Web OS Switch Software



10.0 Command Reference

Part Number: 212778, Revision A, February 2002

NORTEL NETWORKS

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

Alteon WebSystems
Intelligent Webworking™

Copyright 2002 Nortel Networks, Inc., 50 Great Oaks Boulevard, San Jose, California 95119, USA. All rights reserved. Part Number: 212778, 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 Nortel Networks, 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).

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

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

Contents

Preface 11
Who Should Use This Book 11
How This Book Is Organized 11
Typographic Conventions 13
Contacting Alteon WebSystems 14

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 17

Accessing the Switch 19
CLI Versus Setup 21
Command Line History and Editing 21
Idle Timeout 21

Chapter 2: First-Time Configuration 23
Using the Setup Utility 23
- Information Needed For Setup 23
- Starting Setup When You Log In 24
- Stopping and Restarting Setup Manually 25
- Setup Part 1: Basic System Configuration 25
- Setup Part 2: Port Configuration 27
- Setup Part 3: VLANs 30
- Setup Part 4: IP Configuration 32
- Setup Part 5: Final Steps 35
- Optional Setup for SNMP Support 36
- Optional Setup for Telnet Support 36

Setting Passwords 37
- Changing the Default Administrator Password 37
- Changing the Default User Password 39
- Changing the Default Layer 4 Administrator Password 40

Chapter 3: Menu Basics 43
The Main Menu 43
Menu Summary 44
Global Commands 45
Command Line History and Editing 47
Command Line Interface Shortcuts 48
Command Stacking 48
Command Abbreviation 48
Tab Completion 48

Chapter 4: The Information Menu 49
Information Menu 49
SLB Information 53
Show Session Table Information 55
Show All Layer 4 Information 56
IP Routing Information 57
Show All IP Route Information 58
ARP Information 60
Show All ARP Entry Information 61
ARP Address List Information 61
OSPF Information 62
OSPF General Information 63
OSPF Interface Information 63
OSPF Database Information 64
OSPF Information Route Codes 66
FDB Information 67
Show All FDB Information 68
Clearing Entries from the Forwarding Database 68
System Information 69
Show Last 30 Syslog Messages 70
Link Status Information 71
Spanning Tree Information 72
VLAN Information 74
Port Information 75
IP Information 76
VRRP Information 77
Trunk Group Information 78
Bandwidth Management Information 78
Software Enabled Keys 79
Information Dump 79

Chapter 5: The Statistics Menu 81
Statistics Menu 81
Port Statistics Menu 84
Bridging Statistics 85
Ethernet Statistics 86
Interface Statistics 89
Interface Protocol Statistics 91
Link Statistics 92
RMON Statistics 93
CPU Statistics 96
Maintenance Statistics 97
Load Balancing Statistics Menu 100
Port SLB Statistics Menu 103
Port Real Server SLB Statistics 104
Port Real Server Group SLB Statistics 104
Port Virtual Server SLB Statistics 105
Port Filter SLB Statistics 106
Port Maintenance SLB Statistics 106
Real Server SLB Statistics 108
Per Service Octet Counters 108
Real Server Group Statistics 109
Virtual Server SLB Statistics 110
Filter SLB Statistics 110
Global SLB Statistics Menu 111
Real Server Global SLB Statistics 111
Real Server Group Global SLB Statistics 112
Virtual Server Global SLB Statistics 112
Global SLB Maintenance Statistics 113
SLB Layer7 Statistics Menu 114
Layer7 Redirection Statistics 114
Layer 7 SLB String Statistics 115
Layer 7 SLB Maintenance Statistics 116
SLB Secure Socket Layer Statistics 118
File Transfer Protocol SLB and Filter Statistics Menu 119
Active FTP SLB Parsing and Filter Statistics 120
Passive FTP SLB Parsing Statistics 120
FTP SLB Maintenance Statistics 121
FTP SLB Statistics Dump 121
RTSP SLB Statistics 122
DNS SLB Statistics 123
WAP SLB Statistics 124

SLB TCP Rate Limiting Statistics 126
SLB Maintenance Statistics 127
Clearing the SLB Statistics 128
Bandwidth Management Statistics Menu 130
Bandwidth Management Switch Processor Statistics 131
Bandwidth Management Contract Statistics 132
Bandwidth Management Contract Rate Statistics 133
Bandwidth Management History Statistics 134
Management Processor Statistics 135
STEM Memory Statistics 136
All STEM Memory Statistics 137
DMA Statistics 138
Packet Statistics 139
TCP Statistics 140
UCB Statistics 140
UART Statistics 141
CPU Statistics 142
Interface Statistics 143
IP Statistics 145
ICMP Statistics 148
TCP Statistics 150
UDP Statistics 151
SNMP Statistics 152
FDB Statistics 155
Route Statistics 156
ARP Statistics 157
DNS Statistics 158
OSPF Statistics Menu 158
OSPF Global Statistics 159
VRRP Statistics 163
Statistics Dump 163

Chapter 6: The Configuration Menu 165

Configuration Menu 165
Viewing, Applying, and Saving Changes 167
- Viewing Pending Changes 167
- Applying Pending Changes 167
- Saving the Configuration 168

System Configuration 169
- System Host Log Configuration 171
- SSH Server Configuration 172
- RADIUS Server Configuration 173
- NTP Server Configuration 174
- User Access Control Configuration 175

Port Configuration 177
- Port Link Configuration 179
- Temporarily Disabling a Port 180

IP Configuration 181
- IP Interface Configuration 183
- Default IP Gateway Configuration 184
  - Default Gateway Metrics 185
- IP Static Route Configuration 185
- IP Forwarding Configuration 186
  - Local Network Route Caching Definition 186
- Defining IP Address Ranges for the Local Route Cache 187
- Routing Information Protocol Configuration 188
- Open Shortest Path First Configuration 189
  - Area Index Configuration Menu 190
  - OSPF Summary Range Configuration Menu 192
  - OSPF Interface Configuration Menu 193
  - OSPF Virtual Link Configuration Menu 195
  - OSPF Host Entry Configuration Menu 196
- Border Gateway Protocol Configuration 197
  - BGP Peer Configuration 198
  - BGP Filter Configuration 200
- IP Port Configuration 201
- Domain Name System Configuration 201
- Bootstrap Protocol Relay Configuration 202
- Default Gateway Metrics 203

VLAN Configuration 204
Spanning Tree Configuration 206
- Bridge Spanning Tree Configuration 207
- Spanning Tree Port Configuration 209

SNMP Configuration 210
Port Mirroring Menu 212
Port-Mirroring Menu 213
Trunk Configuration 214
VRRP Configuration 215
Virtual Router Configuration 216
Virtual Router Priority Tracking Configuration 219
Virtual Router Group Configuration 221
Virtual Router Group Priority Tracking Configuration 223
VRRP Interface Configuration 225
VRRP Tracking Configuration 226
Bandwidth Management Configuration 228
Bandwidth Management Contract Configuration 229
Bandwidth Management Policy Configuration 230
Bandwidth Management Current Configuration 231
Setup 232
Dump 232
Saving the Active Switch Configuration 233
Restoring the Active Switch Configuration 233

Chapter 7: The SLB Configuration Menu 235

SLB Configuration 235
Filtering and Layer 4 (Server Load Balancing) 237
Real Server SLB Configuration 237
Real Server Layer 7 Configuration 241
Real Server Group SLB Configuration 242
Server Load Balancing Metrics 245
Virtual Server SLB Configuration 246
Virtual Server Service Configuration 248
Virtual Server RTSP Configuration 252
Cookie-Based Persistence 253
SLB Filter Configuration 254
Defining IP Address Ranges for Filters 258
Advanced Filter Configuration 259
Advanced Filter TCP Configuration 261
IP Advanced Menu 263
ICMP Message Types 264
Advanced Layer 7 Deny Filter 265
Port SLB Configuration 266
Global SLB Configuration 268
GSLB Remote Site Configuration 271

GSLB Lookup Configuration 272
GSLB Internet Network Preference Lookup Configuration 273
Layer 7 SLB Resource Definition Menu 274
Web Cache Redirection Configuration 274
RTSP Cache Redirection Configuration Menu 276
Server Load Balance Resource Configuration Menu 276
WAP Configuration 277
Synchronize Peer Switch Configuration 278
Peer Switch Configuration 279
Advanced Layer 4 Configuration 280
Scriptable Health Checks Configuration 283
WAP Health Check Configuration 284
SYN Attack Detection Configuration 285

Chapter 8: The Operations Menu 287

Operations Menu 287
Operations-Level Port Options 289
Operations-Level SLB Options 290
Operations-Level VRRP Options. 291
Operations-Level Bandwidth Management Options 291
Operations-Level IP Options 292
Operations-Level BGP Options 292
Activating Optional Software 293
Removing Optional Software 294

Chapter 9: The Boot Options Menu 295

Boot Menu 295
Updating the Switch Software Image 296
Downloading New Software to Your Web Switch 296
Selecting a Software Image to Run 297
Uploading a Software Image from Your Web Switch 298
Selecting a Configuration Block 299
Resetting the Web Switch 299

Chapter 10: The Maintenance Menu 301
Maintenance Menu 301
System Maintenance Options 302
Forwarding Database Options 303
Debugging Options 304
ARP Cache Options 305
IP Route Manipulation 306
Uuencode Flash Dump 307
TFTP System Dump Put 308
Clearing Dump Information 308
Panic Command 309
Unscheduled System Dumps 309

Appendix A: Web OS Syslog Messages 311
LOG_WARNING 311
LOG_ALERT 312
LOG_CRIT 313
LOG_ERR 313
LOG_NOTICE 319
LOG_INFO 321

Appendix B: Web OS SNMP Agent 323

Appendix C: Performing a Serial Download 327

Glossary 329

Index 333

Preface

The *Web OS 10.0 Command Reference* describes how to configure and use the Web OS software with the Alteon family of Web switches.

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

Who Should Use This Book

This *Command Reference* is intended for network installers and system administrators engaged in configuring and maintaining a network. The administrator 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, Jumbo Frames, Spanning Tree Protocol, SNMP, Port Mirroring, IP Routing, Port Trunking, and more.

Chapter 7 "The SLB Configuration Menu," shows how to configure Server Load Balancing, Filtering, Global Server Load Balancing, and more.

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

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

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

Appendix A, "Web OS Syslog Messages," shows a listing of syslog messages.

Appendix B, "Web OS SNMP Agent," lists the Management Interface Bases (MIBs) supported in the switch software.

Appendix C, "Performing a Serial Download," shows how to directly load a binary software image into the switch for upgrade or maintenance.

"Glossary" includes definitions of terminology used throughout the book.

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]`
<Key>	Non-alphanumeric keyboard items are shown in regular type inside brackets. When directed, press the appropriate key. Do not type the brackets.	Press the <Enter> key.

Contacting Alteon WebSystems

For complete product support and sales information, visit the Nortel Networks website at the following URL:

http://www.nortelnetworks.com

See the contact information on this site for regional support and sales phone numbers and e-mail addresses.

- In North America, call toll-free 1-800-4NORTEL.
- Outside North America, call 987-288-3700.

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

- A built-in, text-based command line interface and menu system for access via local terminal or remote Telnet session
- A Web-based management interface for interactive network access through your Web browser
- SNMP support for access through network management software such as HP-OpenView

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

This chapter explains how to access the Command Line Interface (CLI) to the switch.

Connecting to the Switch

You can access the command line interface in two ways:

- Using a console connection via the console port
- Using a Telnet connection over the network
- Using a SSH connection to securely log into another computer over a 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 guide for specifics).

Procedure

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

   You will next be required to enter a password for access to the switch. (For more information, see “Setting Passwords” on page 37).

Establishing a Telnet Connection

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

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

- Dynamically, from a BOOTP server on your network
- Manually, when you configure the switch IP address (see "Setup Part 1: Basic System Configuration" on page 25).

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 69).

Running Telnet

Once the IP parameters on the Web 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 then be prompted to enter a password as explained on page 17.

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 SSH (Secure Shell) 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.

The switch can do only one session of key/cipher generation at a time. Thus, a SSH/SCP client will not be able to login 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.

The supported SSH encryption and authentication methods 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

The following SSH clients have been tested:

- SSH 1.2.23 and SSH 1.2.27 for Linux (freeware)
- SecureCRT 3.0.2 and SecureCRT 3.0.3 (Van Dyke Technologies, Inc.)
- F-Secure SSH 1.1 for Windows (Data Fellows)

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.

Running SSH

Once the IP parameters are configured and the SSH service is turned on the Web switch, you can access the command line interface using an SSH connection.

To establish an SSH connection with the switch, run the SSH program on your workstation by issuing the SSH command, followed by the switch IP address:

```
>> # ssh <switch IP address>
```

or, if SecurID authentication is required, use the following command:

```
>> # ssh -1 ace <switch IP address>
```

You will then be prompted to enter your user name and password.

Accessing the Switch

To enable better switch management and user accountability, seven levels or *classes* of user access have been implemented on the Web switch. Levels of access to CLI, Web management functions, and screens increase as needed to perform various switch management tasks. Conceptually, access classes are defined as follows:

- User interaction with the switch is completely passive—nothing can be changed on the Web 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 Web 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.
- 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 Web switch. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

Access to switch functions is controlled through the use of unique surnames and passwords. Once you are connected to the switch via local console, Telnet, or SSH, you are prompted to enter a password. The default user names/password for each access level are listed in the following table.

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 37.

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 all switch status information and statistics, but cannot make any configuration changes to the switch.	`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 Server Load Balancing operation menu.	`slboper`

Table 1-2 User Access Levels

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. This user currently has the same access level as the SLB operator. and the access level is reserved for future use, to provide access to operational commands for operators managing traffic on the line leading to the shared Internet services.	`l4oper`
Operator	The Operator manages all functions of the switch. In addition to SLB Operator functions, the Operator can reset ports or the entire switch.	`oper`
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 Server Load Balancing menus, with the exception of not being able to configure filters or bandwidth management.	`slbadmin`
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 Administrator can configure all parameters on the Server Load Balancing menus, including filters and bandwidth management.	`l4admin`
Administrator	The superuser Administrator has complete access to all menus, information, and configuration commands on the Web switch, including the ability to change both the user and administrator passwords.	`admin`

NOTE – With the exception of the “admin” user, access to each user level can be disabled by setting the password to an empty value. All user levels below “admin” will (by default) 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.

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.

The following table shows the Main Menu with administrator privileges.

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

NOTE – If you are accessing a user account or Layer 4 administrator account, some menu options will not be available.

Command Line History and Editing

For a description of global commands, shortcuts, and command line editing functions, see "Menu Basics" on page 43."

Idle Timeout

By default, the switch will disconnect your console or Telnet session after five minutes of inactivity. This function is controlled by the idle timeout parameter, which can be set from 1 to 60 minutes. For information on changing this parameter, see "System Configuration" on page 169.

CHAPTER 2

First-Time Configuration

To help with the initial process of configuring your switch, the Web OS 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 any time after login.

Information Needed For Setup

Setup requests the following information:

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

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

Starting Setup When You Log In

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

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

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

```
Enter Password:
```

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

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

```
Connected to Alteon ACEdirector 4
18:44:05 Wed Jan 3, 2001

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

NOTE – If the default `admin` login is unsuccessful, or if the administrator Main Menu appears instead, the system configuration has probably been changed from the factory default settings. If you are certain that you need to return the switch to its factory default settings, see "Selecting a Configuration Block" on page 299.

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

Stopping and Restarting Setup Manually

Stopping Setup

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

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

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

Restarting Setup

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

```
# /cfg/setup
```

Setup Part 1: Basic System Configuration

When Setup is started, the system prompts:

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

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

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

   If you decide not to configure VLANs during this session, you can configure them later using the configuration menus, or by restarting the Setup facility. For more information on configuring VLANs, see the *Web OS 10.0 Application Guide*.

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

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

   ```
   System Date:
   Enter month [1]:
   ```

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

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

```
Enter day [3]:
```

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

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

```
Enter year [2001]:
```

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

The system displays the date and time settings:

```
System clock set to 18:55:36 Wed Jan 3, 2001.
```

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

```
System Time:
Enter hour in 24-hour format [18]:
```

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

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

```
Enter minutes [55]:
```

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

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

```
Enter seconds [37]:
```

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 8:55:36 Wed Jan 3, 2001.
```

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

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

If available on your network, a BOOTP server can supply the switch with IP parameters so that you do not have to enter them manually. 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 at the prompt:**

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

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

Setup Part 2: Port Configuration

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

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

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

If you wish to change settings for individual ports, enter the number of the port you wish to configure. To skip port configuration, press <Enter> without specifying any port and go to "Setup Part 3: VLANs" on page 30.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

7. **If appropriate, configure Gigabit Ethernet port autonegotiation mode.**

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

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

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

8. **If configuring VLANs, enable or disable VLAN tagging for the port.**

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

```
Port VLAN tagging config (tagged port can be a member of multiple VLANs)
Current TAG flag:               disabled
Enter new TAG status [d/e]:
```

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

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

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

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

Setup Part 3: VLANs

If you chose to skip VLANs configuration back in Part 1, skip to "Setup Part 4: IP Configuration" on page 32.

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

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

If you wish to change settings for individual VLANs, enter the number of the VLAN you wish to configure. To skip VLAN configuration, press <Enter> without typing a VLAN number and go to "Setup Part 4: IP Configuration" on page 32.

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

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

Entering a new VLAN name is optional. To use the pending new VLAN name, press <Enter>.

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

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

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

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

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

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

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

```
>
```

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

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

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

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

Setup Part 4: IP Configuration

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

IP Interfaces

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

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

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

```
IP Config:

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

If you wish to configure individual IP interfaces, enter the number of the IP interface you with to configure. To skip IP interface configuration, press <Enter> without typing an interface number and go to "Default Gateways" on page 33.

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

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

To keep the current setting, press <Enter>.

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

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

To keep the current setting, press <Enter>.

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

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

To keep the current setting, press <Enter>.

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

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

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

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

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

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

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

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

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

Default Gateways

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

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

Enter the number for the default gateway to be configured. To skip default gateway configuration, press <Enter> without typing a gateway number and go to "IP Routing" on page 34.

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

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

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

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

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

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

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

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

IP Routing

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

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

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

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

Enter **y** to enable IP forwarding. To disable IP forwarding, enter **n** and proceed to Step 2.To keep the current setting, press <Enter>.

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

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

Setup Part 5: Final Steps

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NOTE – After initial configuration is complete, it is recommended that you change the default passwords as shown in "Setting Passwords" on page 37.

Optional Setup for SNMP Support

NOTE – This step is optional. Perform this procedure only if you are planning on using SNMP-based tools, such as Alteon EMS.

1. **Enable SNMP and select one of the options.**

```
>> # /cfg/sys/snmp dis|read|write
```

2. **Set SNMP read or write community string. By default, they are public and private respectively.**

```
>> # /cfg/snmp/rcomm|wcomm
```

3. **Apply and save configuration if you are not configuring the switch with Telnet support. Otherwise apply and save after "Optional Setup for Telnet Support" on page 36.**

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

Optional Setup for Telnet Support

NOTE – This step is optional. Perform this procedure only if you are planning on connecting to the switch through any telnet application.

1. **Enable telnet.**

```
>> # /cfg/sys/tnet ena
```

2. **Apply and save SNMP and /or telnet configuration(s).**

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

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

Setting Passwords

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

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

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

Changing the Default Administrator Password

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

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

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

```
Main# /cfg
```

The Configuration Menu is displayed.

```
[Configuration Menu]
      sys     - System-wide Parameter Menu
      port    - Port Menu
      ip      - IP Menu
      vlan    - VLAN Menu
      stp     - Spanning Tree Menu
      snmp    - SNMP Menu
      mirr    - Port Mirroring Menu
      slb     - Server Load Balancing Menu
      trunk   - Trunk Group Menu
      vrrp    - Virtual Router Redundancy Protocol Menu
      bwm     - Bandwidth Management Menu
      setup   - Step by step configuration set up
      dump    - Dump current configuration to script file
      ptcfg   - Backup current configuration to tftp server
      gtcfg   - 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]
     syslog  - Syslog Menu
     sshd    - SSH Server Menu
     radius  - RADIUS Authentication Menu
     ntp     - NTP Server Menu
     date    - Set system date
     time    - Set system time
     idle    - Set timeout for idle CLI sessions
     snmp    - Set SNMP access control
     wport   - Set Web server port number
     notice  - Set login notice
     bannr   - Set login banner
     mnet    - Set management network
     mmask   - Set management netmask
     smtp    - Set SMTP host
     tnport  - Set Telnet server port number
     hprompt - Enable/disable display hostname (sysName) in CLI prompt
     bootp   - Enable/disable use of BOOTP
     http    - Enable/disable HTTP (Web) access
     user    - User Access Control Menu (passwords)
     cur     - Display current system-wide parameters
```

NOTE – The Browser Based Interface (BBI) command (`wport`) is not supported on the Alteon AD3 or 180e Web switches.

4. **Select the administrator password.**

```
System# user/admpw
```

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

```
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 fix-up mode.

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

```
Enter new administrator password:
```

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

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

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

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

Changing the Default User Password

The user login has limited control of the switch. Through a user account, you can view switch information and statistics, but you can't make configuration changes.

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

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

```
Main# cfg
```

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

```
>> Configuration# sys
```

4. **Select the user password.**

```
System# user/usrpw
```

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

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

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

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

```
Enter new user password:
```

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

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

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

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

Changing the Default Layer 4 Administrator Password

The Layer 4 administrator has limited control of the switch. Through a Layer 4 administrator account, you can view all switch information and statistics, but can configure changes only on the Server Load Balancing menus.

The default password for the Layer 4 administrator account is `l4admin`. To change the default password, follow this procedure:

1. **Connect to the switch and log in using the administrator account.**

   To change any switch password, you must login using the administrator password. Passwords cannot be modified from the Layer 4 administrator account or the user account.

2. **From the Main Menu, use the following command to access the System Menu:**

```
Main# /cfg/sys/user
```

3. **Select the Layer 4 administrator password:**

```
System# l4apw
```

4. **Enter the current *administrator* password (not the Layer 4 administrator password) at the prompt:**

```
Changing L4 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 fix-up mode.

5. **Enter the new Layer 4 administrator password at the prompt:**

```
Enter new L4 administrator password:
```

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

```
Re-enter new L4 administrator password:
```

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

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

CHAPTER 3

Menu Basics

The Web switch's Command Line Interface (CLI) is used for viewing switch information and statistics. In addition, the administrator can use the CLI for performing all levels of switch configuration.

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

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

The Main Menu

The Main Menu appears after a successful connection and login. The following table shows the Main Menu for the administrator login. Some features are not available under the user login.

```
[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**

  Provides sub-menus for displaying information about the current status of the switch: from basic system settings to VLANs, Layer 4 settings, and more.

- **Statistics Menu**

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

- **Configuration Menu**

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

- **Operations Command Menu**

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

- **Boot Options Menu**

  This menu is used for upgrading switch software, selecting configuration blocks, and for resetting the switch when necessary.

- **Maintenance Menu**

  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*	Provides more information about a specific command on the current menu. When used without the *command* parameter, a summary of the global commands is displayed.
.	Display the current menu.
..	Go up one level in the menu structure.
/	If placed at the beginning of a command, go to the Main Menu. Otherwise, this is used to separate multiple commands placed on the same line.
diff	Show any pending configuration changes.
apply	Apply pending configuration changes.
save	Write configuration changes to non-volatile flash memory.
revert	Remove pending configuration changes between “apply” commands. Use this command to restore configuration parameters set since last “apply” command.
exit	Exit from the command line interface and log out.
ping	Use this command to verify station-to-station connectivity across the network. The format is as follows: **ping** *address* [*tries* [*delay*]] Where *address* is the hostname or IP address of the device, *tries* (optional) is the number of attempts (1-32), and *delay* (optional) is the number of milliseconds between attempts. The DNS parameters must be configured if specifying hostnames (see “Domain Name System Configuration” on page 201).
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 hostname or IP address of the target station, *max-hops* (optional) is the maximum distance to trace (1-16 devices), and *delay* (optional) is the number of milliseconds for wait for the response. As with `ping`, the DNS parameters must be configured if specifying hostnames.
pwd	Display the command path used to reach the current menu.

Table 3-1 Global Commands

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:

Table 3-2 Command Line History and Editing Options

Option	Description
`history`	Display a numbered list of the last 10 previously entered commands.
`!!`	Repeat the last entered command.
`!`*n*	Repeat the *n*th command shown on the history list.
<Ctrl-p>	(Also the up arrow key.) Recall the *previous* command from the history list. This can be used multiple times to work backward through the last 10 commands. The recalled command can be entered as is, or edited using the options below.
<Ctrl-n>	(Also the down arrow key.) Recall the *next* command from the history list. This can be used multiple times to work forward through the last 10 commands. The recalled command can be entered as is, or edited using the options below.
<Ctrl-a>	Move the cursor to the beginning of command line.
<Ctrl-e>	Move cursor to the *end* of the command line.
<Ctrl-b>	(Also the left arrow key.) Move the cursor *back* one position to the left.
<Ctrl-f>	(Also the right arrow key.) Move the cursor *forward* one position to the right.
<Backspace>	(Also the Delete key.) Erase one character to the left of the cursor position.
<Ctrl-d>	*Delete* one character at the cursor position.
<Ctrl-k>	*Kill* (erase) all characters from the cursor position to the end of the command line.
<Ctrl-l>	Redraw the screen.
<Ctrl-u>	Clear the entire line.
Other keys	Insert new characters at the cursor position.

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 sub-menu. 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, waiting to be entered. 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

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

/info
Information Menu

```
[Information Menu]
      slb     - Layer 4 Information Menu
      route   - IP Routing Information Menu
      arp     - ARP Information Menu
      ospf    - OSPF routing information menu
      fdb     - Forwarding Database Information Menu
      sys     - Show system information
      log     - Show last 30 syslog messages
      link    - Show link status
      stp     - Show STP information
      vlan    - Show VLAN information
      port    - Show port information
      ip      - Show IP information
      vrrp    - Show Virtual Router Redundancy Protocol information
      trunk   - Show Trunk Group information
      bwm     - Show Bandwidth Management information
      swkey   - Show enabled software features
      clrlog  - Clear syslog messages
      dump    - Dump all information
```

The information provided by each menu option is briefly described in Table 4-1 on page 50, with pointers to where detailed information can be found.

Table 4-1 Information Menu Options (/info)

Command Syntax and Usage
slb Displays the Layer 4 Information Menu. For details, see page 53.
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 RIP tagged routes, specifying the number of hops to the destination (1-15 hops, or 16 for infinite hops) ■ The IP interface that the route uses For details, see page 57.
arp Displays the Address Resolution Protocol (ARP) Information Menu. For details, see page 60.
ospf Displays OSPF routing information menu. For details, see page 62.
fdb Displays the Forwarding Database Information Menu. For details, see page 67.
sys Displays system information, including: ■ System date and time ■ Switch model name and number ■ Switch name and location ■ Time of last boot ■ MAC address of the switch management processor ■ IP address of IP interface #1 ■ Hardware version and part number ■ Software image file and version number ■ Configuration name ■ Log-in banner, if one is configured For details, see page 69.
log Displays 30 most recent syslog messages. For details, see page 70.

Table 4-1 Information Menu Options (/info)

Command Syntax and Usage
link Displays configuration information about each port, including: ■ Port number ■ Port speed (10, 100, 10/100, or 1000) ■ Duplex mode (half, full, or auto) ■ Flow control for transmit and receive (no, yes, or auto) ■ Link status (up or down) For details, see page 71.
stp In addition to seeing if STP is enabled or disabled, you can view the following STP bridge information: ■ Priority ■ Hello interval ■ Maximum age value ■ Forwarding delay ■ Aging time You can also see the following port-specific STP information: ■ Port number and priority ■ Cost ■ State For details, see page 72.
vlan Displays VLAN configuration information, including: ■ VLAN Number ■ VLAN Name ■ Status ■ Jumbo Frame usage ■ Port membership of the VLAN For details, see page 74.
port Displays port status information, including: ■ Port number ■ Whether the port uses VLAN Tagging or not ■ Port VLAN ID (PVID) ■ Port name ■ VLAN membership For details, see page 75.

Table 4-1 Information Menu Options (/info)

Command Syntax and Usage
ip Displays IP Information. For details, see page 76. IP information, includes: ■ IP interface information: Interface number, IP address, subnet mask, broadcast address, VLAN number, and operational status. ■ Default gateway information: Metric for selecting which configured gateway to use, gateway number, IP address, and health status ■ IP forwarding information: Enable status, `lnet` and `lmask` ■ Port status
vrrp Displays the VRRP Information Menu. For details, see page 77.
trunk When trunk groups are configured, you can view the state of each port in the various trunk groups. For details, see page 78.
bwm Shows bandwidth management information. For details, see page 78.
swkey Displays a list of all the optional software packages which have been activated or installed on your switch. For details see page 79
clrlog Clears syslog messages displayed with `/info/log`.
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.

/info/slb
SLB Information

```
[Server Load Balancing Information Menu]
      sess    - Session Table Information Menu
      real    - Show real server information
      virt    - Show virtual server information
      filt    - Show filter information
      port    - Show port information
      gslb    - Show GSLB information
      idshash - Show IDS server selected by hash or minmisses metric
      thash   - Show real server selected by thash metric
      cookie  - Decode the HEX value to get the VIP and RIP
      synatk  - Show SYN attack detection information
      dump    - Show all layer 4 information
```

Layer 4 information includes the following:

Table 4-2 Layer 4 Information Menu Options (/info/slb)

Command Syntax and Usage
sess Displays the Session Table Information Menu. To view menu options, see page 55.
real *<real server number (1-255)>* 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 (1-256)>* ■ Displays Virtual Server State: Virtual server number, IP address, virtual MAC address ■ Virtual Port State: Virtual service or port, server port mapping, real server group, group backup server.
filt *<filter ID (1-224)>* Displays the filter number, destination port, real server port, real server group, health check layer, group backup server, URL for health checks, and real server group, IP address, backup server, and status.
port *<port number (1-9)>* 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 the remote switch number, IP address, IP subnet mask, and health status.

Table 4-2 Layer 4 Information Menu Options (/info/slb)

Command Syntax and Usage
idshash Displays the Intrusion Detection System server selected by `hash` or `minmisses` metric.
thash Displays real server selected by `thash` metric.
cookie Decodes the hexadecimal value to get the virtual server IP address and real server IP address.
synatk Displays SYN attack detection information. To identify whether or not the server is under SYN attack, the number of new half open sessions is examined within a set period of time, for example, every two seconds. This feature requires `dbind` to be enabled.
dump Displays all Layer 4 information for the switch. For details, see page 56.

/info/slb/sess

Show Session Table Information

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

Table 4-3 Session Information Menu Options (/info/slb/sess)

Command Syntax and Usage
find *<IP address>* Displays all session entries with source IP address.
port *<port number (1-9)>* Displays all session entries on port.

dump
Displays all session entries. Information similar to the following may appear in a session entry dump:
`4,5:  3.3.3.8 4579, 0.0.0.6 80 -> 6.6.6.5 80  age 10  E`
`4,16:  3.3.3.8 2260, 2.2.2.11 80 -> 6.6.6.5 80  age 10  E`
`9,3:  61.5.193.97 1, 0.0.0.0 1  ALLOW  age 8  EPS`
`4,44:  3.3.3.8 4885, 3.3.3.11 80 -> 0.0.0.0 80  age 0  U`
`4,46:  3.3.3.8 4886, 3.3.3.11 80 -> 0.0.0.0 80  age 0  U`
The codes E, P, and U at the end of the session entries stand for:
E = established
P = persistent
U = pass UP for L7 processing
S = special session, such as passive cookie or SSL

/info/slb/dump

Show All Layer 4 Information

```
Real server state:
  1: 210.1.2.200, 00:01:02:c1:4b:48, vlan 1, port 1, health 3, up
  2: 210.1.2.1, 00:01:02:70:4d:4a, vlan 1, port 8, health 3, up
 26: 20.20.20.102, 00:03:47:07:a4:9e, vlan 1, port 6, health 3, up
 27: 20.20.20.101, 00:01:02:71:9c:a6, vlan 1, port 7, health 3, up

Virtual server state:
  1: 20.20.20.200,    00:60:cf:47:5c:1e
    virtual ports:
    http: rport http, group 88, backup none, dbind
      HTTP Application: urlslb
        real servers:
         26: 20.20.20.102, backup none, 2 ms, up
             exclusionary string matching: disabled
             1: any
             2: urlone
         27: 20.20.20.101, backup none, 1 ms, up
             exclusionary string matching: disabled
             3: urltwo
             4: urlthree

Redirect filter state:
200: group 1, health 3, backup none
     proxy enabled, radius snoop disabled
    real servers:
      1: 210.1.2.200, backup none, 3 ms, up
      2: 210.1.2.1, backup none, 2 ms, up

Port state:
  1: 0.0.0.0
     filt disabled, filters: 80
  2: 0.0.0.0, idslb
     filt  enabled, filters: 200
  3: 0.0.0.0, idslb
     filt  enabled, filters: 200
  4: 0.0.0.0
     filt disabled, filters: 50 200
```

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

Command Syntax and Usage
find *<IP address>* Displays a single route by destination IP address.
gw *<default gateway address>* Displays routes to a single gateway.
type indirect|direct|local|broadcast|martian|multicast Displays routes of a single type. For a description of IP routing types, see Table 4-5 on page 58.
tag fixed|static|snmp|addr|rip|icmp|broadcast|martian|multicast|vip|bgp Displays routes of a single tag. For a description of IP routing types, see Table 4-6 on page 59.
if *<interface number (1-256)>* Displays routes on a single interface.
dump Displays all routes configured in the switch. For more information, see page 58.

