setup (on/off)..... 29
spanning-tree port state 293
switch reset effect 268
troubleshooting..... 292
split horizon 188
stacking commands (CLI) 44
starting switch setup 26
state (STP information) 60

static (IP route tag) 73
statistics
ARP 113, 125
CPU 135, 141
DNS 127
dump 142
filter 103
GSLB maintenance 111
ICMP 118
IP 115, 136
IP interface 114
letter 130
packet count 134
real server group GSLB 109
real server GSLB 109
route 126
SNMP 122
SP 136
SP filter SLB 105
SP maintenance SLB 106
SP STEM thread 141
SP virtual server SLB 105
STEM memory 131
STEM message 131
STEM thread 133
switch processor SLB 103
TCB allocated blocks 134
TCP 120
UART 135
UDP 122
UDP allocated blocks 134
virtual server GSLB 110
VRRP 128
Statistics Menu 85
stopping switch setup 27
subnet mask 34
subnets 34
IP interface 177
SW allocation failures 107
switch
name and location 48, 51
resetting 268
Switch Processor (SP) 277, 285
display trace buffer 278
swkey 260
synchronization
peer switch 237
VRRP switch 236
system
contact (SNMP option) 169
date and time 48, 51
information 51
location (SNMP option) 169
System Maintenance Menu 274
system options
BOOTP 148
cur (current system parameters) 149
date 147
login banner 148
mnet 148
time 148
wport 148
system parameters, current 149

T

tab completion (CLI) 45
TCP 112
ACK flag 225
fragments 106, 219
health checking using 212
health checks 215
source and destination ports 224
Telnet 17
BOOTP 17
troubleshooting 286
terminal emulation 16
text conventions 12
TFTP 264
time
setup 28
system option 148
timeouts
idle connection 23
time-to-live, DNS response (Global SLB option) .. 231
trace buffer 277
Switch Processor 278
traceroute 43
transparent proxies, when used for NAT 225
troubleshooting ..... 285 to 294
type parameters 72
typographic conventions, manual 12

U

UDP 112
  datagrams 106
  server status using 212
  source and destination ports 224
unknown (UNK) port state 65
Unscheduled System Dump 283
upgrade, switch software 264
user access
  levels 21
user account 21
Uuencode Flash Dump 272

V

verbose 44
virtual port state, SLB information about 80
Virtual Router Redundancy Protocol (VRRP)
  authentication parameters for IP interfaces 250
  configuration menu options 242
  operations-level options 259
  password, authentication 251
  priority tracking options 246
  virtual router options 244
virtual routers
  description 243
  HSRP failover 246, 250
  HSRP priority increment value 252
  increasing priority level of 245
  master preemption (prio) 244
  priority increment values (vrs) for VRRP 251, 252
virtual servers 216
  SLB state information 80
  statistics 102
virtual sever IP address (VIP) 80
VLAN tagging
  port restrictions 162
VLANs 34
  ARP entry information 62
  information 58
  interface 35
  name 49, 58
  name setup 32
  number 49
  port membership 49, 58
  port numbers 33
  setup 32
  tagging 162
VRID (virtual router ID) 243

W

watchdog timer 269, 288
wcomm 169
web-based management interface 15
weights
  for SLB real servers 212, 217
  setting virtual router priority values 251
wport 148
write community string (SNMP option) 169

X

Xmodem 290