/info/route/dump

Show All IP Route Information

```
  Destination         Mask          Gateway        Type      Tag     Metr If
--------------- --------------- --------------- --------- --------- ---- 
0.0.0.0         0.0.0.0         172.19.1.1      indirect  rip          2  1
0.0.0.0         0.0.0.0         172.19.1.1      indirect  static          1
127.0.0.0       255.0.0.0       0.0.0.0         martian   martian
172.17.0.0      255.255.0.0     172.19.1.1      indirect  rip          2  1
172.19.1.0      255.255.255.0   172.19.1.201    direct    fixed           1
172.19.1.201    255.255.255.255 172.19.1.201    local     addr            1
172.19.1.255    255.255.255.255 172.19.1.255    broadcast broadcast       1
172.20.0.0      255.255.0.0     172.19.1.1      indirect  rip          2  1
172.23.0.0      255.255.0.0     172.19.1.1      indirect  rip          3  1
172.25.0.0      255.255.0.0     172.19.1.1      indirect  rip          4  1
172.26.0.0      255.255.0.0     172.19.1.1      indirect  rip          3  1
172.27.0.0      255.255.0.0     172.19.1.1      indirect  rip          5  1
172.28.0.0      255.255.0.0     172.19.1.1      indirect  rip          3  1
172.30.0.0      255.255.0.0     172.19.1.1      indirect  rip          3  1
205.178.13.0    255.255.255.0   172.19.1.1      indirect  rip          2  1
205.178.15.0    255.255.255.0   172.19.1.1      indirect  rip          3  1
205.178.16.0    255.255.255.0   172.19.1.1      indirect  rip          3  1
205.178.17.0    255.255.255.0   172.19.1.1      indirect  rip          3  1
205.178.18.0    255.255.255.0   172.19.1.1      indirect  rip          2  1
208.214.245.0   255.255.255.0   172.19.1.1      indirect  rip          5  1
224.0.0.0       224.0.0.0       0.0.0.0         martian   martian
```

The following table describes the `Type` parameters.

Table 4-5 IP Routing Type Parameters

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

The following table describes the `Tag` parameters.

Table 4-6 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 Web switch.
`icmp`	The address was learned via ICMP.
`snmp`	This address was configured through SNMP.
`addr`	The address belongs to one of the switch's IP interfaces.
`rip`	The address was learned by the Routing Information Protocol (RIP).
`broadcast`	Indicates a broadcast address.
`martian`	The address belongs to a filtered group.
`multicast`	Indicates a multicast address.
`vip`	Indicates a route destination that is a virtual server IP address. VIP routes are needed to advertise virtual server IP addresses via BGP.
`bgp`	The address was learned via Border Gateway Protocol (BGP)

/info/arp
ARP Information

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

The ARP information includes IP address and MAC address of each entry, address status flags (see Table 4-8 on page 61), VLAN and port for the address, and port referencing information.

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

Command Syntax and Usage
find *<IP address>* Displays a single ARP entry by IP address.
port *<port number (1-9)>* Displays the ARP entries on a single port.
vlan *<VLAN number (1-4094)>* Displays the ARP entries on a single VLAN.
refpt *<port number (1-9)>* Displays the ARP entries referenced by a single port.
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 which have referenced the address (empty if no port has routed traffic to the IP address shown) For more information, see page 61.
addr Displays the ARP address list: IP address, IP mask, MAC address, and VLAN flags.

/info/arp/dump

Show All ARP Entry Information

```
  IP address     Flags    MAC address     VLAN Port Referenced ports
--------------- ----- ----------------- ---- ---- ----------------
10.10.10.10      P 4  00:60:cf:40:78:ce    1       1-9
172.19.1.1            00:60:cf:42:e4:40    1   8   empty
172.19.1.61           00:10:a4:f0:4c:13    1   8   empty
172.19.1.201     P    00:60:cf:40:78:c0    1       1-9
```

NOTE – If you have VMA turned on, the referenced port will be the designated port. If you have VMA turned off, the designated port will be the normal ingress port.

The `Flag` field is interpreted as follows:

Table 4-8 ARP Dump Flag Parameters

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

/info/arp/addr

ARP Address List Information

```
   IP address        IP mask        MAC address    VLAN Flags
--------------- --------------- ----------------- ---- -----
205.178.18.66   255.255.255.255 00:70:cf:03:20:04        P
205.178.50.1    255.255.255.255 00:70:cf:03:20:06    1
205.178.18.64   255.255.255.255 00:70:cf:03:20:05    1
```

/info/ospf
OSPF Information

```
[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
      sumaddr - Show summary address list
      nsumadd - Show NSSA summary address list
      routes  - Show OSPF routes
```

Table 4-9

Command Syntax and Usage
general Displays general OSPF information. See page 63 for a sample output.
aindex *<area index [0-2]>* Displays area information for a particular area index. If no parameter is supplied, it displays area information for all the areas.
if *<interface number [1-256]>* Displays interface information for a particular interface. If no parameter is supplied, it displays information for all the interfaces. See page 63 for a sample output.
virtual Displays information about all the configured virtual links.
nbr *<nbr router-id [A.B.C.D]>* Displays the status of a neighbor with a particular router ID. If no router ID is supplied, it displays the information about all the current neighbors.
dbase *<interface number [1-256]>* Displays OSPF database menu. To view menu options, see page 64.
sumaddr *<area index [0-2]>* Displays the list of summary ranges belonging to non-NSSA areas.
nsumadd *<area index [0-2]>* Displays the list of summary ranges belonging to NSSA areas.
routes Displays OSPF routing table. See page 66 for a sample output.

/info/ospf/general

OSPF General Information

```
OSPF Version 2
Router ID: 172.25.1.11
Started at 51 and the process uptime is 162162
Area Border Router: no, AS Boundary Router: no
LS types supported are 6
External LSA count 0
External LSA checksum sum 0x0
Number of interfaces in this router is 0
Number of virtual links in this router is 0
0 new lsa received and 0 lsa originated from this router
Total number of entries in the LSDB 0
Database checksum sum 0x0
Total neighbors are 0, of which
                                    0 are >=INIT state,
                                    0 are >=EXCH state,
                                    0 are =FULL state
Number of areas is 0, of which 0-transit 0-nssa
```

/info/ospf/if

OSPF Interface Information

```
Ip Address 172.31.3.179, Area 1.1.1.1, Admin Status UP
   Router ID 172.31.3.179, State DR, Priority 1
   Designated Router (ID) 172.31.3.179, Ip Address 172.31.3.179
   Backup Designated Router (ID) 0.0.0.0, Ip Address 0.0.0.0
   Timer intervals, Hello 10, Dead 40, Wait 52, Retransmit 5,
                Poll interval 0, Transit delay 1
   Neighbor count is 0   If Events 2, Authentication type None
```

/info/ospf/dbase

OSPF Database Information

```
[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 4-10 OSPF Database Information Menu (/info/ospf/dbase)

Command Syntax and Usage
advrtr *<router-id (A.B.C.D)>* Takes advertising router as a parameter. Displays all the Link State Advertisements (LSAs) in the LS database that have the advertising router with the specified router ID, for example: 20.1.1.1.
asbrsum *<adv-rtr (A.B.C.D)>* | *<link_state_id (A.B.C.D>* | *<self>* Displays ASBR summary LSAs. The usage of this command is as follows: a) `asbrsum adv-rtr 20.1.1.1` displays ASBR summary LSAs having the advertising router 20.1.1.1. b) `asbrsum link_state_id 10.1.1.1` displays ASBR summary LSAs having the link state ID 10.1.1.1. c) `asbrsum self` displays the self advertised ASBR summary LSAs. d) `asbrsum` with no parameters displays all the ASBR summary LSAs.
dbsumm Displays the following information about the LS database in a table format: a) the number of LSAs of each type in each area. b) the total number of LSAs for each area. c) the total number of LSAs for each LSA type for all areas combined. d) the total number of LSAs for all LSA types for all areas combined. No parameters are required.
ext *<adv-rtr (A.B.C.D)>* | *<link_state_id (A.B.C.D>* | *<self>* Displays the AS-external (type 5) LSAs with detailed information of each field of the LSAs. The usage of this command is the same as the usage of the command `asbrsum`.

Table 4-10 OSPF Database Information Menu (/info/ospf/dbase)

Command Syntax and Usage
nw *<adv-rtr (A.B.C.D)> | <link_state_id (A.B.C.D> | <self>* Displays the network (type 2) LSAs with detailed information of each field of the LSA.network LS database. The usage of this command is the same as the usage of the command `asbrsum`.
nssa *<adv-rtr (A.B.C.D)> | <link_state_id (A.B.C.D> | <self>* Displays the NSSA (type 7) LSAs with detailed information of each field of the LSAs. The usage of this command is the same as the usage of the command `asbrsum`.
rtr *<adv-rtr (A.B.C.D)> | <link_state_id (A.B.C.D> | <self>* Displays the router (type 1) LSAs with detailed information of each field of the LSAs. The usage of this command is the same as the usage of the command `asbrsum`.
self Displays all the self-advertised LSAs. No parameters are required.
summ *<adv-rtr (A.B.C.D)> | <link_state_id (A.B.C.D> | <self>* Displays the network summary (type 3) LSAs with detailed information of each field of the LSAs. The usage of this command is the same as the usage of the command `asbrsum`.
all Displays all the LSAs.

/info/ospf/routes

OSPF Information Route Codes

```
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
 IA 10.10.0.0/16 via 200.1.1.2
 IA 40.1.1.0/28 via 20.1.1.2
 IA 80.1.1.0/24 via 200.1.1.2
 IA 100.1.1.0/24 via 20.1.1.2
 IA 140.1.1.0/27 via 20.1.1.2
 IA 150.1.1.0/28 via 200.1.1.2
 E2 172.18.1.1/32 via 30.1.1.2
 E2 172.18.1.2/32 via 30.1.1.2
 E2 172.18.1.3/32 via 30.1.1.2
 E2 172.18.1.4/32 via 30.1.1.2
 E2 172.18.1.5/32 via 30.1.1.2
 E2 172.18.1.6/32 via 30.1.1.2
 E2 172.18.1.7/32 via 30.1.1.2
 E2 172.18.1.8/32 via 30.1.1.2
 E2 172.18.1.9/32 via 30.1.1.2
 E2 172.18.1.10/32 via 30.1.1.2
 E2 172.18.1.11/32 via 30.1.1.2
 E2 172.18.1.12/32 via 30.1.1.2
 E2 172.18.1.12/32 via 30.1.1.2
```

/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
      refpt   - Show FDB entries referenced by a single port
      dump    - Show all FDB entries
```

The forwarding database (FDB) contains information that maps the media access control (MAC) 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.

NOTE – The master forwarding database supports up to 2K MAC address entries on the MP per switch. Each SP port supports up to 1K entries.

Table 4-11 FDB Information Menu Options (/info/fdb)

Command Syntax and Usage
find *<MAC address>* [*<VLAN>*] Displays a single database entry by its MAC address. You are prompted to enter the MAC address of the device. Enter the MAC address using the format, `xx:xx:xx:xx:xx:xx`. For example, `08:00:20:12:34:56`. You can also enter the MAC address using the format, `xxxxxxxxxxxx`. For example, `080020123456`.
port *<port number (1-9)>* Displays all FDB entries for a particular port.
vlan *<VLAN number (1-4094)>* Displays all FDB entries on a single VLAN.
refpt *<port number (1-9)>* Displays the FDB entries referenced by a single port.
dump Displays all entries in the Forwarding Database. For more information, see page 68.

/info/fdb/dump

Show All FDB Information

```
    MAC address      VLAN  Port  State  Referenced ports
 -----------------  ----  ----  -----  ----------------
 00:10:4b:76:ab:04     1     9    FWD    9
 00:60:cf:42:08:e0     1     9    FWD    1 9
 00:60:cf:42:39:80     1     9    FWD    9
 00:60:cf:42:72:40     1     9    FWD    9
 00:60:cf:44:59:60     1     9    FWD    9
 00:60:cf:51:97:05     1     1    FWD    1 9
```

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

If the state for the port is listed as an interface (IF), the MAC address is for a standard VRRP virtual router. If the state is listed as a virtual server (VIP), 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 forwarding database (FDB) or to clear the entire FDB, refer to “Forwarding Database Options” on page 303.

/info/sys
System Information

```
System Information at 11:53:07 Thu Nov  8, 2001

Alteon AD4
sysName:     techpubs
sysLocation:

Switch is up 6 days, 0 hours, 59 minutes and 26 seconds.
Last boot: 10:53:38 Fri Nov  2, 2001 (reset from Telnet)

MAC address: 00:60:cf:44:98:80    IP (If 1) address: 172.25.1.11
Hardware Revision: B
Hardware Part No: C04_5A-D_6A-D
Software Version 10.0.11 (FLASH image1), active configuration.

Temperature:   Rear Left Sensor    = 26C    Rear Middle Sensor  = 25C
               Front Middle Sensor = 33C    Front Right Sensor  = 41C
```

NOTE – The display of temperature will come up only if any of the temperatures exceed 50C.

System information includes:

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

/info/log
Show Last 30 Syslog Messages

```
Jan 21 15:12:45 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 21 15:19:11 NOTICE  mgmt: admin idle timeout from Telnet/SSH
Jan 21 15:20:45 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 21 15:27:16 NOTICE  mgmt: admin idle timeout from Telnet/SSH
Jan 21 15:47:07 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 21 15:57:24 NOTICE  mgmt: admin idle timeout from Telnet/SSH
Jan 21 17:30:57 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 21 17:31:31 INFO    mgmt: new configuration applied
Jan 21 17:37:05 NOTICE  mgmt: admin idle timeout from Telnet/SSH
Jan 22 12:04:03 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 22 12:13:36 NOTICE  mgmt: admin idle timeout from Telnet/SSH
Jan 22 12:40:09 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 22 12:44:50 NOTICE  mgmt: admin connection closed from Telnet/SSH
Jan 22 12:45:34 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 22 12:45:47 NOTICE  mgmt: admin connection closed from Telnet/SSH
Jan 22 12:46:33 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 22 12:51:49 NOTICE  mgmt: admin idle timeout from Telnet/SSH
Jan 22 12:52:55 NOTICE  mgmt: admin login from host 47.83.236.121
Jan 22 12:53:15 NOTICE  mgmt: admin login from host 47.83.236.201
Jan 22 12:53:56 NOTICE  mgmt: admin connection closed from Telnet/SSH
Jan 22 12:57:33 NOTICE  mgmt: admin connection closed from Telnet/SSH
Jan 23 11:41:51 NOTICE  mgmt: admin login from host 47.83.236.201
```

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, depending on the condition that the administrator is being notified of, as shown below.

- `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 debut-level message

/info/link
Link Status Information

```
------------------------------------------------------------------
Port   Speed    Duplex     Flow Ctrl      Link
----   -----   --------  --TX-----RX--   ------
  1      100     full     yes    yes       up
  2    10/100     any     yes    yes      down
  3    10/100     any     yes    yes      down
  4    10/100     any     yes    yes      down
  5    10/100     any     yes    yes      down
  6    10/100     any     yes    yes      down
  7    10/100     any     yes    yes      down
  8    10/100     any     yes    yes      down
  9      100     full     yes    yes       up
```

Use this command to display link status information about each port on an Alteon Web switch slot, including:

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

/info/stp
Spanning Tree Information

```
Spanning Tree Group 1: On

Current Root:            Path-Cost  Port Hello MaxAge FwdDel Aging
 7ffe 00:60:cf:46:74:c0       11       5    2     20     15    300

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

Port  Priority  Cost      State       Designated Bridge     Des Port
----  --------  ----   -----------  ----------------------  --------
  1       128      0    DISABLED
  2       128      0    DISABLED
  3       128      0    DISABLED
  4       128      0    DISABLED
  5       128      5   FORWARDING   8000-00:60:cf:46:94:30    32771
  6       128      0    DISABLED
  7       128      0    DISABLED
  8       128      0    DISABLED
  9       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 and priority
- Cost
- State

The following table describes the STP parameters.

Table 4-12 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 reconfigure 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 auto negotiated.
`State`	The state field shows the current state of the port. The state field can be either `BLOCKING`, `LISTENING`, `LEARNING`, `FORWARDING`, or `DISABLED`.

/info/vlan
VLAN Information

```
VLAN                Name                Status Jumbo BWC  Ports
----  --------------------------------- ------ ----- ---- -----------
1     Default VLAN                        ena    n  1024  1 4-9
2     VLAN 2                              ena    n  1024  3
3     VLAN 3                              ena    n  1024  empty
4     VLAN 4                              ena    n  1024  2
```

This information display includes all configured VLANs and all member ports that have an active link state. Port membership is represented in slot/port format.

VLAN information includes:

- VLAN Number
- VLAN Name
- Status
- Jumbo Frame usage (y or n)
- Bandwidth Contract if BWM is enabled
- Port membership of the VLAN

/info/port
Port Information

```
------------------------------------------------------------------
  Port  Tag  RMON  PVID   BWC        NAME              VLAN(s)
----  ---  ----  ----  -----  --------------  ----------------------
  1    n     d      1   1024                      1
  2    n     d      4   1024                      4
  3    n     d      2   1024  servers             2
  4    n     d      1   1024                      1
  5    n     d      1   1024                      1
  6    n     d      1   1024                      1
  7    n     d      1   1024                      1
  8    n     d      1   1024                      1
  9    n     d      1   1024                      1
```

Port information includes:

- Port number
- Whether the port uses VLAN tagging or not (`y` or `n`)
- Port VLAN ID (PVID)
- Port name
- VLAN membership
- Whether RMON is enabled or disabled on the port

/info/ip
IP Information

```
Interface information:
  1: 172.25.1.11     255.255.255.0   172.25.1.255,    vlan 1, up

Default gateway information: metric strict
  1: 172.25.1.1,      up

Current IP forwarding settings: ON, dirbr disabled

Current local networks:

Current IP port settings:
  All other ports have forwarding ON

Current OSPF settings: ON
  Default route none
  Router ID: 172.25.1.11

Current OSPF interface settings:
  1: 172.25.1.11,     area index  0, prio 5, cost 1, disabled
    hello 1, dead 40, trans 1, retra 5
```

/info/vrrp
VRRP Information

Virtual Router Redundancy Protocol (VRRP) support on Alteon Web switches provides redundancy between routers in a LAN. This 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 2, 205.178.18.210, if  1, renter, prio 100, master, server
  2: vrid 1, 205.178.18.202, if  1, renter, prio 100, backup
  3: vrid 3, 205.178.18.204, if  1, renter, prio 100, master, proxy
```

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/trunk
Trunk Group Information

```
Group  Slot  Port  State
------- ----- ------ -------
  1      2      4     DOWN
         2      5     DOWN
         2      6     DOWN
         2      9     DOWN
  2      3      1     forwarding
         4      3     DOWN
```

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`, the remaining ports in the trunk group will also be set to `forwarding`.

/info/bwm <*contract number*>
Bandwidth Management Information

```
Current Bandwidth Management setting: ON
  Policy Enforcement:enabled
  BwM history will be mailed in 10 hour(s) 1 minute(s)
    to '' at host ''
Contract              Policy                                  wTOS          Save
 Num   Name                 Prec  Hard  Soft  Resv oTOS uTOS  Buffer Hist State
   1   mon                   1    10   15m    5m  500k   0D   0    8219   E    E
   2   tue                   1     1   15m    5m  500k   0D   0    8219   E    E
   3   wed                   1     1   15m    5m  500k   0D   0    8219   E    E
   4   thu                   1     1   15m    5m  500k   0D   0    8219   E    E
   5   fri                   1     1   15m    5m  500k   0D   0    8219   E    E
   6   ser80                 1     1   15m    5m  500k   0D   0    8219   E    E
   7   vip                   1   255   15m    5m  500k   0D   0    8219   E    E
  11   bcon_tue              2     1   40m   30m  500k   0D   0   16320   E    E
 256   Default              --     0  --Available BW--    0D   0   16320   E    E
```

NOTE – The contract number is optional.

/info/swkey
Software Enabled Keys

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

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

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). 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

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

/stats

Statistics Menu

```
[Statistics Menu]
      port    - Port Stats Menu
      slb     - Server Load Balancing Stats Menu
      bwm     - Bandwidth Management Stats Menu
      mp      - MP-specific Stats Menu
      if      - Show IP interface (“if”) stats
      ip      - Show IP stats
      icmp    - Show ICMP stats
      tcp     - Show TCP stats
      udp     - Show UDP stats
      snmp    - Show SNMP stats
      fdb     - Show FDB stats
      route   - Show route stats
      arp     - Show ARP stats
      dns     - Show DNS stats
      ospf    - OSPF Statistics Menu
      vrrp    - Show VRRP stats
      dump    - Dump all stats
```

Table 5-1 Statistics Menu Options (/stats)

Command Syntax and Usage
port *<port number (1-9)>* Displays the Port Statistics Menu for the specified port. Use this command to display traffic statistics on a port-by-port basis. Traffic statistics are included in SNMP Management Information Base (MIB) objects. To view menu options, see page 84.
slb Displays the Server Load Balancing (SLB) Menu. To view menu options, see page 100.
bwm Displays the Bandwidth Management Menu. To view menu options, see page 130.
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 135.
if *<interface number (1-256)>* Displays IP interface statistics for the management processors. See page 143 for sample output.
ip Displays IP statistics. See page 145 for sample output.
icmp Displays ICMP statistics. See page 148 for sample output.
tcp Displays TCP statistics. See page 151 for sample output.
udp Displays UDP statistics. See page 151 for sample output.
snmp Displays SNMP statistics. See page 152 for sample output.
fdb Displays FDB statistics. See page 155 for sample output.
route Displays route statistics. See page 156 for sample output.
arp Displays Address Resolution Protocol (ARP) statistics. See page 157 for sample output.

Table 5-1 Statistics Menu Options (/stats)

Command Syntax and Usage
dns Displays Domain Name Server (DNS) statistics. See page 158 for sample output.
ospf Displays OSPF statistics Menu. See page 158 for sample output.
vrrp 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`) See page 163 for sample output.
dump 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. For details, see page 163.

/stats/port <port number>
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
      ip      - Internet Protocol (“IP”) stats
      link    - Link stats
      rmon    - Show RMON stats
      cpu     - Show CPU utilization
      maint   - Maintenance stats
```

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

Command Syntax and Usage
brg Displays bridging (“dot1”) statistics for the port. See page 85 for sample output.
ether Displays Ethernet (“dot1”) statistics for the port. See page 86 for sample output.
if Displays interface statistics for the port. See page 89 for sample output.
ip Displays IP statistics for the port. See page 91 for sample output.
link Displays link statistics for the port. See page 92 for sample output.
rmon Displays RMON statistics for the port. See page 93 for sample output.
cpu Displays CPU statistics for the port. CPU statistics are for periods of 1, 4, and 64 seconds. See page 96 for sample output.
maint Displays maintenance statistics for the port. See page 97 for sample output.

/stats/port <port number>/brg

Bridging Statistics

This menu option enables you to display the bridging statistics of the selected port.

```
Bridging statistics for port 1:
dot1PortInFrames:                       523874
dot1PortOutFrames:                     2497887
dot1PortInDiscards:                          0
dot1TpLearnedEntryDiscards:                  0
dot1BasePortDelayExceededDiscards:           0
dot1BasePortMtuExceededDiscards:             0
dot1StpPortForwardTransitions:               0
```

Table 5-3 Bridging Statistics of a Port (/stats/port/brg)

Statistics	Description
dot1PortInFrames	The number of frames that have been received by this port from its segment. 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	Count of valid frames received which were discarded (that is, filtered) by the Forwarding Process.
dot1TpLearnedEntry Discards	The total number of Forwarding Database entries, which have been or would have been learnt, but have been discarded due to a lack of space to store them 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.
dot1BasePortDelay ExceededDiscards	The number of frames discarded by this port due to excessive transit delay through the bridge. It is incriminated by both transparent and source route bridges.
dot1BasePortMtu ExceededDiscards	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.

/stats/port <port number>/ether

Ethernet Statistics

This menu option enables you to display the ethernet statistics of the selected port

```
Ethernet statistics for port 1:
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 Statistics for Port (/stats/port/ether)

Statistics	Description
`dot3StatsAlignment Errors`	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 Logical Link Control (LLC) (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
`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 obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.

Table 5-4 Ethernet Statistics for Port (/stats/port/ether)

Statistics	Description
`dot3StatsSingle-CollisionFrames`	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`, `ifOutMul-ticastPkts`, or `ifOutBroadcastPkts`, and is not counted by the corresponding instance of the `dot3StatsMultipleCollision-Frame` object.
`dot3StatsMultiple-CollisionFrames`	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`, `ifOutMul-ticastPkts`, or `ifOutBroadcastPkts`, and is not counted by the corresponding instance of the `dot3StatsSingleCollision-Frames` object.
`dot3StatsSQETest-Errors`	A count of times that the SQE TEST ERROR message is generated by the PLS sub layer 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 genera-tion is described in section 7.2.4.6 of the same document.
`dot3StatsDeferred-Transmissions`	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.
`dot3StatsLate-Collisions`	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 microseconds on a 10 Mbit/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for pur-poses of other collision-related statistics.
`dot3StatsExcessive Collisions`	A count of frames for which transmission on a particular interface fails due to excessive collisions.

Table 5-4 Ethernet Statistics for Port (/stats/port/ether)

Statistics	Description
dot3StatsInternal-MacTransmitErrors	A count of frames for which transmission on a particular interface fails due to an internal MAC sub layer 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 dot3Stats-CarrierSenseErrors 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.
dot3StatsCarrier-SenseErrors	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface. The count represented by an instance of this object is incremented at most once per transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
dot3StatsFrameToo-Longs	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 incremented when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
dot3StatsInternal-MacReceiveErrors	A count of frames for which reception on a particular interface fails due to an internal MAC sub layer 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 dot3Stats-AlignmentErrors 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 received errors on a particular interface that are not otherwise counted.
dot3Coll-Frequencies	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.

/stats/port <port number>/if

Interface Statistics

This menu option enables you to display the interface statistics of the selected port.

```
Interface statistics for port 1:
                     ifHCIn Counters       ifHCOut Counters
Octets:                     12046215              86913772
UcastPkts:                    187106                211622
BroadcastPkts:                   248                294478
MulticastPkts:                     8                554238
Discards:                          0                     0
Errors:                            0                     0
ifInUnknownProtos:                 0
```

Table 5-5 Interface Statistics for Port (/stats/port/if)

Statistics	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 to 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 delivered 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 delivered 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.

Table 5-5 Interface Statistics for Port (/stats/port/if)

Statistics	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 to 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 to 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 to 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.

/stats/port <port number>/ip

Interface Protocol Statistics

This menu option enables you to display the interface statistics of the selected port.

```
IP statistics for port 1:
ipInReceives:              0   ipInHdrErrors:           0
ipInAddrErrors:            0   ipForwDatagrams:         0
ipInUnknownProtos:         0   ipInDiscards:            0
ipInDelivers:              0   ipCacheFull:             0
ipTtlExceeds:              0   ipQueueFull:             0
ipFragOKs:                 0   ipFragCreates:           0
ipDontFrags:               0   ipFragFails:             0
```

Table 5-6 Interface Protocol Statistics (/stats/port/ip)

Statistics	Description
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInHdrErrors	The number of input datagrams discarded due to 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, and so forth.
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 entity (the switch). This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities which 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 entity (the switch) 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 entity (the switch), 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 example, for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.

Table 5-6 Interface Protocol Statistics (/stats/port/ip)

Statistics	Description
ipInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
ipCacheFull	The number of ARP entry added to a FULL ARP cache. Currently not supported.
ipTtlExceeds	The number of IP datagram for which an ICMP TTL exceeded message was sent.
ipQueueFull	The number of times an ARP/route cache entry could not be added to an SP due to memory constraints.
ipFragOKs	The number of IP datagrams that have been successfully fragmented at this entity (the switch).
ipFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity (the Switch).
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 entity (the switch) but could not be, for example, because their Don't Fragment flag was set.

/stats/port <*port number*>/link

Link Statistics

This menu enables you to display the link statistics of the selected port.

```
Link statistics for port :   1
linkStateChange:             9
```

Table 5-7 Link Statistics (/stats/port/link)

Statistics	Description
linkStateChange	The total number of link state changes.

/stats/port <port number>/rmon

RMON Statistics

This menu option enables you to display the remote monitor statistics of the selected port.

```
RMON statistics for port 1:
etherStatsDropEvents:                    0
etherStatsOctets:                      192
etherStatsPkts:                          3
etherStatsBroadcastPkts:                 0
etherStatsMulticastPkts:                 0
etherStatsCRCAlignErrors:                0
etherStatsUndersizePkts:                 0
etherStatsOversizePkts:                  0
etherStatsFragments:                     0
etherStatsJabbers:                       0
etherStatsCollisions:                    0
etherStatsPkts64Octets:                  3
etherStatsPkts65to127Octets:             0
etherStatsPkts128to255Octets:            0
etherStatsPkts256to511Octets:            0
etherStatsPkts512to1023Octets:           0
etherStatsPkts1024to1518Octets:          0
```

Table 5-8 Remote Monitor Statistics (/stats/port/rmon)

Statistics	Description
`etherStatsDrop Events`	The total number of events in which packets were dropped by the probe due to 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.
`etherStatsOctets`	The total number of octets of data (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 (which is the percent utilization of the ethernet segment). 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 percent value of utilization.

Table 5-8 Remote Monitor Statistics (/stats/port/rmon)

Statistics	Description
`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 Frame Check Sequence (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.
`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 due to collisions) and noise hits. (A runt is a packet that is less than 64 bytes.)
`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 (10Base-5) and section 10.3.1.4 (10Base-2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 milliseconds and 150 milliseconds.

Table 5-8 Remote Monitor Statistics (/stats/port/rmon)

Statistics	Description
etherStats-Collisions	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.
etherStatsPkts64-Octets	The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including Frame Check Sequence (FCS) octets).
etherStatsPkts65-to127Octets	The total number of packets (including bad packets) received that were between 65 and 127 octets in length (excluding framing bits but including FCS octets).
etherStatsPkts128-to255Octets	The total number of packets (including bad packets) received that were between 128 and 255 octets in length (excluding framing bits but including Frame Check Sequence (FCS) octets).
etherStatsPkts256-to511Octets	The total number of packets (including bad packets) received that were between 256 and 511 octets in length (excluding framing bits but including FCS octets).
etherStatsPkts512-to1023Octets	The total number of packets (including bad packets) received that were between 512 and 1023 octets in length (excluding framing bits but including FCS octets).
etherStatsPkts-1024to1518Octets	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length (excluding framing bits but including FCS octets).

/stats/port <port number>/cpu

CPU Statistics

This menu option enables you to display the CPU statistics of the selected port.

```
CPU utilization for port 1:
cpuAUtil1Second:              2%  cpuBUtil1Second:              1%
cpuAUtil4Seconds:             2%  cpuBUtil4Seconds:             1%
cpuAUtil64Seconds:            2%  cpuBUtil64Seconds:            1%
```

Table 5-9 CPU Statistics of a Port (/stats/port/cpu)

Statistics	Description
cpuAUtil1Second	The utilization of port CPU A over 1 second. It shows the percentage.
cpuBUtil1Second	The utilization of port CPU B over 1 second. It shows the percentage.
cpuAUtil4Seconds	The utilization of port CPU A over 4 seconds. It shows the percentage.
cpuBUtil4Seconds	The utilization of port CPU B over 4 seconds. It shows the percentage.
cpuAUtil64Seconds	The utilization of port CPU A over 64 seconds. It shows the percentage.
cpuBUtil64Seconds	The utilization of port CPU B over 64 seconds. It shows the percentage.

/stats/port <port number>/maint

Maintenance Statistics

This menu option enables you to display the maintenance statistics of the selected port.

```
Maintenance statistics for port 1:
ddwOvflo:                    0   ddwOvfloOvflo:              0
dmaRdOverrun:                0   dmaRdUnderrun:              0
dmaWrOverrun:                0   dmaWrUnderrun:              0
txFlowCntrled:               0   rxFlowCntrled:              0
linkStateChange:             1   macRxBufClean:              1
pfdbFreeEmpty:               0   macRxBufCldma:              0
resolveErrNoddw:             0   macRxBufCldmaOvflo:         0
learnErrNoddw:               0   macRxBufClMacDescr0:        1
deleteMiss:                  0   macRxBufClMacDescr1:        0
txOvfloOvflo:                0   macRxBufClMacDescrN:        0
mac_rx_err:
  CRC               0    Collided          0
  LinkErr           0    PhyErr            0
  Nibble            0    Abort             0
  Runt              0    NoBuffer          0
  Over32k           0    Over16k           0
  Over9k            0
```

Table 5-10 Maintenance Statistics for Selected Port (/stats/port/maint)

Statistics	Description
ddwOvflo	The number of frames that were placed on the software controlled DMA descriptor ring. The frame was *not* dropped.
ddwOvfloOvflo	The number of frames that were dropped because there was no room on either the hardware or software DMA descriptor ring.
dmaRdOverrun	The number of times the Read DMA state machine encountered an overrun error, size of read operation exceeded size of FIFO. It is a fatal error.
dmaRdUnderrun	The number of times the Read DMA state machine encountered an underrun error, read operation read more than what was available. It is a fatal error.
dmaWrOverrun	The number of times the Write DMA state machine encountered an overrun error, write operation tried to write more than what was available. Not fatal, Write DMA will be retried.
dmaWrUnderrun	The number of times the unknown fatal errors occurred.
txFlowCntrled	The number of flow control frames transmitted.

Table 5-10 Maintenance Statistics for Selected Port (/stats/port/maint)

Statistics	Description
rxFlowCntrled	The number of flow control frames received.
linkStateChange	The number of times the port lost or gained link.
macRxBufClean	The number of times a MAC RX buffer attention event (a warning that RX buffer is almost full) occurred.
pfdbFreeEmpty	The number of times the attempts were made to add entries to an already full FDB.
macRxBufCldma	The number of times an RX buffer attention event occurred and no cleanup occurred because frames were waiting to be DMA transferred.
resolveErrNoddw	The number of times MAC address resolution failed because there were no DMA descriptors.
macRxBufCldmaOvflo	The number of times an RX buffer attention occurred and no cleanup occurred because frames were waiting to be DMA transferred on the software DMA descriptor ring.
learnErrNoddw	The number of times MAC address `learn` failed because there were no DMA descriptors.
macRxBufClMac-Descr0	The number of times an RX buffer attention event (a warning that RX buffer is almost full) resulted in zero frames being dropped.
txOvfloOvflo	The number of frames dropped because there were no TX MAC descriptors available.
macRxBufClMac-DescrN	The number of times an RX buffer attention resulted in ALL frames in RX buffer being dropped. Usually occurs when port is in blocking mode, for example, in STP.
mac_rx_err:	
CRC	The number of times frames have been dropped due to CRC errors.
Collided	The number of frames received that contained a collision.
LinkErr	The number of times the link was de-asserted during the reception of a packet.
PhyErr	The number of frames received that the PHY (physical) detected packet error.
Nibble	The number of frames received with an odd number of nibbles (four bits).
Abort	The number of frames received that were aborted by MAC.

Table 5-10 Maintenance Statistics for Selected Port (/stats/port/maint)

Statistics	Description
`Runt`	The number of runt frames (runt is a frame that is less than 64 bytes) received.
`NoBuffer`	The number of frames that would not fit in the RX buffer.
`Over32k`	The number of frames received that were over 32K bytes.
`Over16k`	The number of frames received that were over 16K bytes.
`Over9k`	The number of frames received that were over 9K bytes.

/stats/slb
Load Balancing Statistics Menu

```
[Server Load Balancing Statistics Menu]
      port    - SLB Switch Port Stats Menu
      real    - Show real server stats
      group   - Show real server group stats
      virt    - Show virtual server stats
      filt    - Show filter stats
      gslb    - Show global SLB stats
      layer7  - Show Layer 7 stats
      ssl     - Show SSL SLB stats
      ftp     - Show FTP SLB parsing and NAT stats
      rtsp    - Show RTSP SLB stats
      dns     - Show DNS SLB stats
      wap     - Show WAP SLB stats
      tcp     - Show TCP rate limiting stats
      maint   - Show maintenance stats
      clear   - Clear non-operational Server Load Balancing stats
      dump    - Dump all SLB statistics
```

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

Command Syntax and Usage

port *<port number (1-9)>*

Displays the switch port statistics. See page 103 for sample output.

real *<real server number (1-255)>*

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 the real server was assigned
- Highest number of simultaneous sessions recorded for each real server
- Real server transmit/receive octets

See page 108 for sample output.

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 per-service octet counters, see the procedure on page 108.

See page 109 for sample output.

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

Command Syntax and Usage
virt *<virtual server number (1-256)>* 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-service octet counters, see page 110. See page 110 for sample output.
filt *<filter ID (1-224)>* Displays the total number of times any filter has been used. See page 110 for sample output.
gslb Displays the Global SLB Statistics Menu. For more information, see page 111.
layer7 Displays Layer 7 statistics. See page 114 for sample output.
ssl Displays SSL server load balancing statistics. See page 118 for sample output.
ftp Displays FTP SLB parsing and NAT statistics. See page 119 for sample output.
rtsp Displays RTSP SLB statistics. See page 122 for sample output.
dns Displays DNS SLB statistics. See page 123 for sample output.
wap Displays WAP SLB statistics. See page 124 for sample output.
tcp Displays statistics for TCP rate limiting. See page 124 for sample output.
maint Displays SLB maintenance statistics. See page 127 for sample output.

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

Command Syntax and Usage
clear [**y**|**n**] Clears all non-operating SLB statistics on the Web switch, resetting them to zero. This command does not reset the switch and does ***not*** affect the following counters: ■ Counters required for Layer 4 and Layer 7 operation (such as current real server sessions). ■ All related SNMP counters. To view the statistics reset by this command, refer to Table 5-34 on page 128.
dump Dumps all switch SLB statistics. Use this command to gather data for tuning and debugging switch performance. To save dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump command.

/stats/slb/port <port number>

Port SLB Statistics Menu

```
[Server Load Balancing Port Statistics Menu]
      real    - Show real server stats
      group   - Show real server group stats
      virt    - Show virtual server stats
      filt    - Show filter stats
      maint   - Show maintenance stats
      clear   - Clear port stats
```

Table 5-12 Switch Processor SLB Statistics Menu Options (/stats/slb/port)

Command Syntax and Usage
real *<real server number (1-1023)>* Displays real server statistics for the selected port. See page 104 for sample output.
group *<real server group number (1-256)>* Displays real server group statistics for the selected port. See page 104 for sample output.
virt *<virtual server number (1-256)>* Displays virtual server statistics for the selected port. See page 105 for sample output.
filt *<filter ID (1-224)>* Displays filter statistics for the selected port. See page 106 for sample output.
maint Displays maintenance statistics for the selected port. See page 106 for sample output.
clear Clears the following non-operating SLB statistics for this port, resetting them to zero: Real server stats: Octets, Total sessions Real server group: Octets, Total sessions Virtual server: Octets, Total sessions Total firings: Octets

/stats/slb/port <#>/real <server number>

Port Real Server SLB Statistics

```
Port 1 Real server 1 stats:
Current sessions:                    9
Total sessions:                     24
Octets:                         192000
```

Table 5-13 Port Real Server SLB Statistics (/stats/slb/port/real)

Statistics	Description
Current sessions	The total number of outstanding sessions that are established to the particular real server on the switch port.
Total sessions	The total number of sessions that have been established to the particular real server on the switch port.
Octets	The total number of octets sent by the particular server on the switch port.

/stats/slb/port <#>/group <group number>

Port Real Server Group SLB Statistics

```
Port 1 Real server group 1 stats:
                       Current      Total  Highest
Real IP address       Sessions   Sessions Sessions           Octets
---- --------------- -------- ---------- --------  ---------------
 20  200.100.10.14          9         24       16           192000
 21  200.100.10.15         12         23       15           184000
---- --------------- -------- ---------- --------  ---------------
                           21         47       31           376000
```

Table 5-14 Port Real Server Group SLB Statistics (stats/slb/port/group)

Statistics	Description
Current sessions	The total number of outstanding sessions that are established to the particular real server group on the switch port.
Total sessions	The total number of sessions that have been established to the particular real server group on the switch port.
Highest sessions	The highest number of sessions recorded for the real server group.
Octets	The total number of octets sent by the particular real server group on the switch port.

/stats/slb/port <#>/virt <server number>

Port Virtual Server SLB Statistics

```
Port 1 Virtual server 1 stats:
                       Current      Total  Highest
Real IP address       Sessions   Sessions Sessions           Octets
---- --------------- -------- ---------- --------  ---------------
 20  200.100.10.14          9         24       16            192000
 21  200.100.10.15         12         23       15            184000
---- --------------- -------- ---------- --------  ---------------
     200.100.13.1           21         47       31            376000
```

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

Table 5-15 Port Virtual Server SLB Statistics

Statistics	Description
`Current Sessions`	The total number of outstanding sessions that are established to the particular virtual server on the switch port.
`Total Sessions`	The total number of sessions that have been established to the particular virtual server on the switch port.
`Highest Sessions`	The highest number of sessions ever recorded for the particular virtual server on the switch port.
`Octets Sessions`	The total number of octets sent by the particular virtual server on the switch port.

/stats/slb/port <#>/filter <filter number>

Port Filter SLB Statistics

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

This menu option displays the total number of times a filter has been fired on a specific port.

/stats/slb/port <#>/maint <server number>

Port Maintenance SLB Statistics

```
Port 1 SLB Maintenance stats:
Maximum sessions:                   64512
Current sessions:                       0
  4 second average:                     0
 64 second average:                     0
Terminated sessions:                    0
Allocation failures:                    0
Non TCP/IP frames:                      0
TCP fragments:                          0
UDP datagrams:                          0
Incorrect VIPs:                         0
Incorrect Vports:                       0
No available real server:               0
Filtered (denied) frames:               0
VMA discards:                           0
```

Table 5-16 SLB Port Maintenance Statistics (/stats/slb/port/maint)

Statistics	Description
`Maximum sessions`	The maximum number of simultaneous sessions supported.
`Current sessions`	The current number of sessions.
`Allocation failures`	The number of times a session entry could not be created because the table was full.
`Non TCP/IP frames`	The number of times a frame was dropped because it was a non-TCP/IP frame addressed to a VMAC address received.
`TCP fragments`	The number of TCP fragments dropped.
`UDP datagrams`	The number of UDP datagrams dropped because the virtual service did not support UDP.

Table 5-16 SLB Port Maintenance Statistics (/stats/slb/port/maint)

Statistics	Description
`Incorrect VIPs`	The number of times a frame was dropped because it was addressed to a VMAC address, which was an unknown virtual IP address.
`Incorrect Vports`	The number of frames dropped to a valid virtual IP address but not a valid Vport (destination port).
`No available real server`	The number of times a frame was dropped because there were no real servers available to accept the session.
`Filtered (denied) frames`	The number of frames dropped because they matched a *deny* filter.
`VMA discards`	The number of frames discarded on a VMA enabled Switch as there is not enough memory available on the designated port.

/stats/slb/real <real server number>

Real Server SLB Statistics

```
Real server 1 stats:
Health check failures:              0
Current sessions:                   129
Total sessions:                     65478
Highest sessions:                   4343
Octets                              523824000
```

NOTE – Octets are provided per server, not per service, unless configured as described in "Per Service Octet Counters" on page 108.

Table 5-17 Real Server SLB Statistics (/stats/slb/real)

Statistics	Description
`Current sessions`	The total number of outstanding sessions that are established to the particular real server.
`Total sessions`	The total number of sessions that have been established to the particular real server.
`Highest sessions`	The highest number of sessions ever recorded for the particular real server.
`Octets`	The total number of octets sent by the particular real server.

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 Web switch, configure a real server with a real IP address for each service above.**

   Continuing the example above, 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 Web 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

```
Real server group 1 stats:
                        Current      Total  Highest
Real IP address        Sessions   Sessions Sessions           Octets
---- --------------- -------- ---------- --------  ---------------
  1  200.100.10.14          20         60        9           480000
  2  200.100.10.15          20         77       12           616000
---- --------------- -------- ---------- --------  ---------------
                            40        137       21          1096000
```

Real server group statistics include the following:

- Current and total sessions for each real server in the real server group.
- Current and total sessions for all real servers associated with the real server group.
- Highest number of simultaneous sessions recorded for each real server.
- Real server transmit/receive octets. For per-service octet counters, see the procedure on "Per Service Octet Counters" on page 108.

/stats/slb/virt <virtual server number>

Virtual Server SLB Statistics

```
Virtual server 1 stats:
                      Current      Total  Highest
Real IP address      Sessions   Sessions Sessions           Octets
---- --------------- -------- ---------- --------  ---------------
  1  200.100.10.14         20         60        9           480000
  2  200.100.10.15         20         77       12           616000
---- --------------- -------- ---------- --------  ---------------
     200.100.10.20         40        309       21          1096000
```

NOTE – The virtual server IP address is shown on the last line, below the real server IP addresses.

Virtual server statistics include the following:

- Current and total sessions for each real server associated with the virtual server.
- Current and total sessions for all real servers associated with the virtual server.
- Highest number of simultaneous sessions recorded for each real server.
- Real server transmit/receive octets. For per-service octet counters, see "Per Service Octet Counters" on page 108.

/stats/slb/filt <filter number>

Filter SLB Statistics

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

You can obtain the total number of times any filter has been used.

/stats/slb/gslb

Global SLB Statistics Menu

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

Table 5-18 Global SLB Statistics Menu Options (/stats/slb/gslb)

Command Syntax and Usage
real *<real server number (1-255)>* Where the real server number represents the real server ID on this switch, under which the remote server is configured. To view an example and description of what is displayed on-screen, see page 111.
group *<real server group number (1-256)>* To view an example and description of what is displayed on-screen, see page 112.
virtual *<virtual server number (1-256)>* To view an example and description of what is displayed on-screen, see page 112.
maint To view an example and description of what is displayed on-screen, see page 113.

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

Real Server Global SLB Statistics

```
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 sent by the remote real server/real server group
- Number of HTTP redirects sent by the remote real server/real server group

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

Real Server Group Global SLB Statistics

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

Real server group global statistics include the following:

- 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    205.178.13.55            16             21190
  http     r1    205.178.13.54            10             24120

  telnet   v1    205.178.13.55             4             31032
```

Virtual server global statistics include the following:

- Service: type of service running on the virtual server
- Server: type of server configuration and server ID number.
  - **v**# represents a local virtual server number
  - **r**# represents a remote site. Since each remote sites is configured on its peers as if it were a real server (with certain special properties), the number represents the real server ID on this switch, under which the remote server is configured.

- IP address of the server
- Response time: the average time (current weight) 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:

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

/stats/slb/layer7

SLB Layer7 Statistics Menu

```
[Layer 7 Statistics Menu]
      redir   - Show URL Redirection stats
      str     - Show SLB String stats
      maint   - Show Layer 7 Maintenance stats
```

Table 5-19 SLB Layer 7 Statistics Menu Options (/stats/slb/layer7)

Command Syntax & Usage
redir Displays URL Redirection statistics. See page 114 for a sample output..
str Displays SLB string statistics. See page 115 for a sample output.
maint Displays Layer 7 maintenance statistics. See page 116 for a sample output.

/stats/slb/layer7/redir

Layer7 Redirection Statistics

```
Total URL based web cache redirection stats:
Total cache server hits:                    0
Total origin server hits:                   0
Total none-GETs hits:                       0
Total 'Cookie: ' hits:                      0
Total no-cache hits:                        0
```

Table 5-20 Layer 7 Redirection Statistics (/stats/slb/layer7/redir)

Statistics	Description
`Total cache server hits`	The total number of HTTP requests redirected to the cache server.
`Total origin server hits`	The total number of HTTP requests forwarded to the origin server.
`Total none-GETs hits`	The total number of none GET requests forwarded to the origin server.
`Total 'Cookie:' hits`	The total number of cookie requests forwarded to the origin server.
`Total no-cache hits`	The total number of requests containing *no-cache header* forwarded to the origin server.

/stats/slb/layer7/str

Layer 7 SLB String Statistics

```
SLB String stats:
  ID SLB String                                        Hits
   1 any                                             1527115
   2 www.[abcdefghijklm]*.com                              0
   3 www.[nopqrstuvwxyz]*.com                              0
   4 www.junk.com                                          0
   5 www.abc.com                                           0
   6 www.[abcdefjhijklm]*.org                              0
   7 www.[nopqrstuvwxyz]*.org                              0
```

Table 5-21 Layer 7 SLB String Statistics (/stats/slb/layer7/str)

Statistics	Description
ID SLB String	The user-defined strings being used in URL matching.
Hits	The total number of instances that are load-balanced due to matching of the particular URL ID.

/stats/slb/layer7/maint

Layer 7 SLB Maintenance Statistics

```
Clients reset by switch on client side:            0
Clients reset by switch on server side:            0
Connection Splicing to support HTTP/1.1:           0
Half open connections:                             0
Switch retries:                                    0
Random early drops:                                0
Requests exceeded 4500 bytes:                      0
Invalid 3-way handshakes:                          0
Current SP[1] memory units:   128    Lowest:                   128
Current SP[2] memory units:   128    Lowest:                   128
Current SP[3] memory units:   128    Lowest:                   128
Current SP[4] memory units:   128    Lowest:                   128
Current SP[5] memory units:   128    Lowest:                   128
Current SP[6] memory units:   128    Lowest:                   128
Current SP[7] memory units:   128    Lowest:                   128
Current SP[8] memory units:   128    Lowest:                   128
Current SP memory units:     1024
Current SEQ buffer entries:     0    Highest:                    0
Current Data buffer use:        0    Highest:                    0
Current SP buffer entries:      0    Highest:                    0
Total Nonzero SEQ Alloc:        0
Total SEQ Buffer Allocs:        0    Total SEQ Frees:            0
Total Data Buffer Allocs:        0   Total Data Frees:            0
Alloc Fails - Seq buffers:      0    Alloc Fails - Ubufs:        0
Max sessions per bucket:        0    Max frames per session:     0
Max bytes buffered (sess):      0
```

Table 5-22 SLB Layer 7 Maintenance Statistics (/stats/slb/layer7/maint)

Statistics	Description
`Clients reset by switch on client side`	The number of reset frames sent to the client by the switch during client connection termination.
`Clients reset by switch on server side`	The number of reset frames sent to the server by the switch during server connection termination.
`Connection Splicing to support HTTP/1.1`	The total number of connection swapping between different real servers in supporting multiple HTTP/1.1 client requests.0
`Half open connections`	The total numbers of outstanding TCP connections that are half opened. It is incremented when the switch responds to TCP SYN packet and decremented upon receiving TCP SYN ACK packet from the requester.

Table 5-22 SLB Layer 7 Maintenance Statistics (/stats/slb/layer7/maint)

Statistics	Description
Switch retries	The total number of switch retries to connect to the real server.
Random early drops	The total number of SYN frames dropped when the buffer is low.
Requests exceeded 4500 bytes	The total number of GET requests that exceeded 4500 bytes.
Invalid 3-way hand-shakes	The total number of dropped frames because of invalid 3-way hand shakes.
Current SP memory units	The current available SP memory units.
Current SEQ buffer entries	The number of outstanding sequence buffers used.
Highest SEQ buffer entries	The highest number of sequence buffers ever used.
Current Data buffer use	The number of outstanding data buffers used.
Highest Data buffer use	The highest number of data buffers ever used.
Total Nonzero SEQ Alloc	The total number of sequence buffer allocated.2
Total SEQ Buffer Allocs	The total number of sequence buffer allocations.
Total SEQ Frees	The total number of sequence buffer is freed.
Total Data Buffer Allocs	The total number of buffers allocated to store client request.2
Total Data Frees	The total of number buffers freed.
Alloc Fails - Seq buffers	The number of times sequence buffer allocation failed.
Alloc Fails - Ubufs	The number of times the URL data buffer allocation failed.
Max sessions per bucket	The maximum number of items (sessions) allowed in the session table hash bucket chain.
Max frames per session	The maximum number of frames to be buffered per session.
Max bytes buffered (sess)	The maximum number of bytes to be buffered per session.

/stats/slb/ssl

SLB Secure Socket Layer Statistics

```
SSL SLB maintenance stats:
SessionId allocation fails:                    0
SessionId aged in max time:                    0
                          Current      Total  Highest
                         Sessions   Sessions Sessions
------------------------ -------- ---------- --------
Unique SessionIds               0          0        0
SSL connections                 0          0        0
Persistent Port Sessions        0          0        0
```

Table 5-23 SLB Secure Socket Layer Statistics (/stats/slb/ssl)

Statistics	Description
`SSL SLB maintenance stats`	Debug stats for SSL `SessionId` based persistence.
`SessionId allocation fails`	The number of times allocation of a session table entry failed when attempting to store a `SessionId` in the table.
`SessionId aged in max time`	Debug stat to find out how many sessions were slow aged-out. This should be 0 except when a work around for an SSL session aging bug is used.
The table shows the Current Sessions, the total sessions seen on the switch since last reset and the high water mark of current sessions for the following:	
`Unique SessionIds`	Many SSL sessions can use the same `SessionId`, these should all bind to the same server. This number shows the number of unique SSL sessions seen on the switch.
`SSL connections`	The number of different TCP connections using SSL service.
`Persistent Port Sessions`	The number of `SessionIds` maintained on port 9, to allow for persistence across different client ports.

/stats/slb/ftp

File Transfer Protocol SLB and Filter Statistics Menu

```
[FTP SLB parsing and Filter Statistics Menu]
      active  - Show active FTP NAT filter stats
      parsing - Show FTP SLB parsing server stats
      maint   - Show FTP maintenance stats
      dump    - Dump all FTP SLB/NAT stats
```

Table 5-24 FTP SLB Parsing and Filter Statistics Menu Options (/stats/slb/ftp)

Command Syntax and Usage
active Shows active FTP SLB parsing and filter statistics. See page 120 for sample output.
parsing Shows parsing statistics. See page 120 for sample output.
maint Shows maintenance statistics. See page 121 for sample output.
dump Shows all FTP SLB/NAT statistics. See page 121.

/stats/slb/ftp/active

Active FTP SLB Parsing and Filter Statistics

```
Total Active FTP NAT stats(PORT):
Total FTP:                              0
Total New Active FTP Index:             0
Active FTP NAT ACK/SEQ diff:            0
```

Table 5-25 Active FTP Slb Parsing and Filter statistics (/stats/slb/ftp/active)

Statistics	Description
`Total Active FTP NAT stats (PORT)`	The number of times the switch receives the `port` command from the client.
`Total FTP`	The number of times the switch receives both active and passive FTP connections.
`Total New Active FTP Index`	The number of times the switch creates a new index due to `port` command from the client.
`Active FTP NAT ACK/SEQ diff`	The difference in the numbers of ACK and SEQ that the Switch needs for packet adjustment.

/stats/slb/ftp/parsing

Passive FTP SLB Parsing Statistics

```
Total FTP SLB Parsing Stats(PASV):
Total FTP:                                  0
Total New FTP SLB parsing Index:            0
FTP SLB parsing ACK/SEQ diff:               0
```

Table 5-26 Passive FTP SLB Parsing Statistics (/stats/slb/ftp/parsing)

Statistics	Description
`Total FTP`	The number of times the switch receives both active and passive FTP connections.
`Total New FTP SLB parsing Index`	The number of times the switch creates a new index in response to the `pasv` command from the client.
`FTP SLB parsing ACK/SEQ diff`	The difference in the numbers of ACK and SEQ that the switch needs FTP SLB parsing.

/stats/slb/ftp/maint

FTP SLB Maintenance Statistics

```
FTP mode switch error:                    0
```

Table 5-27 FTP SLB Maintenance Statistics (/stats/slb/ftp/maint)

Statistics	Description
FTP mode switch error	The number of times the switch is not able to switch modes from active to passive and vice versa.

/stats/slb/ftp/dump

FTP SLB Statistics Dump

```
Total FTP :                               0
Total FTP NAT Filtered:                   0
Total new active FTP NAT Index:           0
Total new FTP SLB parsing Index:          0
FTP Active FTP NAT ACK/SEQ diff:          0
FTP SLB parsing ACK/SEQ diff:             0
FTP mode switch error:                    0
```

Table 5-28 FTP SLB Statistics Dump (/stats/slb/ftp/dump)

Statistics	Description
Total FTP	The total number of FTP sessions that occured.
Total FTP NAT Filtered	The total number of FTP NAT filter sessions that occured.
Total new active FTP NAT Index	The total number of new data sessions created for FTP NAT filter in active mode.
Total new FTP SLB parsing Index	The number of times the switch creates a new index in response to the pasv command from the client.
FTP Active FTP NAT ACK/SEQ diff	The total number of times the adjustment between ACK and SEQ occured on the filter.
FTP SLB parsing ACK/ SEQ diff	The difference in the numbers of ACK and SEQ that the switch needs for FTP SLB parsing.
FTP mode switch error	The number of times the switch could not switch mode from active to passive and vice versa.

/stats/slb/rtsp

RTSP SLB Statistics

```
    Control     UDP                          Connection  Buffer      Alloc
SP  Connection  Streams     Redirect         Denied      Allocs      Failures
--  ----------  ----------  ----------       ----------  ----------  ----------
 1           0           0           0                0           0           0
 2           0           0           0                0           0           0
 3           0           0           0                0           0           0
 4           0           0           0                0           0           0
 5           0           0           0                0           0           0
 6           0           0           0                0           0           0
 7           0           0           0                0           0           0
 8           0           0           0                0           0           0
 9           0           0           0                0           0           0
--  ----------  ----------  ----------       ----------  ----------  --------
            0           0           0                0           0           0
```

Table 5-29 RTSP SLB Statistics (/stats/slb/rtsp)

Statistics	Description
`ControlConnection`	The total number of TCP connections for RTSP control connection.
`UDP Streams`	The total number of UDP connections for data channels. The number depends upon the type of media player being used.
`Redirect`	The total number of times the connection got redirected.
`ConnectionDenied`	The total number of times the connections got denied due to shortage of resources or the real server being down.
`BufferAllocs`	The total number of buffer allocations used.
`AllocFailures`	The total number of times the buffer allocation failed.

/stats/slb/dns

DNS SLB Statistics

```
Total number of TCP DNS queries:                                    0
Total number of UDP DNS queries:                                    0
Total number of invalid DNS queries:                                0
Total number of multiple DNS queries:                               0
Total number of domain name parse errors:                           0
Total number of failed real server name matches:                    0
Total number of DNS parsing internal errors:                        0
```

Table 5-30 DNS SLB Statistics (/stats/slb/dns)

Statistics	Description
`Total number of TCP DNS queries`	The total number of DNS queries that received through TCP connections.
`Total number of UDP DNS queries`	The total number of DNS queries received through UDP requests.
`Total number of invalid DNS queries`	The total number of malformed DNS queries received.
`Total number of multiple DNS queries`	The total number of DNS queries that contain more than one domain name to be resolved. Currently only one domain name resolution per request is supported.
`Total number of domain name parse errors`	The total number of DNS queries that have short or invalid domain names to be resolved.
`Total number of failed real server name matches`	The total number of times the user failed to find a real server which has the same l7 strings that match the domain name to be resolved.
`Total number of DNS parsing internal errors`	The total number of out of memory and other unexpected errors the user gets while processing the DNS query.

/stats/slb/wap

WAP SLB Statistics

```
WAP Maintenance stats:
  current sessions:                  0
  allocation failures:               0
  incorrect VIPs:                    0
  incorrect Vports:                  0
  no available real server:          0
  requests to wrong SP:              0
-------------------------------------------------------------------
TPCP External Notification stats:
  add session reqs:            0    del session reqs:            0
  req fails- q full:           0    req fails- q full:           0
  req fails- SP dead:          0    req fails- SP dead:          0
  entries in use:              0    entries in use:              0
  max entries in use:          0    max entries in use:          0
-------------------------------------------------------------------
RADIUS Snooping stats:
  acct reqs:                   0    acct wrap reqs:              0
  acct start reqs:             0    acct update reqs:            0
  acct stop reqs:              0    acct bad reqs:               0
  add session reqs:            0    del session reqs:            0
  req fails- q full:           0    req fails- SP dead:          0
  req fails- DMA:              0    max entries in use:          0
```

Table 5-31 WAP SLB Statistics (/stats/slb/wap)

Statistics	Description
WAP Maintenance stats:	
current sessions	The number of session bindings currently in use.
allocation failures	Indicates instances where the switch ran out of available bindings for a port.
incorrect VIPs	Indicates the number of times the switch received a Layer 4 request for a virtual server which was not configured.
incorrect Vports	This dropped frames counter indicates that the virtual server has received frames for TCP/UDP services that have not been configured. Normally this indicates a mis-configuration on the virtual server or the client.
no available real server	This dropped frames counter indicates that all real servers are either out of service or at their maxcon limit.
requests to wrong SP	The number of session add/delete requests sent to the wrong SP.

Table 5-31 WAP SLB Statistics (/stats/slb/wap)

Statistics	Description
TPCP External Notification stats:	
add session reqs	The number of WAP session add requests via TPCP.
del session reqs	The number of WAP session delete requests via TPCP.
req fails- q full	The number of add-request failures due to request queue being full .
req fails- SP dead	The number of add-request failures due to dead target SP.
entries in use	The number of queue entries in use.
max entries in use	The maximum number of queue entries in use at one time for session delete requests via TPCP.
RADIUS Snooping stats:	
acct reqs	The number of RADIUS Accounting frames received.
acct wrap reqs	The number of wrapped RADIUS Accounting frames received.
acct start reqs	The number of RADIUS Accounting Start frames received.
acct update reqs	The number of RADIUS Accounting Update frames.
acct stop reqs	The number of RADIUS Accounting Stop frames received.
acct bad reqs	The number of bad RADIUS Accounting frames received.
add session reqs	The number of WAP session add requests via RADIUS snooping.
del session reqs	The number of WAP session delete requests via RADIUS snooping.
req fails- q full	The number of add/delete requests failed due to request queue being full.
req fails- SP dead	The number of add/delete request failures due to dead target SP.
req fails- DMA	The number of add/delete requests failed due to DMA write failure.
max entries in use	The maximum number of queue entries in use at a time for session add/delete requests via RADIUS snooping.

/stats/slb/tcp

SLB TCP Rate Limiting Statistics

```
TCP rate limiting stats:
  Total hold downs triggered:                0
  Current per-client state entries:          0
```

Table 5-32 SLB TCP Rate Limiting Statistics (/stats/slb/tcp)

Statistics	Description
`Total hold downs triggered`	The total number of hold downs that occurred since the last stats clear.
`Current per-client state entries`	The current number of per client state entries in the session table.

/stats/slb/maint

SLB Maintenance Statistics

```
SLB Maintenance stats:
Maximum sessions:                516096
Current sessions:                     0
  4 second average:                   0
 64 second average:                   0
Terminated sessions:                  0
Allocation failures:                  0
TCP fragments:                        0
UDP datagrams:                        0
Non TCP/IP frames:                    0
Incorrect VIPs:                       0
Incorrect Vports:                     0
No available real server:             0
Backup server activations:            0
Overflow server activations:          0
Filtered (denied) frames:             0
Filtered max sess exceeded:           0
VMA discards:                         0
Bad buffer copies:                    0
```

SLB Maintenance statistics are described in the following table.

Table 5-33 Server Load Balancing Maintenance Statistics (/stats/slb/maint)

Statistic	Description
`Maximum sessions`	The maximum number of simultaneous sessions supported.
`Current Sessions`	Number of session bindings currently in use (the last 4 and 64 seconds).
`Terminated Sessions`	Number of sessions removed from the session table because the server assigned to them failed and graceful server failure was not enabled.
`Allocation Failures`	Indicates instances where the Switch ran out of available sessions for a port.
`TCP Fragments`	Indicates the number of TCP fragments encountered by the switch. Layer 4 processing might not handle TCP fragments, depending on configuration.
`UDP Datagrams`	Indicates that the virtual server IP address and MAC are receiving UDP frames when UDP balancing is not turned on.
`Non TCP/IP Frames`	Indicates the number of non-IP based frames received by the virtual server.
`Incorrect VIPs`	Indicates the number of times the switch received a Layer 4 request for a virtual server which was not configured.

Table 5-33 Server Load Balancing Maintenance Statistics (/stats/slb/maint)

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

/stats/slb/clear

Clearing the SLB Statistics

The following statistics are reset to zero when the clear command is given and confirmed:

Table 5-34 SLB Statistics Reset (/stats/slb/clear)

Statistics	Description
`Real server stats:`	Health check failures Total sessions Highest sessions Octets
`Real server group stats:`	Total sessions Highest sessions Octets
`Virtual server stats`	Total sessions Highest sessions Octets
`Filter stats`	Total firings
`SLB switch port stats, per port`	Real server stats: Octets, Total sessions Real server group: Octets, Total sessions Virtual server: Octets, Total sessions Total firings: Octets

Table 5-34 SLB Statistics Reset (/stats/slb/clear)

Statistics	Description
`Global SLB stats`	Per real server: DNS handoffs HTTP redirects Per server group: DNS handoffs HTTP redirects
`URL SLB and Redirection stats`	Redir: Total cache server hits Total origin server hits Total none-GETs hits Total 'Cookie: ' hits Total no-cache hits LB: ID SLB String hits
`SSL SLB stats`	Total Sessions Highest Sessions
`FTP SLB parsing and NAT stats`	Total FTP Total FTP NAT Filtered Total new active FTP NAT Index Total new FTP SLB parsing Index FTP Active FTP NAT ACK/SEQ diff FTP SLB parsing ACK/SEQ diff
`Real server stats`	Health check failures Total sessions Highest sessions Octets
`Real server group stats`	Total sessions Highest sessions Octets
`Virtual server stats`	Total sessions Highest sessions Octets

/stats/bwm
Bandwidth Management Statistics Menu

```
[Bandwidth Management Statistics Menu]
      sp      - Switch Processor Contract Stats Menu
      cont    - BW Contract stats
      rcont   - BW Contract rate stats
      hist    - BW History stats
      dump    - Dump all BWM statistics
```

Table 5-35 Bandwidth Management Statistics Menu Options (/stats/bwm)

Command Syntax and Usage
sp *<port number (1-9)>* Displays Switch processor Contract Statistics Menu. To view menu options, see page 131.
cont *<BW Contract number (1-1024)>* Displays bandwidth management contract statistics.
rcont *<BW Contract number (1-1024)>* Displays bandwidth management contract rate statistics.
hist Displays bandwidth management history statistics.
dump Displays all bandwidth management statistics.

/stats/bwm/sp <port number>

Bandwidth Management Switch Processor Statistics

```
[Bandwidth Management SP Statistics Menu]
      cont    - BW Contract stats
      rcont   - BW Contract rate stats
```

Table 5-36 Management Processor Statistics Menu Options (/stats/bwm/sp)

Command Syntax and Usage
cont *<BW Contract number (1-1024)>* Displays bandwidth management contract statistics.
rcont *<BW Contract number (1-1024)>* Displays bandwidth management contract rate statistics.

/stats/bwm/cont <contract number>

Bandwidth Management Contract Statistics

```
BW Contract statistics
Contract Name              Octets    Discards  BufUsed BufMax
-------- --------------- ---------- ---------- ------- ------
     1   ser_80           52600291          0       0  146880
     2   URLSLB                 0          0       0  146880
     5   URLCONT          136713710          0    7168  146880
     6   filter_url             0          0       0  146880
  1024   Default            1878686          0       0  146880
```

The following description of statistics applies on a specific SP (switch processor-equivalent to switch port) for all contracts irrespective of whether the contract is enabled or disabled.

Table 5-37 Bandwidth Management Contract Statistics (/stats/bwm/cont)

Statistics	Description
Contract	The contract number.
Name	The contract name.
Octets	The number of octets that are being transmitted through a particular contract since the switch is booted.
Discards	The number of octets that are being discarded because of seeing more traffic than the bandwidth contract limit permits.
BufUsed	The current amount of buffer space used to store the packets that is waiting to be transmitted.
BufMax	Maximum buffer space that can be used to store the packets before they can be transmitted. The switch starts dropping the packets of a particular contract after the maximum buffer space allocated for that contract is being occupied.

/stats/bwm/rcont

Bandwidth Management Contract Rate Statistics

Use this command to show the rate statistics of all the enabled contracts.

```
BW Contract statistics
Contract Name         Rate(Kbps)   Octets    Discards  BufUsed BufMax
-------- ---------- ----------   --------    -------- ------- ------
     1   ser_80             1773    3089507        0       0  146880
     2   URLSLB                0          0        0       0  146880
     5   URLCONT            3956    6883472        0    8128  146880
     6   filter_url            0          0        0       0  146880
  1024   Default               15     895945        0       0  146880
     1   ser_80             1767    3541927        0       0  146880
     2   URLSLB                0          0        0       0  146880
     5   URLCONT            3950    7894726        0   10048  146880
     6   filter_url            0          0        0       0  146880
  1024   Default               17     900437        0       0  146880
```

Table 5-38 Bandwidth Management Contract Rate Statistics (/stats/bwm/rcont)

Statistics	Description
Contract	The contract number.
Name	The contract name.
Rate(in Kbps)	Rate at which the packets are going out of the switch on a particular contract.
Octets	The number of octets that are being transmitted through a particular contract since the switch is booted.
Discards	The number of octets that are being discarded because of seeing more traffic than the bandwidth contract limits.
BufUsed	The current amount of buffer space used to store the packets that is waiting to be transmitted.
BufMax	Maximum buffer space that can be used to store the packets before they can be transmitted. The switch starts dropping the packets of a particular contract after the maximum buffer space allocated for that contract is being occupied.

/stats/bwm/hist

Bandwidth Management History Statistics

```
BW History statistics
Cont   Octets    Discards  TimeStamp
---- ---------- ---------- ---------
1024  108710098          0 112706:23
1024  108714358          0 112706:24
1024  108718618          0 112706:25
1024  108722594          0 112706:26
1024  108726854          0 112706:27
1024  108731114          0 112706:28
1024  108735374          0 112706:29
1024  108739700          0 112706:30
1024  108743960          0 112706:31
```

You can dump the stats kept in the SMTP history buffer that get dumped periodically when an email is sent. This command is used to keep long term history only for the contracts that are enabled and have `history` command turned on.

Use this command to show the history of all the contracts for which `history` command is enabled. The sampling is done at one-minute intervals.

Table 5-39 Bandwidth Management History Statistics (/stats/bwm/hist)

Statistics	Description
`Contract`	The contract number for which `history` is enabled.
`Octets`	The number of octets sent out on a particular contract.
`Discards`	The number of octets discarded because of seeing more traffic than the bandwidth contract limit permits.
`TimeStamp`	Indicates the time the packets were received or discarded.

/stats/mp
Management Processor Statistics

```
[MP-specific Statistics Menu]
      mem     - Show STEM memory stats
      amem    - Show All STEM memory blocks in use
      dma     - Show DMA exception counts
      pkt     - Show Packet stats
      tcb     - Show All TCP control blocks in use
      ucb     - Show All UDP control blocks in use
      uart    - Show UART counters
      cpu     - Show CPU utilization
```

Table 5-40 Management Processor Statistics Menu Options (/stats/mp)

Command Syntax and Usage
mem Displays STEM memory statistics, showing available memory.
amem Displays all STEM memory blocks in use to check for leaks.
dma Displays DMA exception counts.
pkt Displays packet statistics, to check for leads and load.
tcb Displays all TCP control blocks in use.
ucb Displays all UDP control blocks in use.
uart Displays universal asynchronous receiver/transmitter (UART) statistics.
cpu Displays CPU utilization for periods of up to 1, 4, and 64 seconds.

/stats/mp/mem

STEM Memory Statistics

```
STEM memory stats:
 allocs:          40885   frees:            40412
 alloc_fails:        38   pool_bytes:     2014800
 bytes_curr:     322624   bytes_hiwat:    1376480
 largest:       1048576
```

Table 5-41 Management Processor Memory Statistics (/stats/mp/mem)

Statistics	Description
allocs	The total number of memory allocations.
frees	The total number of bytes of free memory.
alloc_fails	The total number of memory allocation failures.
pool_bytes	The total number of bytes in the memory buffer pool.
bytes_curr	The total number of bytes that have been allocated.
bytes_hiwat	The highest number of bytes that have ever been allocated.
largest	The largest number of bytes that have ever been allocated in one allocation request.

/stats/mp/amem

All STEM Memory Statistics

```
All STEM memory allocated blocks:
Number    Caller    Blocks       Bytes
------   --------   ------  ----------
   1     0005a278       24      108544
   2     0015c0a0       34       96832
   3     000e7964        1          64
   4     000e8a04        1          32
   5     000e7a00       11         704
   6     000e8ae8       11         704
   7     000e7a40       11        1408
   8     000345e4        1          32
   9     0000abf8        1        4096
  10     0000a95c        4       16384
```

Table 5-42 All STEM memory Statistics (/stats/mp/amem)

Statistics	Description
Number	The number of the thread identifier.
Caller	The program counter of the thread.
Blocks	The total number of memory allocations that have been issued by the thread.
Bytes	The total number of bytes that have been allocated by the thread.

/stats/mp/dma

DMA Statistics

Direct Memory Access (DMA) is a capability provided by some computer bus architectures that allows data to be sent directly from an attached device (such as a disk drive) to the memory on the computer's motherboard. The microprocessor is freed from involvement with the data transfer, thus speeding up overall computer operation.

This menu option enables you to display the DMA exception counts. DMA exceptions like `RdOverruns`, `RdUnderrun` and `OvfloOvflos`, and so forth occur when data is being transferred from the device to the memory. See for more information on DMA exceptions.

```
DMA counts:
 RdOverruns:            0   RdUnderrun:           0
 WrOverruns:            0   WrUnderrun:           0
 Ovflos:              553   OvfloOvflos:          0

Mailbox Off: false
 mailbox off:      358961   mailbox on:      358961

DMA Read Off: false
 dma read off:          0   dma read on:          0
```

Table 5-43 DMA Exception Statistics (stats/mp/dma)

Statistics	Description
`RdOverruns`	The number of times the Read DMA state machine encountered an `overrun` error, size of read operation exceeded size of FIFO. It is a fatal error.
`WrOverruns`	The number of times the Write DMA state machine encountered an `overrun` error, write operation tried to write more than what was available. It is not a fatal error. Write DMA will be retried.
`RdUnderrun`	The number of times the Read DMA state machine encountered an `underrun` error, read operation read more than what was available. Fatal error.
`WrUnderrun`	The number of times unknown fatal errors occurred.
`Ovflos`	The number of frames that were placed on the software controlled DMA descriptor ring. The frame was *not* dropped.
`OvfloOvflos`	The number of frames that were dropped because there was no room on either the hardware or software DMS descriptor ring.

Table 5-43 DMA Exception Statistics (stats/mp/dma)

Statistics	Description
`mailbox off`	The number of times mailbox events were disabled to process high incoming flow of letters.
`mailbox on`	The number of times mailbox events were re-enabled.
`dma read off`	The number of times the DMA Read complete event was disabled to allow other operations.
`dma read on`	The number of times the DMA Read complete event was re-enabled.

/stats/mp/pkt

Packet Statistics

```
Packet counts:
 allocs:        2559048   frees:       2559048
 mediums:             0   jumbos:            0
 smalls:              0   failures:          0
```

Table 5-44 Packet Statistics (/stats/mp/pkt)

Statistics	Description
`allocs`	Total number of packet allocations from the packet buffer pool by the TCP/IP protocol stack.
`frees`	Total number of times the packet buffers are freed (released) to the packet buffer pool by the TCP/IP protocol stack.
`mediums`	Total number of packet allocations with size less than 256 bytes from the packet buffer pool by the TCP/IP protocol stack.
`smalls`	Total number of packet allocations with size less than 128 bytes from the packet buffer pool by the TCP/IP protocol stack.
`jumbos`	Total number of packet allocations with size larger than 255 bytes from the packet buffer pool by the TCP/IP protocol stack.
`failures`	Total number of packet allocation failures from the packet buffer pool by the TCP/IP protocol stack.

/stats/mp/tcb

TCP Statistics

```
All TCP allocated control blocks:
0064ca7c:  47.83.233.144    2082 <=> 172.25.1.11        23  established
0064d01c:  0.0.0.0             0 <=> 0.0.0.0            23  listen
```

Table 5-45 MP Specified TCP Statistics (/stats/mp/tcb)

Statistics	Description
0064ca7c/0064d01c	Memory
47.83.233.144/0.0.0.0	Destination IP address
2082/0	Destination port
172.25.1.11/0.0.0.0	Source IP
23/23	Source port
established/listen	State

/stats/mp/ucb

UCB Statistics

```
All UDP allocated control blocks:
   53:  listen
  161:  listen
  520:  listen
 1985:  listen
```

/stats/mp/uart

UART Statistics

```
UART:
 input overflows:    0
 tx hang fix:        1
 Rx discards:        0  Tx discards:        0
 X-OFFs seen:        0  X-ONs seen:         0
 X-OFFs sent:        0  X-ONs sent:         0
 Software RX FIFO discards:        0
UART Info:
 State:
  bRxEmpty 1, bTxEmpty 1, bRxXoff 0, bTxXoff 0, bTxActive 0, bRxXPend 0
  bBlockTx 0

 Buf Info:
  RX Buf Start - 005f12b0, Buf End - 005f16b0, Prod - 005f12b0, Cons -
005f12b0
  TX Buf Start - 005f16b0, Buf End - 005f1ab0, Prod - 005f194b, Cons -
005f194b
 Queue Info:
  Blk RX Thd - 0052ec80,

  TX Blk Queue - 00554224, End - 0055423c, Num Thd - 0, Num Buf Full - 0
  TOQ - 00554224, BOQ - 00554224
```

Table 5-46 UART Statistics (/stats/mp/uart)

Statistics	Description
input overflows	The number of occurrences of receive overflow.
tx hang fix	The number of occurrences of transmits hangs.
Rx discards	The number of receive characters that were dropped.
Tx discards	The number of transmit characters that were dropped.
X-OFFs seen	The number of times X-OFF signals were received.
X-ONs seen	The number of times X-ON signals were received.
X-OFFs sent	The number of times X-OFF signals were transmitted.
X-ONs sent	The number of times an X-ON signals were transmitted.

/stats/mp/cpu

CPU Statistics

This menu option enables you to display the CPU utilization statistics.

```
CPU utilization:
cpuAUtil1Second:            18%  cpuBUtil1Second:            17%
cpuAUtil4Seconds:           17%  cpuBUtil4Seconds:           17%
cpuAUtil64Seconds:          17%  cpuBUtil64Seconds:          17%
```

Table 5-47 CPU Statistics (stats/mp/cpu)

Statistics	Description
cpuAUtil1Second	The utilization of MP CPU A over 1 second. It shows the percentage.
cpuBUtil1Second	The utilization of MP CPU B over 1 second. It shows the percentage.
cpuAUtil4Seconds	The utilization of MP CPU A over 4 seconds. It shows the percentage.
cpuBUtil4Seconds	The utilization of MP CPU B over 4 seconds. It shows the percentage.
cpuAUtil64Seconds	The utilization of MP CPU A over 64 seconds. It shows the percentage.
cpuBUtil64Seconds	The utilization of MP CPU B over 64 seconds. It shows the percentage.

/stats/if <interface number>
Interface Statistics

```
IP interface 1 statistics:
ifInOctets:         48948386   ifInUcastPkts:         220553
ifInNUCastPkts:       167895   ifInDiscards:               0
ifInErrors:                0   ifInUnknownProtos:          0
ifOutOctets:        27100789   ifOutUcastPkts:        441938
ifOutNUcastPkts:      218652   ifOutDiscards:              0
ifOutErrors:               0   ifStateChanges              1
```

Table 5-48 Interface Statistics (/stats/if)

Statistics	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.
`ifInNUCastPkts`	The number of packets, delivered by this sub-layer to a higher (sub-layer), which were addressed to a multicast or broadcast address at this sub-layer. This object is deprecated in favor of `ifInMulticastPkts` and `ifInBroadcastPkts`.
`ifInDiscards`	The number of inbound packets that were chosen to be discarded even though no errors had been detected to prevent their being delivered 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 delivered 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`	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.

Table 5-48 Interface Statistics (/stats/if)

Statistics	Description
`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 to 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.
`ifOutNUcastPkts`	The total number of packets that higher-level protocols requested to be transmitted, and which were addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent. This object is deprecated in favor of `ifOutMulticastPkts` and `ifOutBroadcastPkts`.
`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 an interface has transitioned from either down to up or from up to down.

/stats/ip
IP Statistics

```
IP statistics:
ipInReceives:        3115873   ipInHdrErrors:           1
ipInAddrErrors:        35447   ipForwDatagrams:         0
ipInUnknownProtos:    500504   ipInDiscards:            0
ipInDelivers:        2334166   ipOutRequests:     1010542
ipOutDiscards:             4   ipOutNoRoutes:           4
ipReasmReqds:              0   ipReasmOKs:              0
ipReasmFails:              0   ipFragOKs:               0
ipFragFails:               0   ipFragCreates:           0
ipRoutingDiscards:         0   ipDefaultTTL:          255
ipReasmTimeout:            5
```

Table 5-49 IP Statistics (stats/ip)

Statistics	Description
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInHdrErrors	The number of input datagrams discarded due to 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, and so forth.
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 entity (the switch). This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities which 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 entity (the switch) 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 entity (the switch), 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.

Table 5-49 IP Statistics (stats/ip)

Statistics	Description
`ipInDiscards`	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for example, 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`.
`ipOutDiscards`	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (for example, for 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 which needed to be reassembled at this entity (the switch).
`ipReasmOKs`	The number of IP datagrams successfully re- assembled.
`ipReasmFails`	The number of failures detected by the IP re- assembly algorithm (for whatever reason: timed out, errors, and so forth). 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 entity (the switch).
`ipFragFails`	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity (the switch) but could not be, for example, because their `Don't Fragment` flag was set.
`ipFragCreates`	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity (the switch).

Table 5-49 IP Statistics (stats/ip)

Statistics	Description
`ipRoutingDiscards`	The number of routing entries, which were chosen to be discarded even though they are 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` (TTL) field of the IP header of datagrams originated at this entity (the switch), whenever a TTL value is not supplied by the transport layer protocol.
`ipReasmTimeout`	The maximum number of seconds, which received fragments are held while they are awaiting reassembly at this entity (the switch).

/stats/icmp
ICMP Statistics

```
ICMP statistics:
icmpInMsgs:              245802   icmpInErrors:              1393
icmpInDestUnreachs:          41   icmpInTimeExcds:              0
icmpInParmProbs:              0   icmpInSrcQuenchs:             0
icmpInRedirects:              0   icmpInEchos:                 18
icmpInEchoReps:          244350   icmpInTimestamps:             0
icmpInTimestampReps:          0   icmpInAddrMasks:              0
icmpInAddrMaskReps:           0   icmpOutMsgs:             253810
icmpOutErrors:                0   icmpOutDestUnreachs:         15
icmpOutTimeExcds:             0   icmpOutParmProbs:             0
icmpOutSrcQuenchs:            0   icmpOutRedirects:             0
icmpOutEchos:            253777   icmpOutEchoReps:             18
icmpOutTimestamps:            0   icmpOutTimestampReps:         0
icmpOutAddrMasks:             0   icmpOutAddrMaskReps:          0
```

Table 5-50 ICMP Statistics (/stats/icmp)

Statistics	Description
`icmpInMsgs`	The total number of ICMP messages which the entity (the switch) received. Note that this counter includes all those counted by `icmpInErrors`.
`icmpInErrors`	The number of ICMP messages which the entity (the switch) received but determined as having ICMP-specific errors (bad ICMP `checksums`, bad length, and so forth).
`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 (buffer almost full, stop sending data) messages received.
`icmpInRedirects`	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.

Table 5-50 ICMP Statistics (/stats/icmp)

Statistics	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 which this entity (the switch) attempted to send. Note that this counter includes all those counted by icmpOutErrors.
icmpOutErrors	The number of ICMP messages which this entity (the switch) did not send due to problems discovered within ICMP such as a lack of buffer. 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 errors that 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 (buffer almost full, stop sending data) 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.
icmpOutTimestampReps	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.

/stats/tcp
TCP Statistics

```
TCP statistics:
tcpRtoAlgorithm:          4   tcpRtoMin:                0
tcpRtoMax:           240000   tcpMaxConn:             512
tcpActiveOpens:      252214   tcpPassiveOpens:          7
tcpAttemptFails:        528   tcpEstabResets:           4
tcpInSegs:           756401   tcpOutSegs:          756655
tcpRetransSegs:           0   tcpInErrs:                0
tcpCurBuff:               0   tcpCurConn:               3
tcpOutRsts:             417
```

Table 5-51 TCP Statistics (/stats/tcp)

Statistics	Description
`tcpRtoAlgorithm`	The algorithm used to determine the `timeout` value used for retransmitting unacknowledged octets.
`tcpRtoMin`	The minimum value permitted by a TCP implementation for the retransmission `timeout`, measured in milliseconds. More refined semantics for objects of this type depend upon the algorithm used to determine the retransmission `timeout`. In particular, when the `timeout` 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 `timeout`, measured in milliseconds. More refined semantics for objects of this type depend upon the algorithm used to determine the retransmission `timeout`. In particular, when the `timeout` 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 entity (the switch) 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.
`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.

Table 5-51 TCP Statistics (/stats/tcp)

Statistics	Description
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 (for example, 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 opened.
tcpOutRsts	The number of TCP segments sent containing the RST flag.

/stats/udp
UDP Statistics

```
UDP statistics:
udpInDatagrams:        54   udpOutDatagrams:        43
udpInErrors:            0   udpNoPorts:        1578077
```

Table 5-52 UDP Statistics (/stats/udp)

Statistics	Description
udpInDatagrams	The total number of UDP datagrams delivered to the switch.
udpOutDatagrams	The total number of UDP datagrams sent from this entity (the switch).
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/snmp
SNMP Statistics

```
SNMP statistics:
snmpInPkts:                 54   snmpInBadVersions:              0
snmpInBadC'tyNames:          0   snmpInBadC'tyUses:              0
snmpInASNParseErrs:          0   snmpEnableAuthTraps:            0
snmpOutPkts:                54   snmpInBadTypes:                 0
snmpInTooBigs:               0   snmpInNoSuchNames:              0
snmpInBadValues:             0   snmpInReadOnlys:                0
snmpInGenErrs:               0   snmpInTotalReqVars:           105
snmpInTotalSetVars:          0   snmpInGetRequests:              2
snmpInGetNexts:             52   snmpInSetRequests:              0
snmpInGetResponses:          0   snmpInTraps:                    0
snmpOutTooBigs:              0   snmpOutNoSuchNames:             2
snmpOutBadValues:            0   snmpOutReadOnlys:               0
snmpOutGenErrs:              0   snmpOutGetRequests:             0
snmpOutGetNexts:             0   snmpOutSetRequests:             0
snmpOutGetResponses:        54   snmpOutTraps:                   0
```

Table 5-53 SNMP Statistics (/stats/snmp)

Statistics	Description
`snmpInPkts`	The total number of Messages delivered to the SNMP entity from the transport service.
`snmpInBadVersions`	The total number of SNMP Messages, which were delivered to the SNMP protocol entity and were for an unsupported SNMP version.
`snmpInBadC'tyNames`	The total number of SNMP Messages delivered to the SNMP entity which used an SNMP community name not known to the said entity (the switch).
`snmpInBadC'tyUses`	The total number of SNMP Messages delivered to the SNMP protocol entity which 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 entity when decoding SNMP Messages received. **Note:** OSI's method of specifying abstract objects is called ASN.1 (Abstract Syntax Notation One, defined in X.208), and one set of rules for representing such objects as strings of ones and zeros is called the BER (Basic Encoding Rules, defined in X.209). ASN.1 is a flexible notation that allows one to define a variety of data types, from simple types such as integers and bit strings to structured types such as sets and sequences. BER describes how to represent or encode values of each ASN.1 type as a string of eight-bit octets.

Table 5-53 SNMP Statistics (/stats/snmp)

Statistics	Description
snmpEnableAuth Traps	An object to enable or disable the authentication traps generated by this entity (the switch).
snmpOutPkts	The total number of SNMP Messages which were passed from the SNMP protocol entity to the transport service.
snmpInBadTypes	The total number of SNMP Messages which failed ASN parsing.
snmpInTooBigs	The total number of SNMP Protocol Data Units (PDUs) which were delivered to the SNMP protocol entity and for which the value of the error-status field is *too big*.
snmpInNoSuchNames	The total number of SNMP Protocol Data Units (PDUs) which were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
snmpInBadValues	The total number of SNMP Protocol Data Units (PDUs) which were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
snmpInReadOnlys	The total number of valid SNMP Protocol Data Units (PDUs), which were delivered to the SNMP protocol entity and for which the value of the error-status field is 'read-Only'. It should be noted that it is a protocol error to generate an SNMP PDU, which contains the value 'read-Only' 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 Protocol Data Units (PDUs), which were delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
snmpInTotalReqVars	The total number of MIB objects which have been retrieved successfully by the SNMP protocol entity as a result of receiving valid SNMP Get-Request and Get-Next Protocol Data Units (PDUs).
snmpInTotalSetVars	The total number of MIB objects, which have been altered successfully by the SNMP protocol entity as a result of receiving valid SNMP Set-Request Protocol Data Units (PDUs).
snmpInGetRequests	The total number of SNMP Get-Request Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
snmpInGetNexts	The total number of SNMP Get-Next Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
snmpInSetRequests	The total number of SNMP Set-Request Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.

Table 5-53 SNMP Statistics (/stats/snmp)

Statistics	Description
`snmpInGetResponses`	The total number of SNMP Get-Response Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
`snmpInTraps`	The total number of SNMP Trap Protocol Data Units (PDUs), which have been accepted and processed by the SNMP protocol entity.
`snmpOutTooBigs`	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status field is *too big*.
`snmpOutNoSuchNames`	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status is `noSuchName`.
`snmpOutBadValues`	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status field is `badValue.`
`snmpOutReadOnlys`	Not in use.
`snmpOutGenErrs`	The total number of SNMP Protocol Data Units (PDUs), which were generated by the SNMP protocol entity and for which the value of the error-status field is `genErr`.
`snmpOutGetRequests`	The total number of SNMP Get-Request Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
`snmpOutGetNexts`	The total number of SNMP Get-Next Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
`snmpOutSetRequests`	The total number of SNMP Set-Request Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
`snmpOutGet Responses`	The total number of SNMP Get-Response Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.
`snmpOutTraps`	The total number of SNMP Trap Protocol Data Units (PDUs), which have been generated by the SNMP protocol entity.

/stats/fdb
FDB Statistics

```
FDB statistics:
 creates:          30503   deletes:          30420
 current:             83   hiwat:              855
 lookups:         511889   lookup fails:      1126
 finds:            21801   find fails:           0
 find_or_c's:      36140   overflows:            0
```

This menu option enables you to display statistics regarding the use of the forwarding database, including the number of new entries, finds, and unsuccessful searches.

FDB statistics are described in the following table:

Table 5-54 Forwarding Database Statistics (/stats/fdb)

Statistic	Description
creates	Number of entries created in the Forwarding Database.
current	Current number of entries in the Forwarding Database.
lookups	Number of entry lookups in the Forwarding Database.
finds	Number of successful searches in the Forwarding Database.
find_or_c's	Number of entries found or created in the Forwarding Database.
deletes	Number of entries deleted from the Forwarding Database.
hiwat	Highest number of entries recorded at any given time in the Forwarding Database.
lookup fails	Number of unsuccessful searches made in the Forwarding Database.
find fails	Number of search failures in the Forwarding Database.
overflows	Number of entries overflowing the Forwarding Database.

/stats/route
Route Statistics

```
Route statistics:
ipRoutesCur:                7  ipRoutesHighWater:          7
ipRoutesMax:             1024

RIP statistics:
ripInPkts:                  0  ripOutPkts:                 0
ripBadPkts:                 0  ripRoutesAgedOut:          0

BGP statistics:
bgpInPkts:                  0  bgpOutPkts:                 0
bgpBadPkts:                 0  bgpSessFailures:           0
bgpRoutesAdded:             0  bgpRoutesRemoved:          0
bgpRoutesCur:               0  bgpRoutesFailed:           0
bgpRoutesIgnored:           0  bgpRoutesFiltered:         0
```

Table 5-55 Route Statistics (/stats/route)

Statistics	Description
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.
RIP statistics:	
ripInPkts	The total number of good RIP advertisement packets received.
ripOutPkts	The total number of RIP advertisement packets sent.
ripBadPkts	The total number of RIP advertisement packets received that were dropped.
ripRoutesAgedOut	The total number of routes learned via RIP that has aged out.
BGP statistics:	
bgpInPkts	The total number of BGP packets received.
bgpOutPkts	The total number of BGP packets sent.
bgpBadPkts	The total number of BGP packets dropped.
bgpSessFailures	The total number of failed sessions.

Table 5-55 Route Statistics (/stats/route)

Statistics	Description
bgpRoutesAdded	The total number of routes that were added to the routing table.
bgpRoutesRemoved	The total number of routes that were removed from the routing table.
bgpRoutesCur	The total number of current BGP routes.
bgpRoutesFailed	The total number of BGP routes that failed to add in the routing table.
bgpRoutesIgnored	The total number of routes ignored because the peer was not connected locally or multihop was not configured.
bgpRoutesFiltered	The total number of routes dropped by the filter.

/stats/arp
ARP Statistics

This menu option enables you to display Address Resolution Protocol statistics.

```
ARP statistics:
arpEntriesCur:              3  arpEntriesHighWater:          4
arpEntriesMax:           4096
```

Table 5-56 ARP Statistics (/stats/arp)

Statistics	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/dns
DNS Statistics

This menu option enables you to display Domain Name System statistics.

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

Table 5-57 DNS Statistics (/stats/dns)

Statistics	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/ospf
OSPF Statistics Menu

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

Table 5-58 OSPF Statistics Menu (/stats/ospf)

Command Syntax and Usage
general Displays global statistics. See page 159 for sample output.
aindex Displays area statistics.
if Displays interface statistics.

/stats/ospf/general

OSPF Global Statistics

The OSPF General Statistics contain the sum total of all OSPF packets received on all OSPF areas and interfaces.

```
Rx/Tx Stats:
Rx/Tx Stats:          Rx            Tx
                    --------      --------
  Pkts                   0             0
  hello               1443         90288
  database              66            56
  ls requests         2234             1
  ls acks                0            75
  ls updates             0          2084

Nbr change stats:                 Intf change Stats:
  hello                    4            hello            198
  start                    0            down               0
  n2way                    4            loop               0
  adjoint ok               4            unloop             0
  negotiation done         4            wait timer       197
  exchange done            4            backup             1
  bad requests             0            nbr change         5
  bad sequence             0
  loading done             4
  n1way                    0
  rst_ad                   0
  down                     1

Timers kickoff
  hello                90090
  retransmit          180180
  lsa lock                 0
  lsa ack                  0
  dbage                    0
  summary                  0
  ase export               0
```

Table 5-59 OSPF General Statistics (stats/ospf/general)

Statistics	Description
Rx/Tx Stats:	
Rx Pkts	The sum total of all OSPF packets received on all OSPF areas and interfaces.
Tx Pkts	The sum total of all OSPF packets transmitted on all OSPF areas and interfaces.
Rx Hello	The sum total of all Hello packets received on all OSPF areas and interfaces.
Tx Hello	The sum total of all Hello packets transmitted on all OSPF areas and interfaces.
Rx Database	The sum total of all Database Description packets received on all OSPF areas and interfaces.
Tx Database	The sum total of all Database Description packets transmitted on all OSPF areas and interfaces.
Rx ls Reqs	The sum total of all Link State Request packets received on all OSPF areas and interfaces.
Tx ls Reqs	The sum total of all Link State Request packets transmitted on all OSPF areas and interfaces.
Rx ls Acks	The sum total of all Link State Acknowledgement packets received on all OSPF areas and interfaces.
Tx ls Acks	The sum total of all Link State Acknowledgement packets transmitted on all OSPF areas and interfaces.
Rx ls Updates	The sum total of all Link State Update packets received on all OSPF areas and interfaces.
Tx ls Updates	The sum total of all Link State Update packets transmitted on all OSPF areas and interfaces.

Table 5-59 OSPF General Statistics (stats/ospf/general)

Statistics	Description
Nbr Change Stats:	
hello	The sum total of all Hello packets received from neighbors on all OSPF areas and interfaces.
Start	The sum total number of neighbors in this state (that is, an indication that Hello packets should now be sent to the neighbor at intervals of HelloInterval seconds.) across all OSPF areas and interfaces.
adjoint ok	The sum total number of decisions to be made (again) as to whether an adjacency should be established/maintained with the neighbor across all OSPF areas and interfaces.
negotiation done	The sum total number of neighbors in this state wherein the Master/slave relationship has been negotiated, and sequence numbers have been exchanged, across all OSPF areas and interfaces.
exchange done	The sum total number of neighbors in this state (that is, in an adjacency's final state) having transmitted a full sequence of Database Description packets, across all OSPF areas and interfaces.
bad requests	The sum total number of Link State Requests which have been received for a link state advertisement not contained in the database across all interfaces and OSPF areas.
bad sequence	The sum total number of Database Description packets which have been received that either: a) Has an unexpected DD sequence number b) Unexpectedly has the init bit set c) Has an options field differing from the last Options field received in a Database Description packet. Any of these conditions indicate that some error has occurred during adjacency establishment for all OSPF areas and interfaces.
loading done	The sum total number of link state updates received for all out-of-date portions of the database across all OSPF areas and interfaces.
n1way	The sum total number of Hello packets received from neighbors, in which this router is not mentioned across all OSPF interfaces and areas.
rst_ad	The sum total number of times the Neighbor adjacency has been reset across all OPSF areas and interfaces.
down	The total number of Neighboring routers down (that is, in the initial state of a neighbor conversation.) across all OSPF areas and interfaces.

Table 5-59 OSPF General Statistics (stats/ospf/general)

Statistics	Description
Intf Change Stats:	
hello	The sum total number of Hello packets sent on all interfaces and areas.
down	The sum total number of interfaces down in all OSPF areas.
loop	The sum total of interfaces no longer connected to the attached network across all OSPF areas and interfaces.
unloop	The sum total number of interfaces, connected to the attached network in all OSPF areas.
wait timer	The sum total number of times the Wait Timer has been fired, indicating the end of the waiting period that is required before electing a (Backup) Designated Router across all OSPF areas and interfaces.
backup	The sum total number of Backup Designated Routers on the attached network for all OSPF areas and interfaces.
nbr change	The sum total number of changes in the set of bidirectional neighbors associated with any interface across all OSPF areas.
Timers Kickoff:	
hello	The sum total number of times the Hello timer has been fired (which triggers the send of a Hello packet) across all OPSF areas and interfaces.
retransmit	The sum total number of times the Retransmit timer has been fired across all OPSF areas and interfaces.
lsa lock	The sum total number of times the Link State Advertisement (LSA) lock timer has been fired across all OSPF areas and interfaces.
lsa ack	The sum total number of times the LSA Ack timer has been fired across all OSPF areas and interfaces.
dbage	The total number of times the data base age (Dbage) has been fired.
summary	The total number of times the Summary timer has been fired.
ase export	The total number of times the Autonomous System Export (ASE) timer has been fired.

/stats/vrrp
VRRP Statistics

Virtual Router Redundancy Protocol (VRRP) support on Alteon Web switches provides redundancy between routers in a LAN. This 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`)

The statistics for the VRRP LAN are displayed:

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

Table 5-60 VRRP Statistics (/stats/vrrp)

Statistics	Description
`vrrpInAdvers`	The total number of VRRP advertisements that have been received.
`vrrpBadAdvers`	The total number of VRRP advertisements received that were dropped.
`vrrpOutAdvers`	The total number of VRRP advertisements that have been sent.

/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 to tune or debug 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. Many of the commands, although not new, display more or different information than in the previous version. Important difference are called out in the text.

To make finding information easier, the menu options under the Server Load Balancing Menu (`/cfg/slb`) are in Chapter 7.

/cfg Configuration Menu

```
[Configuration Menu]
      sys     - System-wide Parameter Menu
      port    - Port Menu
      ip      - IP Menu
      vlan    - VLAN Menu
      stp     - Spanning Tree Menu
      snmp    - SNMP Menu
      pmirr   - Port Mirroring Menu
      slb     - Server Load Balancing Menu
      trunk   - Trunk Group Menu
      vrrp    - Virtual Router Redundancy Protocol Menu
      bwm     - Bandwidth Management Menu
      setup   - Step by step configuration set up
      dump    - Dump current configuration to script file
      ptcfg   - Backup current configuration to tftp server
      gtcfg   - Restore current configuration from tftp server
```

Table 6-1 Configuration Menu Options (/cfg)

Command Syntax and Usage
sys Displays the System Configuration Menu. To view menu options, see page 169.
port *<port number (1-9)>* Displays the Port Configuration Menu. To view menu options, see page 177.
ip Displays the IP Configuration Menu. To view menu options, see page 181.
vlan *<VLAN number (1-4094)>* Displays the VLAN Configuration Menu. To view menu options, see page 204.
stp Displays the Spanning Tree Configuration Menu. To view menu options, see page 206.
snmp Displays the SNMP Configuration Menu. To view menu options, see page 210.
pmirr Displays the Mirroring Configuration Menu. To view menu options, see page 212.
slb Displays the Server Load Balancing Configuration Menu. To view menu options, see Chapter 7, "The SLB Configuration Menu".
trunk *<group number (1-4)>* Displays the Trunk Group Configuration Menu. To view menu options, see page 214.
vrrp Displays the Virtual Router Redundancy Configuration Menu. To view menu options, see page 215.
bwm Displays the Bandwidth Management Configuration Menu. To view menu options, see page 228.
setup Step-by-step configuration set-up of the switch. For details, see page 232.
dump Dumps current configuration to a script file. For details, see page 232.
ptcfg *<host name or IP address>* *<filename on host>* Backs up current configuration to TFTP server. For details, see page 233.
gtcfg *<host name or IP address>* *<filename on host>* Restores current configuration from TFTP server. For details, see page 233.

Viewing, Applying, and Saving Changes

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

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

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

Viewing Pending Changes

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

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

Applying Pending Changes

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

```
# apply
```

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

NOTE – All configuration changes take effect immediately when applied, except for starting Spanning Tree Protocol. To turn STP on or off, you must apply the changes, save them (see below), and then reset the switch (see “Resetting the Web Switch” on page 299).

Saving the Configuration

In addition to applying the configuration changes, you can save them to flash memory on the Web 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 n
```

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 299.

/cfg/sys
System Configuration

```
[System Menu]
      syslog  - Syslog Menu
      sshd    - SSH Server Menu
      radius  - RADIUS Authentication Menu
      ntp     - NTP Server Menu
      date    - Set system date
      time    - Set system time
      idle    - Set timeout for idle CLI sessions
      snmp    - Set SNMP access control
      wport   - Set Web server port number
      notice  - Set login notice
      bannr   - Set login banner
      mnet    - Set management network
      mmask   - Set management netmask
      smtp    - Set SMTP host
      tnet    - Enable/disable Telnet access
      tnport  - Set Telnet server port number
      hprompt - Enable/disable display hostname (sysName) in CLI prompt
      bootp   - Enable/disable use of BOOTP
      http    - Enable/disable HTTP (Web) access
      user    - User Access Control Menu (passwords)
      cur     - Display current system-wide parameters
```

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 Menu. To view menu options, see page 171.
sshd Displays the SSH Server Menu. To view menu options, see page 172.
radius Displays the RADIUS Authentication Menu. To view menu options, see page 173.
ntp Displays the Network Time Protocol (NTP) Server Menu. To view menu options, see page 174.
date Prompts the user for the system date.

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

Command Syntax and Usage
time Configures the system time using a 24-hour clock format.
idle *<idle timeout in minutes; affects both console and Telnet>* Sets the idle timeout for CLI sessions, from 1 to 60 minutes. The default is 5 minutes.
snmp Disables or provides read-only/write-read SNMP access.
wport *<TCP port number (1-65535)>* Sets the switch port used for serving switch Web content. The default is HTTP port 80. If Global Server Load Balancing is to be used, set this to a different port (such as 8080).
notice Displays login notice immediately before the “Enter password:” prompt. This notice can contain up to 1024 characters and new lines.
bannr *<string, maximum 80 characters>* Configures a login banner of up to 80 characters. When a user or administrator logs into the switch, the login banner is displayed. It is also displayed as part of the output from the `/info/sys` command.
mnet *<IP subnet (such as 192.4.17.0)>* Sets the base source IP address that allows access to switch management through Telnet, SNMP, RIP, or the Web OS browser-based 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 access to switch management functions. Specify the mask in dotted-decimal notation.
smtp *<SMTP host name or IP address>* Sets the Simple Mail Transfer Protocol (SMTP) host, which is used for sending bandwidth management history information.
tnet Enables or disables telnet access.This command is disabled by default. You will see this command only if you are connected to the switch through the console.
tnport Sets an optional telnet server port number for cases where the server listens for telnet sessions on a non-standard port.
bootp disable|enable Enables or disables the use of BOOTP. If you enable BOOTP, the switch will query its BOOTP server for all of the switch IP parameters. This command is disabled by default.

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

Command Syntax and Usage
http disable|enable Enables or disables HTTP (Web) access to the browser-based interface. It is disabled by default.
user Displays the User Access Control Menu. To view menu options, see page 175.
cur Displays the current system parameters.

/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
      sever   - Set the severity of first syslog host
      sever2  - Set the severity of second syslog host
      console - Enable/disable console output of syslog messages
      log     - Enable/disable sysloggging of features
      cur     - Display current syslog settings
```

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

Command Syntax and Usage
host Sets the IP address of the first syslog host.
host2 Sets the IP address of the second syslog host.
sever *<0–7>* This option sets the severity level of the first syslog host displayed. The default is 7, which means log all the seven severity levels.
sever2 *<0–7>* This option sets the severity level of the second syslog host displayed. The default is 7, which means, log all the seven severity levels.
console disable|enable Enables or disables delivering syslog messages to the console. When necessary, disabling `console` ensures the switch is not affected by syslog messages. It is enabled by default.

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

Command Syntax and Usage
log Displays a list of features for which syslog messages can be generated. You can choose to enable/disable specific features (such as vlans, gslb, filter), or enable/disable syslog on all available features.
cur Displays the current syslog settings.

/cfg/sys/sshd

SSH Server Configuration

```
[SSHD Menu]
      sshport - Set SSH server port number
      ena     - Enable the SCP apply and save
      on      - Turn SSH server ON
      cur     - Display current SSH server configuration
```

For Alteon AD3, AD4, 180e and 184 Web switches, this menu enables Secure Shell access from any SSH client. SSH scripts can be viewed by using the `/cfg/dump` command (see page 232).

NOTE – Except for `cur`, the commands of this menu are only accessible through the console port.

Table 6-4 System Configuration Menu Options (/cfg/sys/sshd)

Command Syntax and Usage
sshport *<port number>* Sets the SSH server port number.
ena Enables the SCP apply and save.
on Enables the SSH server.
cur Displays the current SSH server configuration.

/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  - Enable or disable RADIUS backdoor for telnet
      on      - Turn RADIUS authentication ON
      off     - Turn RADIUS authentication OFF
      cur     - Display current RADIUS configuration
```

Table 6-5 System Configuration Menu Options (/cfg/sys/radius)

Command Syntax and Usage
prisrv *<IP address>* Sets the primary RADIUS server address.
secsrv *<IP address>* Sets the secondary RADIUS server address.
secret *<1-32 character secret>* This is the shared secret between the switch and the RADIUS server(s).
port *<RADIUS port>* Enter the number of the UDP port, between 1500 - 3000. The default is 1645.
retries *<RADIUS server retries (1-3)>* Sets the number of failed authentication requests before switching to a different RADIUS server. The default is 3 requests.
timeout *<RADIUS server timeout 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 Enables or disables the RADIUS backdoor for telnet. `telnet` also applies to SSH/SCP connections.
on Enables the RADIUS server.
off Disables the RADIUS server.
cur Displays the current RADIUS server parameters.

/cfg/sys/ntp

NTP Server Configuration

```
[NTP Server Menu]
      server  - Set NTP server address
      intrval - Set NTP server resync interval
      tzone   - Set NTP timezone offset from GMT
      dlight  - Enable or disable NTP daylight savings time
      on      - Turn NTP service ON
      off     - Turn NTP service OFF
      cur     - Display current NTP configuration
```

This menu enables you to synchronize the switch clock to a Network Time Protocol (NTP) server. By default, this option is disabled.

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

Command Syntax and Usage
server *<IP address>* Specifies the IP address of the NTP server to which you want to synchronize the switch clock.
intrval *<resync interval in minutes>* Specifies the interval, that is, how often, in minutes (1-2880), to re synchronize the switch clock with the NTP server.
tzone *<timezone offset, in hours>* Specifies the timezone offset, in hours, of the switch you are synchronizing from Greenwich Mean Time (GMT).
dlight disable|enable Disables or enables daylight savings time in the system clock. When enabled, the switch will add an extra hour to the system clock so that it is consistent with the local clock. By default, this option is disabled.
on Enables the NTP synchronization service.
off Disables the NTP synchronization service.
cur Displays the current NTP service settings.

/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)
      cur     - Display current user statistics
```

NOTE – Passwords can be a maximum of 15 characters.

Table 6-7 User Access Control Menu Options (/cfg/sys/user)

Command Syntax and Usage
usrpw Sets the user (`user`) password. 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.
sopw Sets the SLB operator (`slboper`) password. The SLB operator manages Web servers and other Internet services and their loads. He or she can view all switch information and statistics and can enable/disable servers using the Server Load Balancing configuration menus. Access includes "`user`" functions.
l4opw Sets the Layer 4 operator (`l4oper`) password. The Layer 4 operator manages traffic on the lines leading to the shared Internet services. He or she can view all switch information and statistics. Access includes "`slboper`" functions.
opw Sets the operator (`oper`) password. The operator password can have a maximum of 15 characters. The operator manages all functions of the switch. He or she can view all switch information and statistics and can reset ports or the entire switch. Access includes "`l4oper`" functions.

Table 6-7 User Access Control Menu Options (/cfg/sys/user)

Command Syntax and Usage
sapw Sets the SLB administrator (`slbadmin`) password. Administrator who configures and manages Web servers and other Internet services and their loads. He or she can view all switch information and statistics, but can configure changes only on the Server Load Balancing menus. Note that the Filter Menu options are not accessible to the SLB administrator. Access includes "`l4oper`" functions.
l4apw Sets the Layer 4 administrator (`l4admin`) password. The Layer 4 administrator configures and manages traffic on the lines leading to the shared Internet services. He or she can view all switch information and statistics and can configure parameters on the Server Load Balancing menus, with the exception of not being able to configure filters. Access includes "`slbadmin`" functions.
admpw Sets the administrator (`admin`) password. The superuser administrator has complete access to all menus, information, and configuration commands on the Web switch, including the ability to change both the user and administrator passwords. Access includes "`oper`" and "`l4admin`" functions.
cur Displays the current user status.

/cfg/port <port number>
Port Configuration

```
[Port 1 Menu]
      fast    - Fast Phy Menu
      gig     - Gig Phy Menu
      pref    - Set preferred phy
      back    - Set backup phy
      pvid    - Set default port VLAN id
      name    - Set port name
      cont    - Set default port BW Contract
      rmon    - Enable/Disable RMON for port
      tag     - Enable or disable VLAN tagging for port
      iponly  - Enable or disable allowing only IP related frames
      ena     - Enable port
      dis     - Disable port
      cur     - Display current port configuration
```

The Port Menu enables you to configure settings for individual switch ports (1-9). This command is enabled by default.

Table 6-8 Port Configuration Menu Options (/cfg/port)

Command Syntax and Usage
fast If a port is configured to support Fast Ethernet, this option displays the Fast Ethernet Physical Link Menu. To view menu options, see page 179.
gig If a port is configured to support Gigabit Ethernet, this option displays the Gigabit Ethernet Physical Link Menu. To view menu options, see page 179.
pref If dual physical connectors are available on the port, this option defines the preferred physical connector. Choices are: ■ Fast Ethernet Port, RJ-45 connector ■ Gigabit Ethernet Port, SC fiber connector (default)
back If dual physical connectors are available on the port, this option defines the physical connector to use when the preferred choice fails or is unavailable. Choices are: ■ Fast Ethernet Port, RJ-45 connector (default) ■ Gigabit Ethernet Port, SC fiber connector ■ None (By default)

Table 6-8 Port Configuration Menu Options (/cfg/port)

Command Syntax and Usage
pvid Sets the default VLAN number which will be used to forward frames which are not VLAN tagged. The default number is 1.
name Sets a name for the port. The assigned port name appears next to the port number on some information and statistics screens. The default is set to None.
cont *<BWM Contract (1-256)>* Sets the default Bandwidth Management Contract for this port.
rmon disable|enable Disables or enables RMON for this port. It is disabled by default.
tag disable|enable Disables or enables VLAN tagging for this port. It is disabled by default.
iponly disable|enable Disables or enables allowing only IP-related frames. It is disabled by default.
ena Enables the port.
dis Disables the port. (To temporarily disable a port without changing its configuration attributes, refer to "Temporarily Disabling a Port" on page 180.)
cur Displays current port parameters.

/cfg/port <port number> fast|gig

Port Link Configuration

```
[Fast Link Menu]
      speed   - Set link speed
      mode    - Set full or half duplex mode
      fctl    - Set flow control
      auto    - Set autonegotiation
      cur     - Display current fast link configuration
```

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

NOTE – Since the `speed` and `mode` parameters cannot be set for Gigabit Ethernet ports, these options do not appear on the Gigabit Link Menu.

Link menu options are described in Table 6-9 and appear on the `fast` and `gig` port configuration menus for the Alteon Web switches. Using these configuration menus, you can set port parameters such as speed, flow control, and negotiation mode for the port link.

Table 6-9 Port Link Configuration Menu Options (/cfg/port <*number*> fast|gig)

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 full|half|any Sets the operating mode. This command is available only in the Fast Link Menu.The choices include: ■ "Any," for autonegotiation (default) ■ Full-duplex ■ Half-duplex
fctl rx|tx|both|none Sets the flow control. This command is available only in the Fast Link Menu.The choices include: ■ Autonegotiation (default) ■ Receive flow control ■ Transmit flow control ■ Both receive and transmit flow control ■ No flow control

Table 6-9 Port Link Configuration Menu Options (/cfg/port <*number*> fast|gig)

Command Syntax and Usage
auto on|off Enables or disables autonegotiation for the port.
cur Displays current port parameters.

Temporarily Disabling a Port

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

```
Main# /oper/port <port number>/dis
```

Because this configuration sets a temporary state for the port, you do not need to use `apply` or `save`. The port state will revert to its original configuration when the Web switch is reset. See the "Operations Menu" on page 287 for other operations-level commands.

/cfg/ip
IP Configuration

```
[IP Menu]
      if      - Interface Menu
      gw      - Default Gateway Menu
      route   - Static Route Menu
      frwd    - Forwarding Menu
      rip1    - Routing Information Protocol Menu
      ospf    - Open Shortest Path First (OSPF) menu
      bgp     - Border Gateway Protocol Menu
      port    - IP Port Menu
      dns     - Domain Name System Menu
      bootp   - Bootstrap Protocol Relay Menu
      rearp   - Set re-ARP period in minutes
      metrc   - Set default gateway metric
      cur     - Display current IP configuration
```

Table 6-10 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage
if *<interface number (1-256)>* Displays the IP Interface Menu. To view menu options, see page 183.
gw *<default gateway number (1-4)>* Displays the IP Default Gateway Menu. To view menu options, see page 184.
route Displays the IP Static Route Menu. To view menu options, see page 185.
frwd Displays the IP Forwarding Menu. To view menu options, see page 186.
rip1 Displays the Routing Interface Protocol version 1 Menu. To view menu options, see page 188.
ospf Displays the OSPF Menu. To view menu options, see page 189.
bgp Displays the Border Gateway Protocol Menu. To view menu options, see page 197.
port *<port number (1-9)>* Displays the IP Port Menu. To view menu options, see page 201.

Table 6-10 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage
dns Displays the IP Domain Name System Menu. To view menu options, see page 201.
bootp Displays the Bootstrap Protocol Menu. To view menu options, see page 202.
rearp *<2-120 minutes>* Sets the re-ARP period in minutes. The switch periodically sends ARP (Address Resolution Protocol) requests to refresh its address database. This command is used for setting the interval between ARP refreshes of the next IP address in the database. The default interval is 10 minutes.
metrc strict|roundrobin Sets the default gateway metric for `strict` or `roundrobin`. The default gateway metric is `strict`.
cur Displays the current IP configuration.

/cfg/ip/if <interface number>

IP Interface Configuration

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

The Web switch can be configured with up to 256 IP interfaces. Each IP interface represents the Web switch on an IP subnet on your network. The Interface option is disabled by default.

Table 6-11 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 255.255.255.0)>* Configures the IP subnet address mask for the interface using dotted decimal notation.
broad *<broadcast address (such as 192.4.17.255)>* Configures the IP broadcast address for the interface using dotted decimal notation.
vlan *<VLAN number>* Configures the VLAN number for this interface. Each interface can belong to one VLAN, though any VLAN can have multiple IP interfaces in it.
relay Enables or disables the BOOTP relay on this interface. It is enabled by default.
ena Enables this IP interface.
dis Disables this IP interface.
del Removes this IP interface.
cur Displays the current interface settings.

/cfg/ip/gw <gateway number>

Default IP Gateway Configuration

```
[Default gateway 1 Menu]
      addr     - Set IP address
      intr     - Set interval between ping attempts
      retry    - Set number of failed attempts to declare gateway DOWN
      arp      - Enable/disable ARP only health checks
      vlan     - Set VLAN number
      ena      - Enable default gateway
      dis      - Disable default gateway
      del      - Delete default gateway
      cur      - Display current default gateway configuration
```

NOTE – The switch can be configured with up to 250 gateways. Gateways one to four are reserved for default gateway load balancing. Gateways five to 250 are used for load-balancing of VLAN-based gateways.

This option is disabled by default.

Table 6-12 Default Gateway Options (/cfg/ip/gw)

Command Syntax and Usage
addr *<default gateway address>* Configures the IP address of the default IP gateway using dotted decimal notation.
intr *<value (0-60 seconds)>* The switch pings the default gateway to verify that it's up. The `intr` option sets the time between health checks. The range is from 1 to 120 seconds. The default is 2 seconds.
retry *<attempts (1-120)>* Sets the number of failed health check attempts required before declaring this default gateway inoperative. The range is from 1 to 120 attempts. The default is 8 attempts.
arp Enables or disables Address Resolution Protocol (ARP) health checks. This command is disabled by default.
vlan *<VLAN number (1-4094)>* Sets the VLAN to be assigned to this default IP gateway.
ena Enables the gateway for use.

Table 6-12 Default Gateway Options (/cfg/ip/gw)

Command Syntax and Usage
dis Disables the gateway.
del Deletes the gateway from the configuration.
cur Displays the current gateway settings.

Default Gateway Metrics

For information about configuring which gateway is selected when multiple default gateways are enabled, see page 203.

/cfg/ip/route

IP Static Route Configuration

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

Up to 128 static routes can be configured.

Table 6-13 IP Static Route Configuration Menu Options (cfg/ip/route)

Command Syntax and Usage
add *<destination> <mask> <gateway> <interface number>* Adds a static route. You will be prompted to enter a destination IP address, destination subnet mask, and gateway address. Enter all addresses using dotted decimal notation.
rem *<destination>* Removes a static route. The destination address of the route to remove must be specified using dotted decimal notation.
cur Displays the current IP static routes.

/cfg/ip/frwd

IP Forwarding Configuration

```
[IP Forwarding Menu]
      local   - Local network definition for route caching menu
      dirbr   - Enable or disable forwarding directed broadcasts
      on      - Globally turn IP Forwarding ON
      off     - Globally turn IP Forwarding OFF
      cur     - Display current IP Forwarding configuration
```

Table 6-14 IP Forwarding Configuration Menu Options (/cfg/ip/frwd)

Command Syntax and Usage
local Displays the menu used to define local network for route caching. Up to five local networks (lnets) can be configured. To view menu options, see page 186.
dirbr Enables or disables forwarding directed broadcasts. This command is disabled by default.
on Enables IP forwarding (routing) on the Web switch.
off Disables IP forwarding (routing) on the Web switch. Forwarding is turned off by default.
cur Displays the current IP forwarding settings.

/cfg/ip/frwd/local

Local Network Route Caching Definition

```
[IP Local Networks Menu]
      add     - Add local network definition
      rem     - Remove local network definition
      cur     - Display current local network definitions
```

This menu is used for adding local networks by setting the local network address and netmask for the route cache, and to remove local networks.

Table 2 IP Local Networks Menu Options (/cfg/ip/frwd/local)

Command Syntax and Usage
add *<Local Network Address> <Local Network Mask>* Adds a definition for a local network. For details, see "Defining IP Address Ranges for the Local Route Cache" below.
rem Removes a definition for a local network.
cur Displays the current local network definitions.

Defining IP Address Ranges for the Local Route Cache

The Local Route Cache lets you use switch resources more efficiently, by reducing the size of the ARP table on the Web switch. The `/cfg/ip/frwd/local/add` parameters define a range of addresses that will be cached on the Web switch. The local network address is used to define the base IP address in the range which will be cached, and the local network mask is the mask which is applied to produce the range. To determine if a route should be added to the memory cache, the destination address is masked (bitwise AND) with the local network mask and checked against the local network address.

By default, the local network address and mask are both set to 0.0.0.0. This produces a range that includes all Internet addresses for route caching: 0.0.0.0 through 255.255.255.255.

Addresses to be cached are subnets that are directly connected and for which there is an interface configured on the Web switch. To limit the route cache to your local hosts, you could configure the parameters as shown in the examples in the following table.

Table 6-15 Local Routing Cache Address Ranges

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

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

/cfg/ip/rip1

Routing Information Protocol Configuration

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

The RIP1 Menu is used for configuring Routing Information Protocol, version 1 (RIP1) parameters. This option is turned off by default.

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

Table 6-16 Routing Information Protocol Menu (/cfg/ip/rip1)

Command Syntax and Usage
updat *<update period (1-120 seconds)>* Sets the RIP update period in seconds. It is set at 30 seconds by default.
spply disable|enable This command is disabled by default. When enabled, the switch supplies routes to other routers.
lsten disable|enable This command is disabled by default. When enabled, the switch learns routes from other routers.
deflt none|lsten|spply|both When enabled, the switch accepts RIP default routes from other routers and gives them priority over configured default gateways. When disabled, the switch rejects RIP default routes. This command is disabled by default.
statc This command is disabled by default. When enabled, the switch supplies static routes.
poisn disable|enable This command is disabled by default. When enabled, the switch uses split horizon with poisoned reverse. When disabled, the switch uses only split horizon.

Table 6-16 Routing Information Protocol Menu (/cfg/ip/rip1)

Command Syntax and Usage
vip disable|enable Enables or disables the advertisement of virtual IP addresses as Host Routes. If a VIP route exists in a routing table, it will always be advertised except when it is included in another network route that is already being advertised. **Note:** If all real servers behind a VIP go down, the route gets removed from the routing table, and will not be advertised. If we disable all the real servers using operation command, the VIP route does not get eliminated from the routing table, and the switch will continue to advertise the route.
on Globally turns RIP ON.
off Globally turns RIP OFF.
cur Displays the current RIP configuration.

/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
      default - Originate default route information
      rtrid   - Set OSPF router ID
      on      - Globally turn OSPF ON
      off     - Globally turn OSPF OFF
      current - Display current OSPF configuration
```

Table 6-17 OSPF Configuration Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
aindex *<area index>* Displays the area index menu. This area index does not represent the actual OSPF area number. See page 190 to view menu options.
range *<range number (1-16)>* Displays summary routes menu for up to 64 IP addresses. See page 192 to view menu options.

Table 6-17 OSPF Configuration Menu Options (/cfg/ip/ospf)

Command Syntax and Usage
if *<interface number>* Displays the OSPF interface configuration menu. See page 193 to view menu options.
virt Displays the Virtual Links menu used to configure OSPF for a Virtual Link. See page 195 to view menu options.
host Displays the menu for configuring OSPF for the host routes. Up to 128 host routes can be configured. Host routes are used for advertising network device IP addresses to external networks to perform server load balancing within OSPF. It also makes Area Border Route (ABR) load sharing and ABR failover possible. See page 196 to view menu options.
default *<metric value>* *<metric type (**1**|**2**)>* Sets one default route among multiple choices in an area.
rtrid *<IP address>* Sets OSPF router ID, which is in IP address format. It is required only when configuring virtual links on the Web switch.
on Enables OSPF on the Web switch.
off Disables OSPF on the Web switch.
current Displays the current OSPF configuration settings.

/cfg/ip/ospf/aindex

Area Index Configuration Menu

```
[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 6-18 Area Index Configuration Menu Options (/cfg/ip/ospf/aindex)

Command Syntax and Usage
areaid *<n.n.n.n>* Defines the OSPF area number.
type Defines the type of the area. For example, when a virtual link has to be established with the backbone, the area type must be defined as transit.
metric *<metric value> <metric type (1 or 2)>* Configures a stub area to send a numeric metric value. All routes received via that stub area carry the configured metric to potentially influence routing decisions. Metric value assigns the priority for choosing the switch for default route. Metric type determines the method for influencing routing decisions for external routes.
auth Authenticates passwords so that only trusted routing devices can participate.
spf Sets time interval between two successive SPF calculations of the shortest path tree using the Dijkstra's algorithm.
enable Enables the OSPF area.
disable Disables the OSPF area.
delete Deletes the OSPF area.
current Displays the current OSPF configuration.

/cfg/ip/ospf/range

OSPF Summary Range Configuration Menu

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

Table 6-19 OSPF Summary Range Configuration Menu Options (/cfg/ip/ospf/range)

Command Syntax and Usage
addr *<IP Address>* Displays the base IP address for the range.
mask *<mask>* Displays the IP address mask for the range.
aindex *<area index>* Displays the area index used by the Web switch.
hide Hides the OSPF summary range.
enable Enables the OSPF summary range.
disable Disables the OSPF summary range.
delete Deletes the OSPF summary range.
current Displays the current OSPF summary range.

/cfg/ip/ospf/if

OSPF Interface Configuration Menu

```
[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
```

Table 6-20 OSPF Interface Configuration Menu Options (/cfg/ip/ospf/if)

Command Syntax and Usage
aindex *<area index>* Displays the OSPF area index.
prio *<priority value (0-127)>* Displays the assigned priority value to the Web switch ‘s OSPF intefaces. (A priority value of 127 is the highest and 1 is the lowest. A priority value of 0 specifies that the interface cannot be used as Designated Router (DR) or Backup Designated Router (BDR).)
cost *<cost value (1-65535)>* Displays cost set for the selected path—preferred or backup. Usually the cost is inversely proportional to the bandwidth of the interface. Low cost indicates high bandwidth.
hello *<1-65535>* Displays the authentication parameters of a `hello` packet, which is set to be in an interval of seconds.
dead *<1-65535>* Displays the health parameters of a `hello` packet, which is set to be in an interval of seconds.
trans *<1-65535>* Displays the transit delay in seconds.
retra *<1-65535>* Displays the retransmit interval in seconds.
key *<password>* Displays the authentication key.

Table 6-20 OSPF Interface Configuration Menu Options (/cfg/ip/ospf/if)

Command Syntax and Usage
enable Enables OSPF interface.
disable Disables OSPF interface.
delete Deletes OSPF interface.
current Displays the current settings for OSPF interface.

/cfg/ip/ospf/virt

OSPF Virtual Link Configuration Menu

```
[OSPF Virtual Link 1  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
```

Table 6-21 OSPF Virtual Link Configuration Menu Options (/cfg/ip/ospf/virt)

Command Syntax and Usage
aindex *<area index>* Displays the OSPF area index.
hello Displays the authentication parameters of a hello packet, which is set to be in an interval of seconds.
dead Displays the health parameters of a hello packet, which is set to be in an interval of seconds.
trans Displays the delay in transit in seconds.
retra Displays the retransmit interval in seconds.
nbr *<router ID (IP address)>* Displays the router ID of the virtual neighbor.
key *<password>* Displays the password (up to eight characters) for each virtual link.
enable Enables OSPF virtual link.
disable Disables OSPF virtual link.

Table 6-21 OSPF Virtual Link Configuration Menu Options (/cfg/ip/ospf/virt)

Command Syntax and Usage
delete Deletes OSPF virtual link.
current Displays the current OSPF virtual link settings.

/cfg/ip/ospf/host

OSPF Host Entry Configuration Menu

```
[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
```

Table 6-22 OSPF Host Entry Configuration Menu Options (/cfg/ip/ospf/host)

Command Syntax and Usage
addr *<IP address>* Displays the base IP address for the host entry.
aindex *<area index>* Displays the area index of the host.
cost *<cost value (1-65535)>* Displays the cost value of the host.
enable Enables OSPF host entry.
disable Disables OSPF host entry.
delete Deletes OSPF host entry.
current Displays the current OSPF host entries.

/cfg/ip/bgp

Border Gateway Protocol Configuration

```
[Border Gateway Protocol Menu]
      peer    - Peer Menu
      filt    - Filter Menu
      on      - Globally turn BGP ON
      off     - Globally turn BGP OFF
      cur     - Display current BGP configuration
```

Border Gateway Protocol (BGP) is an Internet protocol that enables routers on a network to share routing information with each other and advertise information about the segments of the IP address space they can access within their network with routers on external networks. BGP allows you to decide what is the “best” route for a packet to take from your network to a destination on another network, rather than simply setting a default route from your border router(s) to your upstream provider(s). BGP is defined in RFC 1771.

The BGP Menu enables you to configure the switch to receive routes and to advertise static routes, fixed routes and virtual server IP addresses with other internal and external routers. In the current Web OS implementation, the Web switch does not advertise BGP routes that are learned from other BGP “speakers”.

The BGP menu option is turned off by default.

NOTE – Fixed routes are subnet routes. There is one fixed route per IP interface.

When multiple peers advertise the same route, we use the route with the shortest AS path as the preferred route if eBGP or use the local preference if iBGP.

Table 6-23 Border Gateway Protocol Menu (/cfg/ip/bgp)

Command Syntax and Usage
peer *<peer number (1-4)>* Displays the menu used to configure each BGP *peer*. Each border router, within an autonomous system, exchanges routing information with routers on other external networks. To view menu options, see page 198.
filt *<filter number (1-4)>* Displays the menu used to configure the range of IP destinations accepted by each BGP peer filter. To view menu options, see page 200.
on Globally turns BGP on.

Table 6-23 Border Gateway Protocol Menu (/cfg/ip/bgp)

Command Syntax and Usage
off Globally turns BGP off.
cur Displays the current BGP configuration.

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

BGP Peer Configuration

```
[BGP Peer 1 Menu]
      addr    - Set remote IP address
      ras     - Set remote autonomous system number
      if      - Set local IP interface
      las     - Set local autonomous system number
      hold    - Set hold time
      ttl     - Set time-to-live of IP datagrams
      metric  - Set metric of advertised routes
      fixed   - Enable or disable advertising fixed routes
      static  - Enable or disable advertising static routes
      vip     - Enable or disable advertising VIP routes
      ena     - Enable peer
      dis     - Disable peer
      del     - Delete peer
      cur     - Display current peer configuration
```

This menu is used to configure BGP peers, which are border routers that exchange routing information with routers on internal and external networks. The peer option is disabled by default.

Table 6-24 BGP Peer Configuration Options (/cfg/ip/bgp/peer)

Command Syntax and Usage
addr *<IP address (such as 192.4.17.101)>* Defines the IP address for the specified peer (border router), using dotted decimal notation. The default address is 0.0.0.0.
ras *<AS number (0-65535)>* Sets the remote autonomous system number for the specified peer.
if *<interface number (1-256)>* Selects a switch IP interface (between 1 and 256) for the specified peer. The default value is 1.

Table 6-24 BGP Peer Configuration Options (/cfg/ip/bgp/peer)

Command Syntax and Usage
las *<AS number (0-65535)>* Sets the local autonomous system number for the specified peer. It is set at 0 by default.
hold *<hold time (0-65535)>* Sets the period of time, in seconds, that will elapse before the peer session is torn down because the switch hasn't received a "keep alive" message from the peer. It is set at 90 seconds by default.
ttl *<number of router hops (1-255)>* Specifies the number of router hops that the IP datagram can make. This value is used to restrict the number of "hops" the advertisement makes. It is also used to support multi-hops, which allow BGP peers to talk across a routed network. The default number is set at 1.
metric <*metric (1-255)***>** Sets the length of the AS path used when advertising eBGP routes. When advertising iBGP routes, this parameter sets the local preference. The default number is set at 1.
fixed disable|enable Enables or disables advertising fixed routes. This command is disabled by default.
static disable|enable Enables or disables advertising static routes. No default route is advertised. This command is disabled by default.
vip disable|enable Enables or disables advertising virtual server routes.
ena Enables this peer configuration.
dis Disables this peer configuration.
del Deletes this peer configuration.
cur Displays the current BGP peer configuration.

/cfg/ip/bgp/filt <filter number>

BGP Filter Configuration

```
[BGP Filter 1 Menu]
      addr    - Set filter address
      mask    - Set filter mask
      ena     - Enable filter
      dis     - Disable filter
      del     - Delete filter
      cur     - Display current filter configuration
```

This menu enables you to configure filters that specify the routes/range of IP destinations a peer router will accept from other peers. A route must match a filter to be installed in the routing table. By default, the first filter is enabled and the rest of the filters are disabled.

Table 6-25 BGP Filter Configuration Options (/cfg/ip/bgp/filt)

Command Syntax and Usage
addr *<IP address (such as 192.4.17.101)>* Defines the starting IP address for this filter, using dotted decimal notation. The default address is 0.0.0.0.
mask *<IP address>* This IP address mask is used with `addr` to define the range of IP addresses that will be accepted by the peer when the filter is enabled. The default address is 0.0.0.0.
ena Enables this BGP filter.
dis Disables this BGP filter.
del Deletes this BGP filter.
cur Displays the current BGP filter configuration.

/cfg/ip/port <port number>

IP Port Configuration

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

The IP Port Menu allows you to turn IP forwarding on or off on a port-by-port basis. By default, the port forwarding option is turned on.

Table 6-26 IP Forwarding Port Options (/cfg/ip/port)

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

/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
      cur     - 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 6-27 Domain Name Service Menu Options (/cfg/ip/dns)

Command Syntax and Usage
prima *<IP address (such as 192.4.17.101)>* You will be prompted to set the IP address for your primary DNS server. Use dotted decimal notation.
secon *<IP address (such as 192.4.17.101)>* You will be prompted to set the IP address for your secondary DNS server. If the primary DNS server fails, the configured secondary will be used instead. Enter the IP address using dotted decimal notation.
dname *<dotted DNS notation>*|**none** Sets the default domain name used by the switch. For example: `mycompany.com`
cur Displays the current Domain Name System 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
      cur     - Display current BOOTP relay configuration
```

The Bootstrap Protocol (BOOTP) Relay Menu is used to allow hosts to obtain their configurations from a Dynamic Host Configuration Protocol (DHCP) server. The BOOTP configuration enables the switch to forward a client request for an IP address to two DHCP/BOOTP servers with IP addresses that have been configured on the Web switch.

BOOTP relay menu is turned off by default.

Table 6-28 Bootstrap Protocol Relay Configuration Menu Options (/cfg/ip/bootp)

Command Syntax and Usage
addr *<IP address>* Sets the IP address of the BOOTP server.
addr2 *<IP address>* Sets the IP address of the second BOOTP server.
on Globally turns on BOOTP relay.
off Globally turns off BOOTP relay.
cur Displays the current BOOTP relay configuration.

/cfg/ip/metrc *<metric name>*

Default Gateway Metrics

If multiple default gateways are configured and enabled, a metric can be set to determine which primary gateway is selected. There are two metrics, which are described in the table "Default Gateway Metrics (/cfg/ip/metrc)" on page 203.

Table 6-29 Default Gateway Metrics (/cfg/ip/metrc)

Option	Description
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 which point the next in line will take over as the default IP gateway.
roundrobin	This provides basic gateway load balancing. The switch sends each new gateway request to the next healthy, enabled gateway in line. All gateway requests to the same destination IP address are resolved to the same gateway.

/cfg/vlan <VLAN number>
VLAN Configuration

```
[VLAN 1 Menu]
      name    - Set VLAN name
      stg     - Assign VLAN to a Spanning Tree Group
      cont    - Set BW contract
      add     - Add port to VLAN
      rem     - Remove port from VLAN
      def     - Define VLAN as list of ports
      jumbo   - Enable/disable Jumbo Frame support
      ena     - Enable VLAN
      dis     - Disable VLAN
      del     - Delete VLAN
      cur     - 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 30.

By default, the VLAN menu option is disabled except VLAN 1, which is enabled all the time.

Table 6-30 VLAN Configuration Menu Options (/cfg/vlan)

Command Syntax and Usage
name *<name to be assigned to the VLAN, maximum 32 characters>* Assigns a name to the VLAN or changes the existing name. The default VLAN name is the first one.
stg Assigns a VLAN to a Spanning Tree Group.
cont *<BWM Contract (1-1024)>* Sets the Bandwidth Management contract for this VLAN. The default contract number is 256 on AD3 and 1024 on AD4.
add *<port number>* Adds port(s) or trunk group(s) to the VLAN membership.
remove *<port number>* Removes port(s) or trunk group(s) from this VLAN.
def Defines which ports are members of this VLAN. Every port must be a member of at least one VLAN. By default, it defines ports between 1-9 for VLAN 1.

Table 6-30 VLAN Configuration Menu Options (/cfg/vlan)

Command Syntax and Usage
jumbo disable|enable Enables or disables support for Jumbo Frame support on this VLAN. This feature is disabled by default.
ena Enables this VLAN.
dis Disables this VLAN without removing it from the configuration.
del Deletes this VLAN.
cur Displays the current VLAN configuration.

NOTE – All ports must belong to at least one VLAN. Any port which is removed from a VLAN and which is not a member of any other VLAN is automatically added to default VLAN #1. You cannot remove a port from VLAN #1 if the port has no membership in any other VLAN.
Also, you cannot add a port to more than one VLAN unless the port has VLAN tagging turned on (see the `tag` command on page 178).

/cfg/stp
Spanning Tree Configuration

```
[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
      cur     - Display current bridge parameters
```

Web OS supports the IEEE 802.1d Spanning Tree Protocol (STP). STP is used to prevent loops in the network topology. There are 16 spanning tree groups that can be configured on the Switch. This command is turned on by default.

NOTE – When VRRP is used for active/active redundancy, STP must be enabled.

Table 6-31 Spanning Tree Configuration Menu (/cfg/stp)

Command Syntax and Usage
brg Displays the Bridge Spanning Tree Menu. To view menu options, see page 207.
port *<port number (1-9)>* Displays the Spanning Tree Port Menu. To view menu options, see page 209.
add *<VLAN number (1-4094)>* Associates a VLAN with a spanning tree and requires an external VLAN ID as a parameter.
remove *<VLAN number (1-4094)>* Breaks the association between a VLAN and a spanning tree and requires an external VLAN ID as a parameter.
clear Removes all VLANs from a spanning tree.
on Globally enables Spanning Tree Protocol.

Table 6-31 Spanning Tree Configuration Menu (/cfg/stp)

Command Syntax and Usage
off Globally disables Spanning Tree Protocol.
default Restores a spanning tree instance to its default configuration.
cur Displays current Spanning Tree Protocol parameters.

/cfg/stp/brg

Bridge Spanning Tree Configuration

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

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

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

Table 6-32 Bridge Spanning Tree Menu Options (/cfg/stp/brg)

Command Syntax and Usage
prior *<new 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 *<new bridge hello time (1-10 secs)>* Configures the bridge hello time.The hello time specifies how often the root bridge transmits a configuration bridge protocol data unit (BPDU). Any bridge that is not the root bridge uses the root bridge hello value. The range is 1 to 10 seconds, and the default is 2 seconds.
mxage *<new bridge max age (6-40 secs)>* Configures the bridge maximum age. The maximum age parameter specifies the maximum time the bridge waits without receiving a configuration bridge protocol data unit before it reconfigures the STP network. The range is 6 to 40 seconds, and the default is 20 seconds.
frwd *<new bridge Forward Delay (4-30 secs)>* Configures the bridge forward delay parameter. The forward delay parameter specifies the amount of time that a bridge port has to wait before it changes from the listening state to the learning state and from the learning state to the forwarding state. The range is 4 to 30 seconds, and the default is 15 seconds.
aging *<new bridge Aging Time (1-65535 secs, 0 to disable)>* Configures the forwarding database aging time. The aging time specifies the amount of time the bridge waits without receiving a packet from a station before removing the station from the forwarding database. The range is 1 to 65535 seconds, and the default is 300 seconds. To disable aging, set this parameter to 0.
current Displays the current bridge STP parameters.

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

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

/cfg/stp/port <port number>

Spanning Tree Port Configuration

```
[Spanning Tree Port 1 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
      cur     - 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

The **port** option of STP is turned on by default.

Table 6-33 Spanning Tree Port Menu (/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 100Mbps ports, and 1 for Gigabit ports. A value of 0 indicates that the default cost will be computed for an auto negotiated link speed.
on Enables STP on the port.
off Disables STP on the port.
cur Displays the current STP port parameters.

/cfg/snmp
SNMP Configuration

```
[SNMP Menu]
      name    - Set SNMP "sysName"
      locn    - Set SNMP "sysLocation"
      cont    - Set SNMP "sysContact"
      rcomm   - Set SNMP read community string
      wcomm   - Set SNMP write community string
      trap1   - Set first SNMP trap host address
      trap2   - Set second SNMP trap host address
      t1comm  - Set community string for first trap host
      t2comm  - Set community string for second trap host
      auth    - Enable or disable SNMP "sysAuthenTrap"
      linkt   - Enable or disable SNMP link up/down trap
      cur     - Display current SNMP configuration
```

The Web OS 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 6-34 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 name of the system location. The location can have a maximum of 64 characters.
cont *<new string, maximum 64 characters>* Configures the name of the system contact. The 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. The default read community string is *public*.
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. The default write community string is *private*.
trap1 *<new SNMP trap host IP address>* Configures the IP address of the first SNMP trap host using dotted decimal notation. The SNMP trap host is the device that receives SNMP trap messages from the switch.
trap2 *<new SNMP trap host IP address>* Configures the IP address of the second SNMP trap host using dotted decimal notation.
t1com *<new trap host community string, maximum 32 characters>* Configures the community string for the first trap host. The default community string for the first trap host is *public*.
t2com *<new trap host community string, maximum 32 characters>* Configures the community string for the second trap host. The default community string for the second trap host is *public*.
auth disable|enable Enables or disables the use of the system authentication trap facility. The default setting is disabled.
linkt *<port>* [**disable|enable**] Enables or disables the sending of SNMP link up and link down traps. The default setting is enabled.
cur Displays the current STP port parameters.

/cfg/pmirr
Port Mirroring Menu

```
[Port Mirroring Menu]
      mirror  - Enable/Disable Mirroring
      monport - Monitoring Port based PM Menu
      cur     - Display All Mirrored and Monitoring Ports
```

Port mirroring is disabled by default.

NOTE – Port mirroring menu options are supported only on the Alteon AD4 and Alteon 184 Web switches.

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

Table 6-35 Port Mirroring menu options (/cfg/pmirr)

Command Syntax and Usage
mirror enable|disable Enables or disables port mirroring
monport *<port number (1-9)>* Displays port-mirroring menu. To view menu options, see page 213.
cur Displays current settings of the mirrored and monitoring ports.

/cfg/pmirr/monport

Port-Mirroring Menu

```
[Port 1 Menu]
      add     - Add "Mirrored" port
      rem     - Rem "Mirrored" port
      cur    - Display current Port-based Port Mirroring configuration
```

Table 6-36 Port-Based Port-Mirroring Menu Options (/cfg//pmirr/monport)

Command Syntax and Usage
add *<mirrored port> <direction (in, out, or both)>* Adds the port to be mirrored. This command also allows you to enter the direction of the traffic. It is necessary to specify the direction because: If the source port of the frame matches the mirrored port and the mirrored direction is ingress or both (ingress and egress), the frame is sent to the mirrored port. If the destination port of the frame matches the mirrored port and the mirrored direction is egress or both, the frame is sent to the monitoring port.
rem Removes the mirrored port.
cur Displays the current settings of the monitoring port.

/cfg/trunk <trunk group number>
Trunk Configuration

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

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

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

By default, the trunk group is empty and disabled.

Table 6-37 Trunk Configuration Menu Options (/cfg/trunk)

Command Syntax and Usage
cont *<BWM Contract (1-1024)>* Sets the default Bandwidth Management Contract for this trunk group. By default, the contract number is 256 for AD3 and 1024 for AD4.
add *<port number (1-9)>* Adds a physical port to the current trunk group.
rem *<port number (1-9)>* Removes a physical port from the current trunk group.
ena Enables the current trunk group.
dis Turns the current trunk group off.
del Removes the current trunk group configuration.
cur Displays current trunk group parameters.

/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 or disable hot-standby processing
      on       - Globally turn VRRP ON
      off      - Globally turn VRRP OFF
      cur      - Display current VRRP configuration
```

Virtual Router Redundancy Protocol (VRRP) support on Alteon Web switches provides redundancy between routers in a LAN. This 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.

By default, VRRP is disabled. Alteon WebSystems has extended VRRP to include virtual servers as well, allowing for full active/active redundancy between its Layer 4 switches. For more information on VRRP, see the "High Availability" chapter in the *Web OS 10.0 Application Guide*.

Table 6-38 Virtual Router Redundancy Protocol Options (/cfg/vrrp)

Command Syntax and Usage
vr *<virtual router number (1-256>* Displays the VRRP Virtual Router Menu. This menu is used for configuring up to 256 virtual routers on this switch. To view menu options, see page 216.
group Displays the VRRP virtual router group menu, used to combine all virtual routers together as one logical entity. Group options must be configured when using two or more Alteon switches in a hot-standby failover configuration where only one switch is active at any given time. To view menu options, see page 221.
if *<interface number (1-256)>* Displays the VRRP Virtual Router Interface Menu. To view menu options, see page 225.

Table 6-38 Virtual Router Redundancy Protocol Options (/cfg/vrrp)

Command Syntax and Usage
track *<interface number (1-256)>* 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 menu options, see page 226.
hotstan disable|enable Enables or disables hot standby processing, in which two or more switches provide redundancy for each other. By default, this option is disabled.
on Globally enables VRRP on this switch.
off Globally disables VRRP on this switch.
cur Displays the current VRRP parameters.

/cfg/vrrp/vr *<router number>*

Virtual Router Configuration

```
[VRRP Virtual Router 1 Menu]
      track   - Priority Tracking Menu
      vrid    - Set virtual router ID
      addr    - Set IP address
      if      - Set interface number
      prio    - Set renter priority
      adver   - Set advertisement interval
      preem   - Enable or disable preemption
      share   - Enable or disable sharing
      ena     - Enable virtual router
      dis     - Disable virtual router
      del     - Delete virtual router
      cur     - Display current VRRP virtual router configuration
```

This menu is used for configuring up to 256 virtual routers for this switch. A virtual router is defined by its virtual router ID and an 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.

Virtual routers are disabled by default.

Table 6-39 VRRP Virtual Router Options (/cfg/vrrp/vr)

Command Syntax and Usage
track Displays the VRRP Priority Tracking Menu for this virtual router. Tracking is an Alteon Web-Systems proprietary extension to VRRP, used for modifying the standard priority system used for electing the master router. Tracking is not needed if sharing (`share`) is enabled. To view menu options, see page 219.
vrid *<virtual router ID (1-255)>* Defines the virtual router ID. This is used in conjunction with `addr` (below) to define a virtual router on this switch. To create a pool of VRRP-enabled routing devices which can provide redundancy to each other, each participating VRRP device must be configured with the same virtual router: one that shares the same `vrid` and `addr` 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. The default value is 1. All `vrid` values must be unique within the VLAN to which the virtual router's IP interface belongs.
addr *<IP address>* Defines the IP address for this virtual router using dotted decimal notation. This is used in conjunction with the `vrid` (above) to configure the same virtual router on each participating VRRP device. The default address is 0.0.0.0.
if *<interface number (1-256)>* Selects a switch IP interface (between 1 and 256). If the IP interface has the same 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 which has assumed master routing authority. This preemption occurs even if the `preem` option below is disabled. The default value is 1.
prio *<priority (1-254)>* Defines the election priority bias for this virtual server. This can be any integer between 1 and 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 (highest). 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. This can be any integer between 1 and 255 seconds. The default value is 1.

Table 6-39 VRRP Virtual Router Options (/cfg/vrrp/vr)

Command Syntax and Usage
preem disable|enable 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. Note that 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 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.
ena Enables this virtual router.
dis Disables this virtual router.
del Deletes this virtual router from the switch configuration.
cur Displays the current configuration information for this virtual router.

/cfg/vrrp/vr <router number>/track

Virtual Router Priority Tracking Configuration

```
[VRRP Virtual Router 1 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
      hsrv    - Enable/disable tracking HSRP by VLAN
      cur     - 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 226).

Criteria are tracked dynamically, continuously updating virtual router priority levels when enabled. If the virtual router preemption option (see `preem` in Table 6-39 on page 217) is enabled, this virtual router can assume master routing authority when its priority level rises above that of the current master.

Some tracking criteria (`vrs`, `ifs`, and `ports` below) apply to standard virtual routers, otherwise called "virtual interface routers." Other tracking criteria (`l4pts`, `reals`, and `hsrp`) apply to "virtual server routers," which perform Layer 4 Server Load Balancing functions. A virtual *server* router is defined as any virtual router whose IP address (`addr`) is the same as any configured virtual server IP address.

Table 6-40 VRRP Priority Tracking Options (/cfg/vrrp/vr #/track)

Command Syntax and Usage
vrs disable|enable 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 making sure that traffic for any particular client/server pairing are handled by the same switch, increasing routing and load balancing efficiency. This command is disabled by default.
ifs disable|enable 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.

Table 6-40 VRRP Priority Tracking Options (/cfg/vrrp/vr #/track)

Command Syntax and Usage
ports disable|enable When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. A port is considered "active" if it has a link and is forwarding traffic. This helps elect the virtual routers with the most available ports as the master. This command is disabled by default.
l4pts disable|enable 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 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 *<priority (1-254)>* Hot Standby Router Protocol (HSRP) is used with some types of routers for establishing router failover. In networks where 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.
hsrv disable|enable Hot Standby Router on VLAN (HSRV) is used to work in VLAN-tagged environments. Enable this switch option to increment only that **vrrp** instance that is on the *same* VLAN as the tagged hsrp master flagged packet. This command is disabled by default.
cur 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 or disable preemption
      share   - Enable or disable sharing
      ena     - Enable virtual router
      dis     - Disable virtual router
      del     - Delete virtual router
      cur     - Display current VRRP virtual router configuration
```

The Virtual Router Group menu is used for associating all virtual routers into a single logical virtual router, which forces all virtual routers on the Web switch to either be master or backup as a group. A virtual router is defined by its virtual router ID and an 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.

NOTE – This option is required to be configured only when using at least two Alteon Web switches in a hot-standby failover configuration, where only one switch is active at any time.

Table 6-41 VRRP Virtual Router Group Options (/cfg/vrrp/group)

Command Syntax and Usage
track Displays the VRRP Priority Tracking Menu for the virtual router group. Tracking is an Alteon WebSystems proprietary extension to VRRP, used for modifying the standard priority system used for electing the master router. Tracking is not needed if sharing (`share`) is enabled. To view menu options, see page 226.
vrid *<virtual router ID (1-255)>* Defines the virtual router ID. 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 (see `if` below) belongs. The default virtual router ID is 1.
if *<interface number (1-256)>* Selects a switch IP interface (between 1 and 256). The default switch IP interface number is 1.

Table 6-41 VRRP Virtual Router Group Options (/cfg/vrrp/group)

Command Syntax and Usage
prio *<priority (1-254)>* Defines the election priority bias for this virtual router group. This can be any integer between 1 and 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 (highest). 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 *<1-255 seconds>* Defines the time interval between VRRP master advertisements. This can be any integer between 1 and 255 seconds. The default is 1.
preem disable|enable Enables or disables master preemption. When enabled, if the virtual router group 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. Note that 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 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.
ena Enables the virtual router group.
dis Disables the virtual router group.
del Deletes the virtual router group from the switch configuration.
cur Displays the current configuration information for the virtual router group.

/cfg/vrrp/group/track

Virtual Router Group Priority Tracking Configuration

```
[Virtual Router Group 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
      hsrv    - Enable/disable tracking HSRP by VLAN
      cur     - Display current VRRP Group Tracking configuration
```

NOTE – If *Virtual Router Group Tracking* is enabled, then the tracking option will be available only under *group* option. The tracking setting for the other individual virtual routers will be ignored.

Table 6-42 Virtual Router Group Priority Tracking Options (/cfg/vr/group/track)

Command Syntax and Usage
vrs disable|enable 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 making sure that traffic for any particular client/server pairing are handled by the same switch, increasing routing and load balancing efficiency. This command is disabled by default.
ifs disable|enable 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 When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. A port is considered “active” if it has a link and is forwarding traffic. This helps elect the virtual routers with the most available ports as the master. This command is disabled by default.
l4pts disable|enable 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 6-42 Virtual Router Group Priority Tracking Options (/cfg/vr/group/track)

Command Syntax and Usage
reals disable|enable When enabled for virtual server routers, the priority for this virtual router will be increased for each healthy real server. 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 Enables Hot Standby Router Protocol (HSRP) for this virtual router group. HSRP is used with some types of routers for establishing router failover. In networks where 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. This helps elect the switch closest to the master HSRP router as the master, optimizing routing efficiency. This command is disabled by default.
hsrv disable|enable Hot Standby Router on VLAN (HSRV) is used to work in VLAN-tagged environments. Enable this switch option to increment only that **vrrp** instance that is on the *same* VLAN as the tagged hsrp master flagged packet. This command is disabled by default.
cur Displays the current configuration for priority tracking for this virtual router.

/cfg/vrrp/if <interface number>

VRRP Interface Configuration

NOTE – The *interface-number* (1 to 256) represents the IP interface on which authentication parameters must be configured.

```
[VRRP Interface 1 Menu]
      auth    - Set authentication types
      passw   - Set plain-text password
      del     - Delete interface
      cur     - Display current VRRP interface configuration
```

This menu is used for configuring VRRP authentication parameters for the IP interfaces used with the virtual routers.

Table 6-43 VRRP Interface Menu Options (/cfg/vrrp/if)

Command Syntax and Usage
auth none|password Defines the type of authentication that will be used: `none` (no authentication), or `password` (password authentication).
passw *<key>* 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).
del Clears the authentication configuration parameters for this IP interface. The IP interface itself is not deleted.
cur 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
      hsrv    - Set priority increment for HSRP by VLAN tracking
      cur     - 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 (see "VRRP Virtual Router Priority Tracking Menu" on page 219), the priority level for the virtual router is increased by an amount defined through this menu.

Table 6-44 VRRP Tracking Options (/cfg/vrrp/track)

Command Syntax and Usage
vrs *<0-254>* Defines the priority increment value (1 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 (1 through 254) for active IP interfaces detected on this switch. The default value is 2.
ports *<0-254>* Defines the priority increment value (1 through 254) for active ports on the virtual router's VLAN. The default value is 2.
l4pts *<0-254>* Defines the priority increment value (1 through 254) for physical switch ports with active Layer 4 processing. The default value is 2.
reals *<0-254>* Defines the priority increment value (1 through 254) for healthy real servers behind the virtual server router. The default value is 2.
hsrp *<0-254>* Defines the priority increment value (1 through 254) for switch ports with Layer 4 client-only processing that receive HSRP broadcasts. The default value is 10.

Table 6-44 VRRP Tracking Options (/cfg/vrrp/track)

Command Syntax and Usage
hsrv *<0-254>* Defines the priority increment value (1 through 254) for vrrp instances that are on the same VLAN. The default value is 10.
cur 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 219) are enabled.

/cfg/bwm
Bandwidth Management Configuration

```
[Bandwidth Management Menu]
      cont    - Contract Menu
      policy  - Policy Menu
      user    - Set SMTP server user name
      force   - Enable or disable enforce policies
      on      - Globally turn Bandwidth Management processing ON
      off     - Globally turn Bandwidth Management processing OFF
      cur     - Display current Bandwidth Management configuration
```

By default, BWM is turned off.

NOTE – Up to 1024 bandwidth management contracts can be configured on the Alteon AD4 and Alteon 184 Web switches. The Alteon AD3 and Alteon 180e Web switches do not support Bandwidth Management.

Table 6-45 Bandwidth Management Options (/cfg/bwm)

Command Syntax and Usage
cont *<contract number (1-1024)>* Displays the Bandwidth Management Contract Menu. By default, this option is disabled. To view menu options, see page 229.
policy *<policy number (1-64)>* Displays the Bandwidth Management Policy Menu. To view menu options, see page 230.
user *<user name>* Sets the SMTP user name to whom the history statistics will be mailed. The default is set to None.
force disable|enable Enables or disables enforces policies. When disabled, no bandwidth limits will be applied on queues. By default, this option is enabled.
on Globally enables Bandwidth Management on this switch.
off Globally disables Bandwidth Management on this switch.
cur Displays the current Bandwidth Management configuration.

/cfg/bwm/cont <contract number>

Bandwidth Management Contract Configuration

```
[BW Contract 1 Menu]
      name    - Set Contract name
      policy  - Set Contract Policy
      prec    - Set Contract Precedence
      history - Enable/disable Saving Contract stats history
      wtos    - Enable/disable overwriting IP TOS for this Contract
      ena     - Enable BW Contract
      dis     - Disable BW Contract
      del     - Delete BW Contract
      cur     - Display current BW Contract configuration
```

Table 6-46 Bandwidth Management Policy Menu Options (/cfg/bwm/cont)

Command Syntax and Usage
name *<15 character name>* Sets the name for this Bandwidth Management contract.
policy *<bandwidth policy number (1-64)>* Sets the policy number for this Bandwidth Management contract. The default policy number is 64.
prec *<bandwidth precedence value (1-255)>* Sets the precedence value for this Bandwidth Management contract. The default precedence value is 1.
history disable|enable Disables or enables saving statistics for this contract on the server. By default, it is enabled.
wtos disable|enable Disables or enables overwriting the IP Type of Service (TOS) for this contract. By default, it is disabled.
ena Enables this Bandwidth Management contract.
dis Disables this Bandwidth Management contract.
del Removes this contract from the switch.
cur Displays the current Bandwidth Management contract configuration.

/cfg/bwm/pol <policy number>

Bandwidth Management Policy Configuration

```
[Policy 1 Menu]
      hard    - Set hard Limit
      soft    - Set soft Limit
      resv    - Set Reservation Limit
      utos    - Set underlimit (soft limit) TOS
      otos    - Set overlimit (soft limit) TOS
      buffer  - Set Buffer Limit
      del     - Delete BW Policy
      cur     - Display current Policy configuration
```

Table 6-47 Bandwidth Management Policy Menu Options (/cfg/bwm/pol)

Command Syntax and Usage
hard *<250K-5000K|1M-1000M>* Sets the hard bandwidth limit for this policy. This is the highest amount of bandwidth available to this policy. The default value is 2000 kbps.
soft *<250K-5000K|1M-1000M>* Sets the soft bandwidth limit for this policy. The default value is 1000 kbps.
resv *< 250K-5000K|1M-1000M>* Sets the reserve limit for this policy. This is the amount of bandwidth always available to this policy. The default value is 500Kbytes.
utos *<BW Policy TOS (0-255)>* Sets the new `utos` value to overwrite the original TOS value if the traffic for this contract is under the soft limit. With this option set to the default value of "0," the switch will not overwrite the TOS value.
otos *<BW Policy TOS (0-255)>* Sets the new `otos` value to overwrite the original TOS value if the traffic for this contract is over the soft limit. With this option set to the default value of "0," the switch will not overwrite the TOS value.
buffer *<Maximum buffer space (bytes) (8192-512000)>* Sets the buffer limit for this policy. The default value is 8192 bytes.
del Deletes the bandwidth management policy.
cur Displays the current bandwidth policy configuration.

/cfg/bwm/cur

Bandwidth Management Current Configuration

```
Current Bandwidth Management setting: ON
  Policy Enforcement: enabled
  SMTP server user name:

Contract Name             Policy Prec Hist TOS State
     4                       4      4   E   D    E
     6                       4      6   E   D    E
   256   Default             --     0   E   D    E
*Default contract gets all the BW that is available on
 a port after the active contracts reserved BW is taken.

Policy     Hard  Soft  Resv oTOS uTOS Buffer
    1     2000k 1000k  500k    0    0   32640
    2       40m   35m   30m    0    0   32640
    3     2000k 1000k  500k    0    0   32640
    4       10m    9m    8m    0    0   32640
    5     2000k 1000k  500k    0    0   32640
    6     2000k 1000k  500k    0    0   32640
    7     2000k 1000k  500k    0    0   32640
    8     2000k 1000k  500k    0    0   32640
    9     2000k 1000k  500k    0    0   32640
   10     2000k 1000k  500k    0    0   32640
   11     2000k 1000k  500k    0    0   32640
   12     2000k 1000k  500k    0    0   32640
   13     2000k 1000k  500k    0    0   32640
   14     2000k 1000k  500k    0    0   32640
   15     2000k 1000k  500k    0    0   32640
   16     2000k 1000k  500k    0    0   32640
   17     2000k 1000k  500k    0    0   32640
   18     2000k 1000k  500k    0    0   32640
   19     2000k 1000k  500k    0    0   32640
   20     2000k 1000k  500k    0    0   32640
   21     2000k 1000k  500k    0    0   32640
   22     2000k 1000k  500k    0    0   32640
   23     2000k 1000k  500k    0    0   32640
   24     2000k 1000k  500k    0    0   32640
   25     2000k 1000k  500k    0    0   32640
   26     2000k 1000k  500k    0    0   32640
   27     2000k 1000k  500k    0    0   32640
   28     2000k 1000k  500k    0    0   32640
   29     2000k 1000k  500k    0    0   32640
   30     2000k 1000k  500k    0    0   3264
```

/cfg/setup
Setup

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

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

```
Configuration# setup
```

For a complete description of how to use setup, see Chapter 2, "First-Time Configuration."

/cfg/dump
Dump

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

```
Configuration# dump
```

The configuration is displayed with parameters that have been changed from the default values. The screen display can be captured, edited, and placed in a script file, which can be used to configure other switches through a Telnet connection. When using Telnet to configure a new switch, paste the configuration commands from the script file at the command line prompt of the switch. The active configuration can also be saved or loaded via TFTP, as described on page 233.

/cfg/ptcfg <TFTP server> <filename>
Saving the Active Switch Configuration

When the ptcfg command is used, the switch's active configuration commands (as displayed using /cfg/dump) will be uploaded to the specified script configuration file on the TFTP server. To start the switch configuration upload, at the Configuration# prompt, enter:

```
Configuration# ptcfg <TFTP server> <filename>
```

Where *server* is the TFTP server IP address or hostname, and *filename* is the name of the target script configuration file.

NOTE – The output file is formatted with line-breaks but no carriage returns—the file cannot be viewed with editors that require carriage returns (such as Microsoft Notepad).

NOTE – If the TFTP server is running SunOS or the Solaris operating system, the specified ptcfg file must exist prior to executing the ptcfg 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 configuration data.

/cfg/gtcfg <TFTP server> <filename>
Restoring the Active Switch Configuration

When the gtcfg command is used, the active configuration will be replaced with the commands found in the specified configuration file. The file can contain a full switch configuration or a partial switch configuration. The configuration loaded using gtcfg is not activated until the apply command is used. If the apply command is found in the configuration script file loaded using this command, the apply action will be performed automatically.

To start the switch configuration download, at the Configuration# prompt, enter:

```
Configuration# gtcfg <TFTP server> <filename>
```

Where *server* is the TFTP server IP address or hostname, and *filename* is the name of the target script configuration file.

CHAPTER 7

The SLB Configuration Menu

This chapter discusses how to use the Command Line Interface (CLI) for configuring Server Load Balancing (SLB) on the Web switch.

/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
      layer7  - Layer 7 Resource Definition Menu
      wap     - WAP Menu
      sync    - Config Synch Menu
      adv     - Layer 4 Advanced Menu
      on      - Globally turn Layer 4 processing ON
      off     - Globally turn Layer 4 processing OFF
      cur     - Display current Layer 4 configuration
```

Table 7-1 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage
real *<real server number (1-1023)>* Displays the menu for configuring real servers. To view menu options, see page 237.
group *<real server group number (1-256)>* Displays the menu for placing real servers into real server groups. To view menu options, see page 242.
virt *<virtual server number (1-256)>* Displays the menu for defining virtual servers. To view menu options, see page 246.

Table 7-1 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage
`filt` *<filter ID (1-2048)>* Displays the menu for Filtering and Application Redirection. To view menu options, see page 254.
`port` *<port number (1-9)>* Displays the menu for setting physical switch port states for Layer 4 activity. To view menu options, see page 266.
`gslb` Displays the menu for configuring Global Server Load Balancing. To view menu options, see page 268.
`layer7` Displays later 7 Resource Definition Menu. To view menu options, see page 274.
`wap` Displays WAP Menu. To view menu options, see page 277
`sync` Displays the Synch Peer Switch Menu. To view menu options, see page 278.
`adv` Displays the Layer 4 Advanced Menu. To view menu options, see page 280.
`on` Globally turns on Layer 4 software services for Server Load Balancing and Application Redirection. This option can be performed only after the optional Layer 4 software is enabled (see "Activating Optional Software on page 293). Enabling Layer 4 services is not necessary for using filters only to allow, deny, or NAT traffic.
`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
`cur` Displays the current Server Load Balancing configuration.

Filtering and Layer 4 (Server Load Balancing)

Filters configured to allow, deny, or perform Network Address Translation (NAT) on 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.

/cfg/slb/real <server number>
Real Server SLB Configuration

```
[Real server 1  Menu]
   layer7  - Real Server Layer 7 Command Menu
   rip     - Set IP addr of real server
   name    - Set server name
   weight  - Set server weight
   maxcon  - Set maximum number of connections
   tmout   - Set minutes inactive connection remains open
   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
   addport - Add real port to server
   remport - Remove real port to server
   remote  - Enable/disable remote site operation
   proxy   - Enable/disable client proxy operation
   submac  - Enable/disable source MAC address substitution
   ena     - Enable real server
   dis     - Disable real server
   del     - Delete real server
   cur     - Display current real server configuration
```

This menu is used for configuring information about real servers that participate in a server pool for Server Load Balancing or Application Redirection. The required parameters are:

- Real server IP address
- Real server enabled (disabled by default)

Table 7-2 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
layer7 Displays the Layer 7 Menu. To view menu options, see page 241.
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| none>* 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.
weight *<server weight (1-48)>* Sets the weighting value (1 to 48) that this real server will be given in the load balancing algorithms. Higher weighting values force the server to receive more connections than the other servers configured in the same real server group. By default, each real server is given a weight setting of 1. A setting of 10 would assign the server roughly 10 times the number of connections as a server with a weight of 1. Weights are not applied when using the hash or minmisses metrics (see "Server Load Balancing Metrics" on page 245).
maxcon *<maximum connections (0-200000)>* Sets the maximum number of connections that this server should simultaneously support. By default, the number of maximum connections is set at 20,000. This option sets a threshold as an artificial barrier, such that new connections will not be issued to this server if the 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 servers in a real server group for a virtual server reach their maxcon limit at the same time, client requests will be sent to the backup/overflow server or backup/overflow server group. If no backup servers/server group are configured, client requests will be dropped by the virtual server.

Table 7-2 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
tmout *<even number of minutes (2-30)>* 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 Table. When a client makes a request, the session is recorded in the table, the data is transferred until the client ends the session, and the session table entry is then removed. In certain circumstances, such as when a client application is abnormally terminated by the client's system, TCP/UDP connections will remain registered in the switch's binding table. In order to prevent table overflow, these orphaned entries must be aged out of the binding table. Using the tmout option, you can set the number of minutes to wait before removing orphan table entries. Settings must be specified in even numbered increments between 2 and 30 minutes. The default setting is 10. This option is also used with the Persistent option (see `/cfg/slb/virt/pbind`). When persistent is activated, this option sets how long an idle client is allowed to remain associated with a particular server.
backup *<real server number (1-255)>*| **none** Sets the real server used as the backup/overflow server for this real server. To prevent loss of service if a particular real server fails, use this option to assign a backup real server number. Then, if the real server becomes inoperative, the switch will activate the backup real server until the original becomes operative again. The backup server is also used in overflow situations. If the real server reaches its `maxcon` (maximum connections) limit, the backup comes online to provide additional processing power until the original server becomes desaturated. The same backup/overflow server may be assigned to more than one real server at the same time
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 `inter` option lets you choose the time between health checks. The range is from 1 to 60 seconds. The default interval is 2 seconds. An interval of "0" disables health checking for the server.
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
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

Table 7-2 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
addport *<real server port (2–65534)>* Add multiple service ports to the server.
remport *<real server port (2–65534)>* Remove multiple service ports from the server.
remote disable|enable Enables or disables remote site operation for this server. This option should be enabled when the real IP address supplied above represents a remote server (real or virtual) that this switch will access as part of its Global Server Load Balancing network. By default, this option is disabled.
proxy disable|enable 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).
submac disable|enable Enables or disables source MAC address substitution. By default, this option is disabled.
exclude disable|enable Enables or disables exclusionary string matching.
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/ena` on page 290 for an operations-level command.
dis 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` is are used, disables this real server until it is explicitly re-enabled. This option *does not* perform a graceful server shutdown. See `/oper/slb/dis` on page 290 for an operations-level command.
del 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.
cur Displays the current configuration information for this real server.

/cfg/slb/real <server number>/layer7

Real Server Layer 7 Configuration

```
[Layer 7 Commands  Menu]
     addlb    - Add URL path for URL load balance
     remlb    - Remove URL path for URL load balance
     cookser  - Enable/disable cookie assignment server
     exclude  - Enable/disable exclusionary string matching
     cur      - Display current real server configuration
```

This menu is used for entering commands and strings for Layer 7 processing.

Table 7-3 Layer 7 Commands Menu Options (/cfg/slb/real/layer7)

Command Syntax and Usage
addlb *<URL path ID (1-128)>* Adds the predefined URL loadbalance string ID to the real server.
remlb *<URL path ID (1-128)>* Removes the predefined URL loadbalance string ID from the real server.
cookser disable|enable Enables or disables the real server to handle client requests that don't contain a cookie. This option is used if you want to designate a specific server to assign cookies only. This server gets the client request, assigns the cookie, and embeds the IP address of the real server that will handle the subsequent requests from the client. By default, this option is disabled.
exclude disable|enable Enables or disables *exclusionary string* matching. By default, this option is disabled.
cur Displays the current real server configuration.

/cfg/slb/group <real server group number>
Real Server Group SLB Configuration

```
[Real server group 1 Menu]
      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
      viphlth - Enable/disable VIP health checking in DSR mode
      add     - Add real server
      rem     - Remove real server
      del     - Delete real server group
      cur     - Display current group configuration
```

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

Table 7-4 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
metric leastconns|**roundrobin**|**minmisses**|**hash**|**response**|**bandwidth** Set the load balancing metric used for determining which real server in the group will be the target of the next client request. The default setting is `leastconns`. See "Server Load Balancing Metrics" on page 245.
content <*filename*>|**//**<*host*>**/**<*filename*>|**none** This option defines the specific content which is examined during health checks. The content depends on the type of health check specified in the `health` option (see below).

Table 7-4 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage	
health	**link|arp|icmp|tcp|http|dns|pop3|smtp|nntp|ftp|imap|radius|sslh|script<n>|udpdns|wsp|wtls|ldap**
Sets the type of health checking performed. The default is `tcp`. The options are as follows:	
`link`	For IDSLB group only, checks status of port for each server.
`arp`	For Layer 2 health checking, sends an ARP request.
`icmp`	For Layer 3 health checking, `pings` the server.
`tcp`	For TCP service, opens and closes a TCP/IP connection to the server.
`http`	For HTTP service, use HTTP 1.1 `GETS` when a `HOST:` header is required to check that the URL content is specified in `content` command. Otherwise, an `HTTP/1.0 GET` occurs. **Note:** If the content is not specified, the health check will revert back to TCP on the port that is being load balanced.
`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 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 serve
`radius`	For RADIUS remote access server authentication, check that the *user:password* value specified in `content` exists on the Web switch and the server. To perform application health checking to a RADIUS server, the network administrator must also configure the `/cfg/slb/secrt` parameter. The `secrt` 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) and by the RADIUS server to decrypt the password during verification.
`sslh`	Enables the switch to query the health of the SSL servers by sending an SSL client "Hello" packet and then verify the contents of the server's "Hello" response. During the handshake, the user and server exchange security certificates, negotiate an encryption and compression method, and establish a session ID for each session.
`script`	Enables the use of script-based health checks in send/expect format to check for application and content availability. <n> denotes the health script number (1-8).
`wsp`	Enables connectionless WSP content health checks for WAP gateways. The content under `/cfg/slb/adv/waphc` (see page 284) must also be configured.
`udpdns`	Allows the user to perform health checking using UDP DNS queries.
`wtls`	Provides Wireless Transport Layer Security (WTLS) Hello-based health check for encrypted and connection-oriented WTLS traffic on port 9203.
`ldap`	Sets the health check type to LDAP.

Table 7-4 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
backup r<*real server number (1-256)*>|**g**<*group number*>|**none** Sets the real server or real server group used as the backup/overflow server/server group for this real server group. To prevent loss of service if the entire real server group fails, use this option to assign a backup real server/real server group number. Then, if the real server group becomes inoperative, the switch will activate the backup real server /server group until one of the original real servers becomes operative again. The backup server/server group is also used in overflow situations. If all the servers in the real server group reach their `maxcon` (maximum connections) limit, the backup server/server group comes online to provide additional processing power until one of the original servers becomes desaturated. The same backup/overflow server/server group may be assigned to more than one real server group at the same time.
name <*string, maximum 31 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.
realthr <*real server failure threshold (1-31)*> Specifies a minimum number of real servers available. If any time, the number reaches this minimum limit, a SYSLOG ALERT message is sent to the configured SYSLOG servers stating that the real server threshold has been reached for the concerned server load balancing group. The default threshold is 0, which also means the option is disabled
viphlth Enables or disables VIP health checking in a service. This feature is enabled by default. However, it works only when the service has DSR (Direct Server Return) feature enabled. When **viphlth** is disabled, the switch uses RIP to perform all health checks, whether DSR is enabled or disabled.
add <*real server number (1-255)*> Adds a real server to this real server group. You will be prompted to enter the number (1 to 256) of the real server to add to this group.
rem <*real server number (1-255)*> Remove a real server from this real server group. You will be prompted for the ID number for the real server to remove from this group.
del Deletes this real server group from the Layer 4 software configuration. This removes the group from operation under all virtual servers it is assigned to. Use this command with caution: if you remove the only group that is assigned to a virtual server, the virtual server will become inoperative.
cur Displays the current configuration parameters for this real server group.

Server Load Balancing Metrics

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 7-5 Real Server Group Metrics (/cfg/slb/group/metric)

Option and Description
`minmisses` Minimum misses. This 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. Minmisses can also be used for Server Load Balancing. When specified for a real server group performing Server Load Balancing, all requests from a specific client will be sent to the same server. This is useful for applications where client information must be retained on the server between sessions. Server load with this metric becomes most evenly balanced as the number of active clients increases.
`hash` Like `minmisses`, the `hash` metric uses IP address information in the client request to select a server. For Application Redirection, all requests for a specific IP destination address will be sent to the same server. This is particularly useful for maximizing successful cache hits. For Server Load Balancing, all requests from a specific client will be sent to the same server. This is useful for applications where client information must be retained between sessions. The `hash` metric should be used if the statistical load balancing achieved using `minmisses` is not as optimal as desired. Although the `hash` metric can provide more even load balancing at any given instance, it is not as effective as `minmisses` when servers leave and reenter service. If the Load Balancing statistics indicate that one server is processing significantly more requests over time than other servers, consider using the `hash` metric.
`leastconns` Least connections. With this option, the number of connections currently open on each real server is measured in real time. The server with the fewest current connections is considered to be the best choice for the next client connection request. This option is the most self-regulating, with the fastest servers typically getting the most connections over time, due to their ability to accept, process, and shut down connections faster than slower servers.
`roundrobin` Round robin. With this option, new connections are issued to each server in turn: the first real server in this group gets the first connection, the second real server gets the next connection, followed by the third real server, and so on. When all the real servers in this group have received at least one connection, the issuing process starts over with the first real server.

Table 7-5 Real Server Group Metrics (/cfg/slb/group/metric)

Option and Description
`response` Real server response time. With this option, the switch monitors and records the amount of time that each real server takes to reply to a health check. The response time is used to adjust the real server weights. The weights are adjusted so they are inversely proportional to a moving average of response time.
`bandwidth` Bandwidth Metric. With this option, the real server weights are adjusted so they are inversely proportional to the number of octets that the real server processes during a given interval. The higher the bandwidth used, the smaller is the weight assigned to that server.

NOTE – Under the `leastconns` and `roundrobin` metrics, when real servers are configured with weights (see the `weight` option on page 238), a higher proportion of connections are given to servers with higher weights. This can improve load balancing among servers of different performance levels. Weights are not applied when using the `hash` or `minmisses` metrics.

/cfg/slb/virt *<virtual server number>*
Virtual Server SLB Configuration

```
[Virtual Server 1 Menu]
      service - Virtual Service Menu
      vip     - Set IP addr of virtual server
      dname   - Set domain name of virtual server
      cont    - Set BW Contract
      layr3   - Enable/disable layer 3 only balancing
      ena     - Enable virtual server
      dis     - Disable virtual server
      del     - Delete virtual server
      cur     - Display current virtual configuration
```

This menu is used for configuring the virtual servers which will be the target for client requests for Server Load Balancing. Configuring a virtual server requires the following parameters:

- Creating a virtual server IP address
- Adding TCP/UDP port and real server group
- Enabling the virtual server (disabled by default)

Table 7-6 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
service *<virtual port or name, from 2 - 65534>* Displays the Virtual Services Menu. The virtual port name can be a well-known port name, such as http, ftp, the service number, and so on. To get more information about well-known ports, see the **sport** command on page 256. To view services menu options, see page 248.
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 forth). An example would be foocorp.com. It does not include the hostname portion (www, www2, ftp, and so forth). The maximum number of characters that can be used in a domain name is 34. To define the hostname, see hname below. To clear the dname, specify the name as **none**.
cont <*BWM contract (1-1024)***>** Enter a new Bandwidth Management Contract for this virtual service. By default, all services under this virtual server are assigned this BW contract. However, the BW contract can be changed for a selected virtual server with /cfg/slb/virt *<number>*/service *<number>*/cont. All the frames that match this virtual server services are assigned this BW contract if the previously assigned contract for the frame has lower or equal precedence of the virtual server contract. The default number of contracts is set at 256 for AD3/180e and 1024 for AD4/184.
layr3 disable|enable Normally, the client IP address is used with the client Layer 4 port number to produce a session identifier. When the layr3 option is enabled (disabled by default), the switch uses only the client IP address as the session identifier. It associates all the connections from the same client with the same real server while any connection exists between them. This option is necessary for some server applications where state information about the client system is divided across different simultaneous connections, and also in applications where TCP fragments are generated. If the real server to which the client is assigned becomes unavailable, the Layer 4 software will allow the client to connect to a different server.
ena 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.
dis This option disables the virtual server so that it no longer services client requests.

Table 7-6 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
del This command removes this virtual server from operation within the switch and deletes it from the Layer 4 switching software configuration. Use this command with caution, as it will delete the options that have been set for this virtual server.
cur Displays the current configuration of the specified virtual server.

/cfg/slb/virt <server number>/service <virtual port or name>

Virtual Server Service Configuration

This menu is used for configuring services assigned to a virtual server. The following example shows a menu for http (port 80) services.

NOTE – Select virtual service port 554 to configure RTSP traffic. See page 252 to view the menu options for configuring virtual services on port 554 for RTSP.

```
[Virtual Server 1 http Service Menu]
      group   - Set real server group number
      rport   - Set real port
      hname   - Set hostname
      httpslb - Set HTTP SLB processing
      cont    - Set BW contract for this virtual service
      urlcont - Set BW cont of an SLB string specific to this service
      pbind   - Set persistent binding type
      rcount  - Set multi response count
      dbind   - Enable/disable delayed binding
      udp     - Enable/disable UDP balancing
      frag    - Enable/disable remapping UDP server fragments
      nonat   - Enable/disable only substituting MAC addresses
      dnsslb  - Enable/disable DNS query load balancing
      ftpp    - Enable/disable FTP SLB parsing for virtual server
      del     - Delete virtual service
      cur     - Display current virtual service configuration
```

Table 7-7 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage
group *<real server group number (1-256)>* Sets a real server group for this service. The default is set at 1. You will be prompted to enter the number (1 to 256) of the real server group to add to this service.
rport *<real server port (0-65534)>* Defines the real server TCP or UDP port assigned to this service. By default, this 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` *<number>*`/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 dname (above) to create a full host/domain name for individual services. The format for this command is: # **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 (above), "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**
dbind Enables or disables Layer 4 Delayed Binding for TCP service and ports. Enabling this command protects the server from Denial of Service (DoS) attacks. This option is disabled by default.
httpslb urlslb|host|cookie|browser|urlhash|headerhash|others|none Load balances on the following applications: ■ `urlslb`: Enable or disable URL SLB ■ `host`: Enable or disable for virtual hosting ■ `cookie`: Enable or disable cookie-based SLB for cookie-based preferential load balancing. You will be prompted for the following: Cookie name, starting point of the cookie value, number of bytes to be extracted, enable/disable checking for cookie in URI ■ `browser`: Enable or disable SLB, based on browser type ■ `urlhash`: Enable or disable URL hashing based on URI ■ `headerhash`: Hashes on any HTTP header value. ■ `others`: Requires inputs for a particular header field ■ `none`: Removes any HTTPSLB configuration. You may choose to combine or select applications to load balance using the commands *and* and/or *or*. For example: ■ **httpslb** *<none>* ■ **httpslb** *<application>* ■ **httpslb** *<application>* and|or *<application>*

Table 7-7 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage
cont *<URL path ID (1-128)> <BWM Contract (0-1024)>* Sets a Bandwidth Management contract for this virtual service. The default number of contracts is set at 256 for Alteon AD3/180e and 1024 for Alteon AD4/184 Web switches. **Note:** If you enter 0 for the service contract, it will carry the value entered for the Virtual Server IP (`vip`) contract.
urlcont *<URL path ID (1-128)> <BWM contract <1-1024)>* Sets the Bandwidth Management contract of a string specific to this virtual service. Only use this command when a string is shared by multiple virtual services and each service requires a separate bandwidth. The default is set at 1024.
pbind clientip|cookie|sslid|disable Enables or disables persistent bindings for a real server (disabled by default). This may be necessary for some server applications where state information about the client system is retained on the server over a series of sequential connections, such as with SSL (Secure Socket Layer, HTTPS), Web site search results, or multi-page Web forms. ■ The `clientip` option uses the client IP address as an identifier, and associates all connections from the same client with the same real server until the client becomes inactive and the connection is aged out of the binding table. The connection timeout value (set in the Real Server Menu) is used to control how long these inactive but persistent connections remain associated with their real servers. When the client resumes activity *after* their connection has been aged out, they will be connected to the most appropriate real server based on the load balancing metric. An alternative approach may be to use the real server group metrics `minmisses` or `hash` (see Server Load Balancing Metrics). ■ The `cookie` option uses a cookie defined in the HTTP header or placed in the URI for hashing. For more information on `cookie` option, see "Cookie-Based Persistence" on page 253. For detailed information on Cookie-Based Persistence, see the *Persistence* chapter in the *Web OS Application Guide*. ■ The `sslid` option is for Secure Sockets Layer (SSL), which is a set of protocols built on top of TCP/IP that allow an application server and user to communicate over an encrypted HTTP session. SSL provides authentication, non-repudiation, and security. The session ID is a value comprising 32 random bytes chosen by the SSL server that gets stored in a session hash table. By enabling the `sslid` option, all subsequent SSL sessions which present the same session ID will be directed to the same real server. ■ The `disable` option enables you to disable presistent binding, if it has previously been enabled for a particular application.
rcount *<1–16>* Sets the maximum response counter for cookie-based persistence. The Web switch will examine each server response until the cookie is found, or until the maximum count is reached. The default number is 1.

Table 7-7 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage
udp disable|enable|stateless Enables or disables UDP load balancing for a virtual port (disabled by default). 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. **Note:** If applying a filter to the same virtual server IP address on which UDP load balancing is enabled, *disable caching on that filter for optimal performance*. For more information, see the **cache** command in Table 7-12 on page 260.
frag disable|enable Enables or disables remapping server fragments for virtual port. This option is enabled by default.
nonat disable|enable Enables or disables substituting only the MAC address of the real server (disabled by default). This option does not substitute IP addresses. This option is used for Direct Server Return (DSR) in an one-armed load balancing setup, so that frames returning from server to the client do not have to pass through the switch.
dnsslb Enables or disables DNS-based Layer 7 content load balancing.
ftpp disable|enable Enables or disables FTP SLB parsing for this virtual server (disabled by default). When this option is enabled, the switch modifies the appropriate FTP method/command to support FTP servers on a private network for both active and passive FTP modes. To do this, the switch looks deeper into the packet and modifies the `port` command for active FTP or the "entering the passive mode" command for passive FTP.
del 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.
cur Displays the current configuration of services on the specified virtual server.

/cfg/slb/virt <server number>/service 554

Virtual Server RTSP Configuration

This menu displays virtual services configured on service port 554 for RTSP traffic. See "/cfg/slb/virt <server number>/service <virtual port or name>" on page 248 to view the menu options for configuring virtual services on port 80.

```
[Virtual Server 1 rtsp Service Menu]
      group   - Set real server group number
      rtspslb - Enable/disable RTSP URL load balancing
      del     - Delete virtual service
      cur     - Display current virtual service configuration
```

Table 7-8 Virtual Server Service Configuration Menu Options (/cfg/slb/virt 1/service 554)

Command Syntax and Usage
group Sets real server group number.
rtspslb hash|pattern|disable This Layer 7 load balancing option sets the type of `rtspslb` (`hash|pattern`, thereby enabling the service), or disables `rtspslb` service altogether with `disable` command. To enable Layer 7 load balancing for RTSP service, `group` must be configured under the menu `/cfg/slb/virt` <*virtual server number*>`/service 80`. If you don't configure `group`, service 80 and service RTSP will load balance the default group, which is `group 1`. See command `group` in the table "Virtual Server Service Configuration Options (/cfg/slb/virt/service)" on page 249 for details on how to configure `group`.
del Deletes this virtual service.
cur Displays the current virtual service configuration.

/cfg/slb/virt/service 80/pbind cookie

Cookie-Based Persistence

The `cookie` option is used to establish cookie-based persistence, and has the following command syntax and usage:

`pbind cookie` *<mode> <name> <offset> <length> <URI>*

Each parameter is explained in the following table.

Table 7-9 Command Syntax and Usage for `pbind cookie` Options (/cfg/slb/virt/service/pbind cookie)

Option	Description
<mode>	Specify the mode for cookie-based persistence. The following three modes are available: ■ **`p:`** Passive mode. In this mode, the network administrator configures the Web server to embed a cookie in the server response that the switch looks for in subsequent requests from the same client. ■ **`r:`** Rewrite mode. In active cookie mode (or cookie rewrite mode), the switch, and not the network administrator, generates the cookie value on behalf of the server. The switch intercepts this persistence cookie and rewrites the value to include server-specific information before sending it to the client. ■ **`i:`** Insert mode. When a client sends a request *without* a cookie, the server responds with the data, and the switch inserts an *Alteon persistence cookie* into the data packet. The switch uses this cookie to bind to the appropriate server. Insert cookie mode expiration parameters are as follows: Enter insert-cookie expiration as either: ■ ... a date *<MM/dd/yy[@hh:mm]>* (e.g. 12/31/01@23:59) ■ ... a duration *<days[:hours[:minutes]]>* (e.g. 45:30:90) ■ ... or none *<return>*
<name>	Enter the name of the cookie.
<offset>	Enter the starting point of the cookie value (1-64)
<length>	Enter number of bytes to extract (1-64). For cookie rewrite, the extracting length must be 8 or 16.
<URI>	Look for cookie in the URI. If you want to look for cookie name or value in the URI, enter **`e`** to enable this option. To look for cookie in the HTTP header, enter **`d`** to disable this option.

For more information on Cookie-Based Persistence, see the *Web OS 10.0 Application Guide*.

/cfg/slb/filt <filter number>
SLB Filter Configuration

```
[Filter 1   Menu]
      adv     - Filter Advanced Menu
      name    - Set filter name
      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
      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
      vlan    - Set vlan id
      invert  - Enable/disable filter inversion
      ena     - Enable filter
      dis     - Disable filter
      del     - Delete filter
      cur     - Display current filter configuration
```

The switch supports up to 2048 traffic filters. Each filter can be configured to allow, deny, redirect or perform Network Address Translation on traffic according to a variety of address and protocol specifications, and each physical switch port can be configured to use any combination of filters. This command is disabled by default.

There are several options available in the Filter Advanced Menu (`/cfg/slb/filt/adv`, page 259) that can be used to provide more information through syslog. The types of information include:

- IP protocol
- TCP/UDP ports
- TCP flags
- ICMP message type

The following parameters are required for filtering:

- Set the address, masks, and/or protocol that will be affected by the filter
- Set the filter action (allow, deny, redirect, nat)
- Enable the filter
- Add the filter to a switch port
- Enable filtering on the Web switch port

Table 7-10 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
adv Displays the Filter Advanced Menu. To view menu options, see page 259.
name Allows the user to assign a name to a filter.
smac any|*<MAC address>* Sets the source MAC address. The default is any.
dmac any|*<MAC address>* Sets the destination MAC address. The default is any.
sip any|*<IP address>* 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. The default is any if the source MAC address is any.
smask *<IP address>* This IP address mask is used with the sip to select traffic which this filter will affect. See details below for more information on producing address ranges. For more information, see "Defining IP Address Ranges for Filters" on page 258.
dip any|*<IP address>* If defined, traffic with this destination IP address will be affected by this filter. Specify an IP address in dotted decimal notation, or "**any**". A range of IP addresses is produced when used with the dmask below. The default is any if the destination MAC address is any. For more information, see "Defining IP Address Ranges for Filters" on page 258.
dmask *<IP subnet mask (such as 255.255.255.0)>* This IP address mask is used with the dip to select traffic which this filter will affect.

Table 7-10 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
proto any|<*number*>|<*name*> If defined, traffic from the specified protocol is affected by this filter. Specify the protocol number, name, or "**any**". The default is any. Listed below are some of the well-known protocols. Number Name 1 icmp 2 igmp 6 tcp 17 udp 89 ospf 112 vrrp
sport any|<*name*>|<*port*>|<*port*>-<*port*> If defined, traffic with the specified TCP or UDP source port will be affected by this filter. Specify the port number, range, name, or "**any**". The default is any. Listed below are some of the well-known ports:: Number Name 20 ftp-data 21 ftp 22 ssh 23 telnet 25 smtp 37 time 42 name 43 whois 53 domain 69 tftp 70 gopher 79 finger 80 http 109 pop2 110 pop3
dport any|<*name*>|<*port*>|<*port*>**-**<*port*> If defined, traffic with the specified real server TCP or UDP destination port will be affected by this filter. Specify the port number, range, name, or "**any**", just as with sport above. The default is set at any.

Table 7-10 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
action Specify the action this filter takes: `allow` Allow the frame to pass (by default). `deny` Discard frames that fit this filter's profile. This can be used for building basic security profiles. `redir` Redirect frames that fit this filter's profile, such as for web cache redirection. In addition, Layer 4 processing must be activated (see the `/cfg/slb/on` command on page 235). `nat` Perform generic Network Address Translation (NAT). This can be used to map the source or destination IP address and port information of a private network scheme to/from the advertised network IP address and ports. This is used in conjunction with the `nat` option below and can also be combined with proxies.
group *<real server group number (1-16)>* This option applies only when `redir` is specified at the filter action. Define a real server group (1 to 16) to which redirected traffic will be sent. The default is group 1
rport *<real server port (0-65535)>* This option applies only when `redir` is specified at the filter action. This defines the real server TCP or UDP port to which redirected traffic will be sent. For valid Layer 4 health checks, this must be configured whenever TCP protocol traffic is redirected. Also, if transparent proxies are used for Network Address Translation (NAT) on the Web switch (see the `pip` option in Table 7-17 on page 266), `rport` must be configured for all Application Redirection filters. The default is set at 0.
nat source|dest When `nat` is set as the filter action (see above), this command specifies whether Network Address Translation (NAT) is performed on the source or the destination information. Destination (`dest`) is set as the default filter. If `source` is specified, the frame's source IP address (`sip`) and port number (`sport`) are replaced with the `dip` and `dport` values. If `dest` is specified, the frame's destination IP address (`dip`) and port number (`dport`) are replaced with the `sip` and `sport` values.
vlan any|<*VLAN ID (1-4094)***>** Sets the ID of the VLAN that is to be filtered. This option allows you to match the VLAN ID of the switch against the VLAN ID of the incoming packet. The default is `any`, which means the switch will match any VLAN ID of the incoming packet This command allows filters to be configured on per VLAN basis, and applies a filter to a VLAN that already has been configured. The VLAN has a set of member ports. By applying this filter to a VLAN, the filter gets applied to all member ports of this VLAN.
inver disable|enable Inverts the filter logic. If the conditions of the filter are met, *don't* act. If the conditions for the filter are *not met*, perform the assigned action. This option is disabled by default.

Table 7-10 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
ena Enables this filter.
dis Disables this filter.
del Deletes this filter.
cur Displays the current configuration of the filter.

Defining IP Address Ranges for Filters

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

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

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

Table 7-11 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

/cfg/slb/filt <filter number>/adv

Advanced Filter Configuration

```
[Filter 1 Advanced Menu]
      tcp      - TCP Advanced Menu
      ip       - IP Advanced Menu
      l7deny   - Layer 7 Deny Menu
      icmp     - Set ICMP message type
      cont     - Set BW contract
      tmout    - Set NAT session timeout
      proxy    - Enable/disable client proxy
      cache    - Enable/disable caching sessions that match filter
      log      - Enable/disable logging
      fwlb     - Enable/disable firewall redirect hash method
      linklb   - Enable/disable WAN link load balancing
      idshash  - Set hash parameter for intrusion detection SLB
      urlp     - Enable/disable URL parsing
      urlcont  - Set BW cont of an URL path specific to this filter
      ftpa     - Enable/disable active FTP NAT
      rdsnp    - Enable/disable WAP RADIUS Snooping
      thash    - Set hash parameter for Filter
      cur      - Display current advanced filter configuration
```

Table 7-12 Advanced Filter Menu (/cfg/slb/filt/adv)

Command Syntax and Usage
tcp Displays the TCP Flags Advanced Menu. To view menu options, see page 261.
ip Sets IP advanced menu. To view menu options, see page 263.
l7deny Displays Layer7 Deny Filter menu. To view menu options, see page 265.
icmp any|<*number*>|**name** Sets the ICMP message type. The default is set at any. For a list of ICMP message types, see Table 7-15 on page 264. For a detailed description of filtering and ICMP, see the *Web OS 10.0 Application Guide.*

Table 7-12 Advanced Filter Menu (/cfg/slb/filt/adv)

Command Syntax and Usage
proxy disable|enable 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 is enabled. If disabled, any proxy defined for the switch port using the `pip` command (see page 266) is not performed for traffic meeting the filter criteria. This is useful when certain traffic must retain original IP address information, or when other forms of translation (such as Application Redirection or NAT) are preferred.
cache disable|enable Enables or disables caching sessions that match the filter. Exercise caution while applying cache-enabled and cache-disabled filters to the same switch port. A cache-enabled filter creates a session entry in the switch, so that the switch can bypass checking for subsequent frames that match the same criteria. Cache is enabled by default. **Note:** Cache should be disabled if applying a filter to virtual server IP address while performing UDP load balancing (see "`udp disable|enable|stateless`" on page 251).
cont *<BWM Contract (1-1024)>* Sets the Bandwidth Management Contract. By default, the contract number is set at 256 for the Alteon AD3/180e and 1024 for Alteon AD4/184 Web switches.
tmout Sets the Network Address Translation (NAT) session timeout in an even number of minutes (4–30). The default is set at 4 minutes.
log disable|enable Enables or disables logging filter messages. This option is disabled by default.
fwlb disable|enable To ensure that the *stateful inspection* behavior of firewalls is maintained, a hashing algorithm is used to ensure that inbound packets and outbound packets for a pair of IPSA/IPDA traverse through the same firewall. If the dport is 80 or 21, enabling this option changes the hash of the filter from a WCR hash to a FWLB hash. By default, this option is disabled.
linklb Enables or disables WAN Link Load Balancing. By default, this option is disabled.
idshash sip|dip|both Sets the hash metric parameter for Intrusion Detection System Server Load Balancing
urlp disable|enable Enables or disables URL parsing. By default, this option is disabled.
urlcont *<URL path ID (1-128)> <BW contract (1-1024)>* Sets the URL path BW contract for this filter. Only use this command when a string is shared by multiple filters and each filter requires a separate bandwidth.

Table 7-12 Advanced Filter Menu (/cfg/slb/filt/adv)

Command Syntax and Usage
ftpa disable|enable Enables or disables active FTP Client Network Address Translation (NAT). When a client in active FTP mode sends a PORT command to a remote FTP server, the switch will look into the data part of the frame and replace the client 's private IP address with a proxy IP (PIP) address. The real server port (RPORT) will be replaced with a proxy port (PPORT), that is PIP:PPORT. By default, this option is disabled.
rdsnp Enables or disables WAP RADIUS Snooping capability of a filter. By default, this option is disabled.
thash auto|sip|dip|both|sip+sport Allows you to choose hash parameter to use for filter redirection. The Default is auto. The sip option allows you to perform tunable hash on source IP address for this filter. The option dip allows you to perform tunable hash on destination IP address for this filter. The option both allows you to perform tunable hash on both source IP address and the destination IP address at the same time. The option sip+sport allows you to perform tunable hash on both source IP address and source port at the same time.
cur Displays the current advanced filter configuration.

/cfg/slb/filt <*filter number*>/adv/tcp

Advanced Filter TCP Configuration

```
[TCP advanced menu Menu]
      urg     - Enable/disable TCP URG matching
      ack     - Enable/disable TCP ACK matching
      psh     - Enable/disable TCP PSH matching
      rst     - Enable/disable TCP RST matching
      syn     - Enable/disable TCP SYN matching
      fin     - Enable/disable TCP FIN matching
      ackrst  - Enable/disable TCP ACK or RST matching
      tcplim  - Enable/disable TCP connection rate limiting
      maxconn - Set maximum connections for TCP rate limiting
      cur     - Display current TCP configuration
```

These commands can be used to configure packet filtering for specific TCP flags.

Table 7-13 Advanced Filter TCP Menu (/cfg/slb/filt/adv/tcp)

Command Syntax and Usage
urg disable|enable Enables or disables TCP URG (urgent) flag matching. By default, this option is disabled.
ack disable|enable Enables or disables TCP ACK (acknowledgement) flag matching. By default, this option is disabled.
psh disable|enable Enables or disables TCP PSH (push) flag matching. By default, this option is disabled.
rst disable|enable Enables or disables TCP RST (reset) flag matching. By default, this option is disabled.
syn disable|enable Enables or disables TCP SYN (synchronize) flag matching. By default, this option is disabled.
fin disable|enable Enables or disables TCP FIN (finish) flag matching. By default, this option is disabled.
ackrst Enables or disables TCP acknowledgement or reset flag matching. By default, this option is disabled.
tcplim Enables or disables TCP connection rate limiting. By default, this option is disabled.
maxconn Sets the maximum limit for new TCP connections in units of 10. To set the maximum number of connections (2,550), enter 250. To set the minimum number of connections (10, from the same user), enter 1. The default is 10 (100 connections).
cur Displays the current Access Control List TCP filter configuration.

/cfg/slb/filt <filter number> /adv/ip

IP Advanced Menu

```
[IP advanced menu]
      tos     - Set IP Type of Service
      tmask   - Set IP TOS mask
      newtos  - Set new IP TOS
      option  - Enable/disable IP option matching
      cur     - Display current IP configuration
```

Table 7-14 IP Advanced Menu Options (/cfg/slb/filt #/adv/ip)

Command Syntax and Usage
tos *<0-255>* Sets IP type of service (ToS) and the value of the type of service. For more information on ToS, refer to RFC 1340 and 1349.
tmask *<0-255>* Sets IP type of service mask.
newtos *<0-255>* Sets new IP type of service.
option Enables or disables IP option matching.
cur Displays current advanced IP settings for the selected filter.

ICMP Message Types

The following ICMP message types are used with the `/cfg/slb/filt/adv/icmp` command. You can list all ICMP message types with the `/cfg/slb/filt/adv/icmp list` command.

Table 7-15 ICMP Message Types

Type #	Message Type	Description
0	`echorep`	ICMP echo reply
3	`destun`	ICMP destination unreachable
4	`quench`	ICMP source quench
5	`redir`	ICMP redirect
8	`echoreq`	ICMP echo request
9	`rtradv`	ICMP router advertisement
10	`rtrsol`	ICMP router solicitation
11	`timex`	ICMP time exceeded
12	`param`	ICMP parameter problem
13	`timereq`	ICMP timestamp request
14	`timerep`	ICMP timestamp reply
15	`inforeq`	ICMP information request
16	`inforep`	ICMP information reply
17	`maskreq`	ICMP address mask request
18	`maskrep`	ICMP address mask reply

/cfg/slb/filt <filter number> /adv/l7deny

Advanced Layer 7 Deny Filter

```
[Layer7 Deny Menu]
      ena     - Enable l7 deny
      dis     - Disable l7 deny
      addstr  - Add string for layer 7 deny
      remstr  - Remove string for layer 7 deny
      cur     - Current string for layer 7 deny
```

Table 7-16 Advanced Layer 7 Deny Filter (/cfg/slb/filt/adv/l7deny)

Command Syntax and Usage
ena Enables Layer 7 deny filter. This filter blocks the incoming packet with the virus string and prevents it from entering the network.
dis Disables Layer 7 deny filter.
addstr Adds the virus string to the Layer 7 deny filter.
remstr Removes the virus string from the Layer 7 deny filter.
cur Displays the current strings for Layer 7 deny filter.

/cfg/slb/port <port number>
Port SLB Configuration

```
[SLB port 1 Menu]
      client  - Enable/disable client processing
      server  - Enable/disable server processing
      rts     - Enable/disable RTS processing
      hotstan - Enable/disable hot-standby processing
      intersw - Enable/disable inter-switch processing
      proxy   - Enable/disable use of PIP for ingress traffic
      pip     - Set Proxy IP address for port
      filt    - Enable/disable filtering
      add     - Add filter to port
      rem     - Remove filter from port
      idslb   - Enable/disable intrusion detection server load balancing
      cur     - Display current port configuration
```

Web OS switch software allows you to enable or disable processing independently for each type of Layer 4 traffic (client and server) on a *per port* basis, expanding your topology options.

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 8-3 on page 290).

Table 7-17 Port Configuration Menu Options (/cfg/slb/port)

Command Syntax and Usage
client disable|enable 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, re-mapping 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. This option is disabled by default.
server disable|enable 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 re-maps 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. This option is disabled by default.

Table 7-17 Port Configuration Menu Options (/cfg/slb/port)

`rts disable|enable`

Enables or disables Return to Sender (RTS) load balancing on this port. This option is used for firewall load balancing or VPN load balancing applications. Enable `rts` on all client-side ports to ensure that traffic ingresses and egresses through the same port. This option is disabled by default.

For more information on using `rts`, see the "Firewall Load Balancing" and "VPN Load Balancing" chapters in the *Web OS 10.0 Application Guide*.

`hotstan disable|enable`

Enables or disables hot-standby processing. Use this option and the `intersw` option in conjunction with VRRP hot-standby failover. This option is disabled by default.

`intersw disable|enable`

Enables or disables inter-switch processing. This option is enabled for ports connected to a peer switch and is disabled by default.

`proxy disable|enable`

Enables or disables a proxy for traffic that ingresses this port. When the PIP is defined, client address information in Layer 4 requests is replaced with this proxy IP address.

In Server Load Balancing applications, this forces response traffic to return through the switch, 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 255). This option is disabled by default.

`pip` *<proxy IP address>*

Sets the proxy IP address for this port, using dotted decimal notation. When the `pip` is defined, client address information in Layer 4 requests is replaced with this proxy IP address.

`filter disable|enable`

Enables or disables filtering on this port. Enabling the filter sets up the Real Server to look into the VPN session table. This option is disabled by default.

`add` *<filter ID (1 to 2048)|block of IDs (first-last)>*

Adds a filter or a block of filters for use on this port. Enter filter ID (1 to 2048) or a contiguous block of filter IDs. For example, 1-100.

`rem` *<filter ID (1 to 2048)|block of IDs (first-last)>*

Removes a filter or a block of filters from use on this port. Enter filter ID (1 to 2048) or a contiguous block of filter IDs. For example, 1-100.

`idslb`

Enables or disables Intrusion Detection System Server Load Balancing. This option is disabled by default.

`cur`

Displays current system parameters.

/cfg/slb/gslb
Global SLB Configuration

```
[Global SLB Menu]
   site    - Remote Site Menu
   lookup  - Network Preference Lookup Menu
   ttl     - Set Time To Live of DNS resource records
   mincon  - Set minimum number of site connections
   inter   - Set interval between remote site updates
   weight  - Set local weight
   dns     - Enable or disable DNS handoffs
   local   - Enable or disable DNS responses with only local
             addresses
   one     - Enable or disable DNS responses with only one address
   alway   - Enable or disable DNS responses at least one address
   geo     - Enable or disable geographic awareness
   http    - Enable or disable HTTP redirects
   usern   - Enable or disable HTTP redirect to real server name
   on      - Globally turn Global SLB ON
   off     - Globally turn Global SLB OFF
   cur     - Display current Global SLB configuration
```

NOTE – The `local`, `one`, `alway`, and `geo` options have no effect on lookup.

Table 7-18 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage
`site` *<remote site (1-64)>* Displays the Remote Site Menu for one of up to 64 remote sites. To view menu options, see page 271. By default, this option is disabled.
`lookup` Displays the Global SLB Lookup Menu. The options in this menu will overwrite the geographic awareness (IANA table) during DNS queries. To view menu options, see page 272.
`ttl` *<time to live in seconds (0-65535)>* Specifies the duration (from 0 to 65535 seconds, with default at 60) 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.

Table 7-18 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage
mincon *<minimum connections, 0-65535>* Sets the minimum number of available site connections. The default is set at 1024. 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 *<interval in minutes (1-120)>* Sets the time between Distributed Site State Protocol (DSSP) updates between this switch and its peers. The range is between 1 and 120 minutes. The default is 1 minute.
weight *<server weight (1-48)>* Sets the local weight. 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.
dns disable|enable Enables or disables DNS handoffs 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. This option is enabled by default.
local disable|enable Enables or disables switch responses to DNS queries with local virtual server IP addresses. This option is disabled by default. When enabled, the switch will always respond to DNS queries by providing a local virtual server IP address, as long as the virtual server IP address has healthy real servers with an aggregate number of available connections equal to 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 Enables or disables DNS responses with only one address. At most one IP address is included in each DNS response. This option is disabled by default.
alway disable|enable 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 VIP is returned in DNS response to eliminate long DNS timeouts caused by an empty response. This option is disabled by default.
geo disable|enable Enables or disables geographic awareness, such as the IANA table. This option is enabled by default. 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.

Table 7-18 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage	
`http disable`|**`enable`**	Enables or disables HTTP redirects to peer sites by this switch. When enabled (default), 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`**	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. This option is disabled by default. If `usern` is enabled, the client will be redirected to the domain name specified by the remote real server name plus virtual server domain name: *<remote real server name>* **.** *<virtual server domain name>*
`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 293).
`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. By default, GSLB is turned off.
`cur`	Displays current Global SLB configuration.

/cfg/slb/gslb/site <site number>

GSLB Remote Site Configuration

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

Up to 64 remote sites can be configured.

Table 7-19 GSLB Remote Site Menu Options (/cfg/slb/gslb/site)

Command Syntax and Usage
prima *<server IP address>* 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 *<server IP address>* If the remote site is configured with a redundant switch, enter the IP address of the IP interface for the remote secondary switch here. If the remote site primary switch fails, the local switch will address the remote site secondary switch instead.
name *<31 character name* | **"none"***>* Sets the name of the remote site. The default is set at **none**.
update disable|**enable** Enables or disables remote site updates. If enabled (default), 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. **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 Browser-Based Interface (BBI) to use a different service port (see the `/cfg/sys/wport` option on page 170).
ena Enables this remote site for use with Global Server Load Balancing.
dis Disables this remote site. The switch will no longer use this remote site for Global Server Load Balancing.

Table 7-19 GSLB Remote Site Menu Options (/cfg/slb/gslb/site)

Command Syntax and Usage
del Removes this remote site from operation and deletes its configuration.
cur Displays the current remote site configuration.

/cfg/slb/gslb/lookup

GSLB Lookup Configuration

```
[Global SLB Lookup Menu]
      network - Internet Network Preference Menu
      dname   - Set domain name for internal lookup table
      lookups - Enable or disable network preference lookups
      cur     - Display current lookup configuration
```

Table 7-20 GSLB Lookup Menu Options (/cfg/slb/gslb/lookup)

Command Syntax and Usage
network *<preference number (1-128)>* Displays the Internet Network Preference Menu. If enabled, the switch responds to DNS requests based on the configured `dname` and Internet Preference Menu option settings. To view menu options, see page 273.
dname *<domain name>*|**none** Sets the domain name for the internal lookup table. The maximum number of characters that a domain name can use is 34. The default is set at `none`.
lookups disable|enable Enables or disables network preference lookups. This option is disabled by default.
cur Displays the current lookup configuration.

/cfg/slb/gslb/lookup/network <preference#>

GSLB Internet Network Preference Lookup Configuration

```
[Network 1 Menu]
      sip     - Set Source IP address
      mask    - Set net mask
      vip1    - Set VIP address
      vip2    - Set VIP address
      del     - Delete internet network entry
      cur     - Display current internet network entry configuration
```

Up to 128 network preference numbers can be set. You can overwrite the IANA table by defining client networks, using the options in this menu. You should use regular GSLB to respond to a DNS request under the following conditions:

- Queried domain is not matched.
- Client IP address doesn't 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 number of 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 VIP 2 are configured. All client networks not in the Network Preference Menu will use this entry to respond to a DNS request.

Table 7-21 GSLB Internet Network Preference Menu Options

Command Syntax and Usage
sip *<IP address>* 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>* This IP address mask is used with the source IP SIP address to find a correct virtual server IP address to respond to a DNS request.
vip1 *<IP address>* Sets the first virtual server IP address. The address can either be a local or remote virtual server. The switch returns the VIP address with the least response time that is over the `mincon` (minimum number of available connections).
vip2 *<IP address>* Sets the second virtual server IP address.
del Deletes the specified network entry.
cur Displays the current Internet network entry configuration.

/cfg/slb/layer7
Layer 7 SLB Resource Definition Menu

```
[Layer 7 Resource Definition Menu]
      redir   - Web Cache Redirection Menu
      rtspred - RTSP Cache Redirection Menu
      slb     - Server Load Balancing Menu
```

Table 7-22 Layer 7 Resource Definition Menu Options (/cfg/slb/layer7)

Command Syntax and Usage
redir Displays the Web Cache Redirection Menu. To view menu options, see page 274.
rtspred Displays the RTSP Cache Redirection Menu. To view menu options, see page 276.
slb Displays the Server Load Balancing Menu. To view menu options, see page 276.

/cfg/slb/layer7/redir
Web Cache Redirection Configuration

```
[Web Cache Redirection Menu]
   add     - Add URL expression
   rem     - Remove URL expression
   urlal   - Enable or disable auto-ALLOW for non-GETs to origin servers
   cookie  - Enable or disable auto-ALLOW for Cookie to origin servers
   nocache - Enable or disable no-cache control header to origin servers
   hash    - Enable or disable URL hashing based on URI
   header  - Enable or disable server load balance based on HTTP header
   cur     - Display current URL expression table
```

Table 7-23 Web Cache Redirection Menu Options (/cfg/slb/layer7/redir)

Command Syntax and Usage
add *<string>* Adds the URL SLB redirected expression strings.
rem *<string>* Removes the URL SLB redirected expression strings.

Table 7-23 Web Cache Redirection Menu Options (/cfg/slb/layer7/redir)

Command Syntax and Usage
urlal disable|enable Enables or disables auto-ALLOW for non-GETs to origin servers. ■ If this command is enabled, the switch will redirect all non-GET requests to the origin server. ■ If this command is disabled, the switch will compare the URI against the expression table to determine whether all non-GET requests should be redirected to a cache server or origin server. This option is enabled by default.
cookie disable|enable Enables or disables auto-ALLOW for cookie to origin servers. ■ If this command is enabled, the switch will redirect all requests that contain *Cookie:* in the HTTP header to the origin server. ■ If this command is disabled, the switch will compare the URI against the expression table to determine whether it should redirect all requests that contain *Cookie:* in the HTTP header to a cache server or origin server. This option is disabled by default.
nocache disable|enable Enables or disables no-cache control header to origin servers. ■ If this command is enabled, the switch will redirect all requests that contain *Cache-Control: no-cache* in HTTP/1.1 header, or *Pragma: no-cache* in HTTP/1.0 header to the origin server. ■ If this command is disabled, the switch will compare the URI against the expression table to determine whether it should redirect requests that contain *Cache-Control: no-cache* in HTTP/1.1 header, or *Pragma: no-cache* in HTTP/1.0 header to a cache server or origin server. This option is enabled by default.
hash disable|enable Enables or disables URL hashing based on the URI. ■ If hashing is enabled, you can set the length of URI that will be used to hash into the cache server. ■ If hashing is disabled, the switch will only use the host header field to calculate the hash key. This option is disabled by default.
header disable|enable Enables or disables server load balancing based on HTTP header. This option is disabled by default.
cur Displays the current URL expression table.

cfg/slb/layer7/rtspred

RTSP Cache Redirection Configuration Menu

```
[RTSP Cache Redirection Menu]
      add     - Add non-cacheable RTSP URL expression
      rem     - Remove non-cacheable RTSP URL expression
      cur     - Display current non-cacheable RTSP URL expression table
```

Table 3 RTSP Cache redirection Configuration Menu Options (/cfg/slb/layer7/rtspred)

Command Syntax and Usage
add Adds non-cacheable RTSP URL expression to the RTSP expression table.
rem Removes non-cacheable RTSP URL expression from the RTSP expression table.
cur Displays the current non-cacheable RTSP URL expression table.

/cfg/slb/layer7/slb

Server Load Balance Resource Configuration Menu

```
[Server Loadbalance Resource Menu]
      message - Set HTTP error message
      add     - Add SLB string for load balance
      rename  - Rename SLB string for load balance
      rem     - Remove SLB string for load balance
      cont    - Set BW contract for the SLB string
      cur     - Display current configuration
```

Table 7-24 Server Load Balance Resource Menu Options (/cfg/slb/layer7/slb)

Command Syntax and Usage
message *<64 byte error message>* Sets the message that will be displayed when an error occurs. The default message is “No available server to handle this request.”
add *<SLB string ID>* Adds the SLB string for load balancing.

Table 7-24 Server Load Balance Resource Menu Options (/cfg/slb/layer7/slb)

Command Syntax and Usage
rename *<SLB string ID>* Renames the SLB string for load balancing.
rem *<SLB string ID>* Removes the SLB string from load balancing.
cont *<SLB string ID (1-128)> <BW contract number (1-1024)>* Sets the Bandwidth Management contract for a specified string for the SLB string ID.
cur Displays the current configuration of SLB string.

/cfg/slb/wap
WAP Configuration

```
[WAP Options Menu]
      tpcp    - Enable/disable WAP TPCP external notification
      debug   - WAP debug level
      cur     - Display current WAP configuration
```

Table 7-25 WAP Configuration Menu Options (/cfg/slb/wap)

Command Syntax and Usage
tpcp disable|enable Enables or disables the TPCP external notification for `Add`/`Delete` session requests. This option is disabled by default.
debug Sets the debug level for tracing the WAP related messages. The default is set at 0.
cur Displays the current WAP configuration

/cfg/slb/sync
Synchronize Peer Switch Configuration

```
[Config Synchronization Menu]
      peer    - Synch peer switch menu
      filt    - Enable or disable syncing filter configuration
      ports   - Enable or disable syncing port configuration
      prios   - Enable or disable syncing VRRP priorities
      pips    - Enable or disable syncing proxy IP addresses
      bwm     - Enable or disable syncing BWM configuration
      state   - Enable or disable syncing persistent session state
      update  - Set stateful failover update period
      cur     - Display current Layer 4 sync 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. Peers are sent SLB, FILT, and VRRP configuration updates using `/oper/slb/synch`.

Table 7-26 Synchronization Menu Options (/cfg/slb/sync)

Command Syntax and Usage
peer *<peer switch number (1-2)>* Displays the Sync Peer Switch Menu. This option is enabled by default. To view menu options, see page 279.
filt disable|enable Enables or disables synchronizing filter configuration.
ports disable|enable Enables or disables synchronizing Layer 4 port configuration. This option is enabled by default.
prios disable|enable Enables or disables syncing VRRP priorities. This option is enabled by default.
pips disable|enable Enables or disables synchronizing proxy IP addresses. This option is disabled by default.
bwm disable|enable Enables or disables synchronizing Bandwidth Management configuration between Master and backup switches. This option is enabled by default.
state disable|enable Enables or disables stateful failover for synchronizing the persistent session state. This option is disabled by default.

Table 7-26 Synchronization Menu Options (/cfg/slb/sync)

Command Syntax and Usage
update *<seconds, 1–60>* Sets the stateful failover update interval. The active server sends update packets of persistent binding entries to the backup switch at the specified update interval. The default value is 30 seconds.
cur Displays the current Layer 4 synchronization configuration.

/cfg/slb/sync/peer <peer switch number>

Peer Switch Configuration

```
[Peer Switch 1 Menu]
      addr    - Set peer switch IP address
      ena     - Enable peer switch
      dis     - Disable peer switch
      del     - Delete peer switch
      cur     - Display current peer switch 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 7-27 Peer Switch Configuration Menu Options (/cfg/slb/sync/peer)

Command Syntax and Usage
addr *<IP address>* Sets the peer switch IP address. The default is 0.0.0.0
ena Enables the peer for this switch. By default, this option is disabled.
dis Disables the peer for this switch.
del Deletes the peer for this switch
cur Displays the current peer switch configuration.

/cfg/slb/adv
Advanced Layer 4 Configuration

```
[Layer 4 Advanced Menu]
      script  - Scriptable Health Check Menu
      waphc   - WAP Health Check Menu
      aphttp  - Enable/disable Allow HTTP Health Check on any port
      synatk  - SYN Attack Detection Menu
      ldapver - LDAP version
      imask   - Set virtual and real IP address mask
      mnet    - Set managment network
      mmask   - Set management subnet mask
      pmask   - Set persistent mask
      secret  - Set RADIUS secret
      minter  - Set interval of response and bandwidth metric updates
      idslb   - Set real server group number for IDSLB
      timewin - Set time window for TCP rate limiting
      holddur - Set hold down duration for TCP rate limiting
      direct  - Enable/disable Direct Access Mode
      grace   - Enable/disable graceful real server failure
      matrix  - Enable/disable Virtual Matrix Architecture
      tpcp    - Enable/disable Transparent Proxy Cache Protocol
      fastage - Session table fast-age (1 sec) period bit shift
      slowage - Session table slow-age (2 min) period bit shift
      cur     - Display current Layer 4 advanced configuration
```

Table 7-28 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage
script *<health script number (1-16)>* Displays the Scriptable Health Check Menu. To view menu options, see page 283.
waphc Displays the WAP Health Check Menu. To view menu options, see page 284.
aphttp disable|enable Enables or disables HTTP health checks on any port. By default, this option is disabled. When disabled, you can use HTTP health checks only for HTTP service. Enabling it will allow you to use it on any port, like HTTPs.
synatk Displays SYN Attack Detection Menu. To view menu options, see page 285.
ldapver Sets the LDAP version to 2 or 3. The default is 2.

Table 7-28 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage
imask *<IP subnet mask (such as 255.255.255.0)>* Configures the real and virtual server IP address mask using dotted decimal notation. The default is 255.255.255.255.
mnet *<IP address>* 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)>* This IP address mask is used with the `mnet` to select management traffic which is allowed direct access to real servers. The default is 255.255.255.255.
pmask *<IP subnet mask (such as 255.255.255.0)>* Sets persistent mask. The default is 255.255.255.255.
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/secret` value and the `cntnt` 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) and by the RADIUS server to decrypt the password during verification. The default is **none**.
minter *<1-256 (number of seconds between updates)>* This command sets the interval of response and bandwidth metric updates. The default is set at 10.
idslb *<1-256>* Sets the real server group number for Intrusion Detection System server load balancing. The copied incoming sessions will be load-balanced across the servers in this real server group. The default is set at 1.
timewin Sets the parameter for time window for TCP rate limiting, which is a multiple of the `fastage` period. For example, if the `fastage` parameter is 2 seconds, and the `timewin` is 3, then the resulting time window is 6 seconds.
holddur This command configures the hold down duration, which is a multiple of `slowage`. Hold down (blocking of new TCP connections) occurs when the number of received SYN packets exceeds the threshold of a pre-defined time window. For example, the parameter of `slowage` is 2 minutes, and `holddur` is 5, then the resulting hold down duration is 10 minutes.
direct disable|enable Enable/disables Direct Access Mode to real servers/services. This option also allows any virtual server to load balance any real server. By default, this option is disabled.

Table 7-28 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage
grace disable|enable Enables or disables graceful real server failure. Allows existing connections to newly failed server to gracefully continue. By default, this option is disabled.
matrix disable|enable Enables or disables the use of Virtual Matrix Architecture on the Web switch. By default, this option is enabled.
tpcp disable|enable Enables or disables the TPCP (Transparent Proxy Cache Protocol). This command is used for security reasons—the UDP port can be closed. By default, this option is disabled.
fastage *<0–7>* Controls how frequently a *fastage scan* is performed. The default interval is two seconds. Each incremental increase of the value doubles the length of the interval. The `fastage` scan is used to remove TCP sessions that have been closed with a FIN and sessions that have been identified by the `slowage` scan as idle for the maximum allowed period. If a large value of `fastage` is used, a session can remain in the session table for a few minutes. The default is 0.
slowage *<0–15>* Controls how frequently a *slowage scan* is performed. The default interval is two minutes. Each incremental increase of the value doubles the length of the interval. (Value is set in bits rather than seconds, which causes the time to double per increment). The `slowage` scan is used to remove idle or non-TCP sessions from the session at the specified intervals. If a large value of slowage is used, a session can remain in the session table for months. The default is 0.
cur Displays the current Layer 4 advanced configuration.

/cfg/slb/adv/script <health script number>

Scriptable Health Checks Configuration

```
[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
      rem     - Remove last command from script
      del     - Delete script
      cur     - Display current script configuration
```

The Health Script menu provides commands that can be used to define the health "script." The total number of characters cannot exceed 1024 bytes. Up to eight scripts can be configured.

Table 7-29 Scriptable Health Check Menu Options (/cfg/slb/adv/script)

Command Syntax and Usage
open *<real port or name, for example: http>* Sets the TCP port to be opened.
send *<text string>* Sends an ASCII string through open TCP port. For example, an HTTP request, such as, `"GET /default.asp HTTP/1.1\\r\\nHOST: www.alteon.com\\r\\n\\r\\n."`
expect *<text string>* Expects an ASCII string for successful health check on open TCP port, such as an HTTP response: `HTTP/1.1 200`
close Closes TCP connection.
rem Removes the last entered line from the script.
del Deletes the current script.
cur Lists the current script configuration.

/cfg/slb/adv/waphc

WAP Health Check Configuration

```
[WAP Health Check Menu]
      wspport - WSP port number to health check
      wtlsprt - WTLS port number to health check
      offset  - Offset in received WSP packet
      sndcnt  - Content to be sent to the WAP gateway
      rcvcnt  - Content to be received from the WAP gateway
      cur     - Display current WAP health check configuration
```

Table 7-30 WAP Health Check Menu Options (/cfg/slb/adv/waphc)

Command Syntax and Usage
wspport *<port number (0-65534)>* Enter the port number on which WSP health checks will be performed. The default port number is 9200.
wtlsprt *<port number (0-65534)>* Enter the port number on which WTLS health checks will be performed. The default port number is 9203.
offset *<WSP receive offset (0-256)>* Enter the receive offset value content of the received WSP packages. An offset value of 0 (default) sets the switch to start comparisons from the beginning of the content of the received packet.
sndcnt *<hexadecimal string>* Enter a hexidecimal string that represents a connectionless WSP request to a WSP gateway. This string will be delivered to the WSP gateway.
rcvcnt *<hexadecimal string>* Enter a hexadecimal string that represents the content that the switch expects to receive from the WSP gateway.
cur Displays the current WAP Health Check configuration.

/cfg/slb/adv/synatk

SYN Attack Detection Configuration

```
[SYN Attack Detection Menu]
      intrval - Set SYN attack detection interval
      thrshld - Set SYN attack alarm threshold
      cur     - Display current SYN attack detection configuration
```

Table 7-31 SYN Attack Detection Menu Options (/cfg/slb/adv/synatk)

Command Syntax and Usage
intrval *<seconds (2-3600)>* Sets the interval of SYN attack inspection.
thrshld *<1-100000]<new half-open sessions/second)>* Sets the threshold of SYN attack alarm.
cur Displays the current SYN attack detection configuration.

CHAPTER 8

The Operations Menu

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

/oper
Operations Menu

```
[Operations Menu]
      port    - Operational Port Menu
      slb     - Operational Server Load Balancing Menu
      vrrp    - Operational Virtual Router Redundancy Menu
      bwm     - Operational Bandwidth Management 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.

Port Mirroring menu options are accessible only to the Alteon AD4 and Alteon 184 Web-Switches.

Table 8-1 Operations Menu Options (/oper)

Command Syntax and Usage
port *<port as number (1-9)>* Displays the Operational Port Menu. To view menu options, see page 289.
slb Displays the Operational Layer 4 Menu. To view menu options, see page 290.

Table 8-1 Operations Menu Options (/oper)

Command Syntax and Usage
vrrp Displays the Operational Virtual Router Redundancy Menu. To view menu options, see page 291.
bwm Operational Bandwidth Management Menu. To view menu options, see page 291.
ip Displays the IP Operations Menu, which has one sub-menu/option, the Operational Border Gateway Protocol Menu. To view menu options, see page 291.
swkey *<16-hex-digit key to enable software feature>* Enter key to enable software feature. For details, see page 293.
rmkey *<software feature to be removed>* Enter software feature to be removed. For details, see page 294.

/oper/port <port number>
Operations-Level Port Options

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

Operations-level port options are used for temporarily disabling or enabling a port, and for changing Remote Monitoring (RMON) status on a port.

Table 8-2 Operations-Level Port Menu Options (/oper/port)

Command Syntax and Usage
rmon disable|enable Temporarily enables/disables RMON on the port. The port will be returned to its configured operation mode when the switch is reset.
ena Temporarily enables the port. The port will be returned to its configured operation mode when the switch is reset.
dis Temporarily disables the port. The port will be returned to its configured operation mode when the switch is reset.
cur Displays the current settings for the port.

/oper/slb
Operations-Level SLB Options

```
[Server Load Balancing Operations Menu]
      sync    - Synchronize SLB, VRRP and other configurations on peers
      ena     - Enable real server
      dis     - Disable real server
      clear   - Clear session table
      cur     - 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 active/active switches.

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

Command Syntax and Usage
sync Synchronizes the SLB, filter, and VRRP configuration on a peer switch (a switch that owns the IP address). To take effect, peers must be configured on the Web switches and the administrator password on the Web switches must be identical.
ena *<real server number (1-1023)>* Temporarily enables a real server. The real server will be returned to its configured operation mode when the switch is reset.
dis *<real server number (1-1023)>* **p**|**n** The disable command is used to temporarily disable real servers as follows: ■ Using the **n** (none) option, disable the real server entirely, removing it from operation within its real server group and virtual server ■ Using the **p** (persistent) option, temporarily disable sessions except for persistent http 1.0 sessions. The real server will be returned to its configured operation mode when the switch is reset.
clear Clears all session tables and allows port filter changes to take effect immediately. **Note:** This command disrupts current Server Load Balancing and Application Redirection sessions.
cur Displays the current SLB operational state.

/oper/vrrp
Operations-Level VRRP Options.

```
[VRRP Operations Menu]
        back    - Set virtual router to backup
```

Table 8-4 Virtual Router Redundancy Operations Menu Options (/oper/vrrp)

Command Syntax and Usage
back *<virtual router number (1-256)>* 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 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/bwm
Operations-Level Bandwidth Management Options

```
[Bandwidth Management Operations Menu]
      sndhist - Send BW History to SMTP server
```

Table 8-5 Bandwidth Operations Menu Options (/oper/bwm/sndhist)

Command Syntax and Usage
sndhist Sends the bandwidth history to a system administrator specified under `/cfg/bwm/user` (see page 228).

/oper/ip
Operations-Level IP Options

```
[IP Operations Menu]
      bgp     - Operational Border Gateway Protocol Menu
```

Table 8-6 IP Operations Menu Options (/oper/ip)

Command Syntax and Usage
bgp Displays the Border Gateway Protocol Operations Menu, see.

/oper/ip/bgp
Operations-Level BGP Options

```
[Border Gateway Protocol Operations Menu]
      start   - Start peer session
      stop    - Stop peer session
      cur     - Current BGP operational state
```

Table 8-7 IP Operations Menu Options (/oper/ip)

Command Syntax and Usage
start *<peer number (1-4)>* Starts the peer session.
stop *<peer number (1-4)>* Stops the peer session.
cur Displays the current BGP operational state.

/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 your License Password, perform the following actions:

1. **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: [SLB]|GSLB|WCR: SLB
```

CHAPTER 9

The Boot Options Menu

To use the Boot Options Menu, you must be logged in to the switch as the administrator. The Boot Options Menu provides options for:

- Selecting a switch software image to be used when the switch is next reset
- Selecting a configuration block to be used when the switch is next reset
- Downloading or uploading a new software image to the switch via TFTP

/boot
Boot Menu

```
[Boot Options Menu]
        image - Select software image to use on next boot
        conf  - Select config block to use on next boot
        gtimg - Download new software image via TFTP
        ptimg - Upload selected 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 Web switch. A version of the image ships with the switch, and comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch.

Upgrading the software image on your switch requires the following:

- Loading the new image onto a TFTP server on your network
- Downloading the new image from the TFTP server to your switch
- Selecting the new software image to be loaded into switch memory the next time the switch is reset

Downloading New Software to Your Web Switch

The switch can store up to two different software images, called `image1` and `image2`, as well as boot software, called `boot`. When you download new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

To download a new software to your switch, you will need the following:

- The image or boot software loaded on a TFTP server on your network
- The hostname or IP address of the TFTP server
- The name of the new software image or boot file

NOTE – The DNS parameters must be configured if specifying hostnames. See "Domain Name System Configuration" on page 201).

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# gtimg
```

2. **Enter the name of the switch software to be replaced:**

```
Enter name of switch software image to be replaced
 ["image1"/"image2"/"boot"]: <image>
```

3. **Enter the hostname or IP address of the TFTP server.**

```
Enter hostname or IP address of TFTP server: <server name or IP address>
```

4. **Enter the name of the new software file on the server.**

```
Enter name of file on TFTP server: <filename>
```

The exact form of the name will vary by TFTP server. However, the file location is normally relative to the TFTP directory (usually `/tftpboot`).

5. **The system prompts you to confirm your request.**

You should next select a software image to run, as described below.

Selecting a Software Image to Run

You can select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# image
```

2. **Enter the name of the image you want the switch to use upon the next boot.**

The system informs you of which image is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use switch software "image1" on next reset.
Specify new image to use on next reset ["image1"/"image2"]:
```

Uploading a Software Image from Your Web Switch

You can upload a software image from the switch to a TFTP server.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# ptimg
```

2. **The system prompts you for information. Enter the desired image:**

```
Enter name of switch software image to be uploaded
["image1"|"image2"|"boot"]: <image>
```

3. **Enter the name or the IP address of the TFTP server:**

```
Enter hostname or IP address of TFTP server: <server name or IP address>
```

4. **Enter the name of the file into which the image will be uploaded on the TFTP server:**

```
Enter name of file on TFTP server: <filename>
```

5. **The system then requests confirmation of what you have entered. To have the file uploaded, enter Y.**

```
image2 currently contains Software Version 8.3.12
 that was downloaded at 15:05:47 Fri Dec  1, 2000.
Upload will transfer image2 (778656 bytes) to file "test"
 on TFTP server test.
Confirm upload operation [y/n]: y
```

Selecting a Configuration Block

When you make configuration changes to the Web 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 Web switch was manufactured. Under certain circumstances, it may be desirable to reset the switch configuration to the default. This can be useful when a custom-configured Web switch is moved to a network environment where it will be reconfigured for a different purpose.

Use the following procedure to set which configuration block you want the switch to load the next time it is reset:

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# conf
```

2. **Enter the name of the configuration block you want the switch to use:**

The system informs you of which configuration block is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use active configuration block on next reset.
Specify new block to use ["active"/"backup"/"factory"]:
```

Resetting the Web Switch

You can reset the switch to make your software image file and configuration block changes occur.

NOTE – Resetting the switch causes the Spanning Tree Protocol to restart. This process can be lengthy, depending on the topology of your network.

To reset the switch, at the `Boot Options#` prompt, enter:

```
>> Boot Options# reset
```

You are prompted to confirm your request.

CHAPTER 10

The Maintenance Menu

The Maintenance Menu is used to manage dump information and forward database information. It also includes a debugging menu to help with troubleshooting.

/maint
Maintenance Menu

NOTE – To use the Maintenance Menu, you must be logged in to the switch as the administrator.

```
[Maintenance Menu]
      sys     - System Maintenance Menu
      fdb     - Forwarding Database Manipulation Menu
      debug   - Debugging Menu
      arp     - ARP Cache Manipulation Menu
      route   - IP Route Manipulation Menu
      uudmp   - Uuencode FLASH dump
      ptdmp   - tftp put FLASH dump to tftp server
      cldmp   - Clear FLASH dump
      panic   - Dump state information to FLASH and reboot
      tsdmp   - Tech support dump
```

Dump information contains internal switch state data that is written to flash memory on the Web 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 that is attached to the console port. The switch reset key combination is <Shift><Ctrl><->.
- 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 10-1 Maintenance Menu Options (/maint)

Command Syntax and Usage
sys Displays the System Maintenance Menu. To view menu options, see page 302.
fdb Displays the Forwarding Database Manipulation Menu. To view menu options, see page 303.
debug Displays the Debugging Menu. To view menu options, see page 304.
arp Displays the ARP Cache Manipulation Menu. To view menu options, see page 305.
route Displays the IP Route Manipulation Menu. To view menu options, see page 306.
uudmp Displays dump information in uuencoded format. For details, see page 307.
ptdmp Saves the system dump information via TFTP. For details, see page 308.
cldmp Clears dump information from flash memory. For details, see page 308.
panic Dumps MP information to FLASH and reboots. For details, see page 309.
tsdmp Dumps all Web switch information, statistics, and configuration.You can log the tsdump output into a file, and send it to Alteon WebSystems Tech Support for debugging purposes.

/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]
      flags   - Set NVRAM flag word
      tmask   - Set MP trace mask word
```

/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
      refpt   - Show FDB entries referenced by a single port
      dump    - Show all FDB entries
      del     - Delete an FDB entry
      clear   - Clear entire FDB
```

The Forwarding Database Manipulation Menu can be used to view information and to delete a MAC address from the forwarding database or clear the entire forwarding database. This is helpful in identifying problems associated with MAC address learning and packet forwarding decisions.

Table 10-2 FDB Manipulation Menu Options (/maint/fdb)

Command Syntax and Usage
find *<MAC address>* [*<VLAN>*] Displays a single database entry by its MAC address. You are prompted to enter the MAC address of the device. Enter the MAC address using the `xx:xx:xx:xx:xx:xx` format (such as `08:00:20:12:34:56`) or `xxxxxxxxxxxx` format (such as `080020123456`).
port *<port number (0-9)>* Displays all FDB entries for a particular port. Use "0" for unknown port number.
vlan *<VLAN number (1-4094)>* Displays all FDB entries on a single VLAN.
refpt *<port number>* Displays all FDB entries reference by a single port.
dump Displays all entries in the Forwarding Database. For details, see page 67.
del *<MAC address>* Removes a single FDB entry.
clear Clears the entire Forwarding Database from switch memory.

/maint/debug
Debugging Options

```
[Miscellaneous Debug Menu]
      tbuf    - Show MP trace buffer
      snap    - Show MP snap (or post-mortem) trace buffer
      sptb    - Show SP trace buffer
      clrcfg  - Clear all flash configs
```

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 the Switch Processor (SP)
- Events traced to a buffer area when a reset occurs

If the switch resets for any reason, the MP trace buffer and SP trace buffers are saved into the snap trace buffer area. The output from these commands can be interpreted by the Alteon WebSystems Customer Support organization.

Table 10-3 Miscellaneous Debug Menu Options (/maint/debug)

Command Syntax and Usage
tbuf Displays the Management Processor trace buffer. Header information similar to the following is shown: `MP trace buffer at 13:28:15 Fri May 25, 2001; mask: 0x2ffdf748` The buffer information is displayed after the header.
snap Displays the Management Processor snap (or post-mortem) trace buffer. This buffer contains information traced at the time that a reset occurred.
sptb *<port number>* Displays the Switch Processor trace buffer. Header information similar to the following is shown: `Port 1 trace buffer at 13:37:22 Fri May 25, 2001; mask: 0x00800008` The buffer information is displayed after the header.
clrcfg Deletes all flash configuration blocks.

/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
      vlan    - Show ARP entries on a single VLAN
      refpt   - Show ARP entries referenced by a single port
      dump    - Show all ARP entries
      add     - Add a permanent ARP entry
      del     - Delete an ARP entry
      clear   - Clear ARP cache
      addr    - Show ARP address list
```

Table 10-4 Address Resolution Protocol Menu Options (/maint/arp)

Command Syntax and Usage
find *<IP address>* Shows a single ARP entry by IP address.
port *<port number>* Shows ARP entries on a single port.
vlan *<VLAN ID>* Shows ARP entries on a single VLAN.
refpt *<port number>* Shows all ARP entries referenced by a single port.
dump Shows all ARP entries.
add *<IP address> <MAC address> <VLAN number> <port>* Adds a single ARP entry from switch memory.
del *<IP address>* Removes a single ARP entry from switch memory.
clear Clears the entire ARP list from switch memory.
addr Shows the list of IP addresses which the switch will respond to 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 60.

/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
        dump  - Show all routes
        clear - Clear route table
```

Table 10-5 IP Route Manipulation Menu Options (/maint/route)

Command Syntax and Usage

find *<IP address>*
Shows a single route by destination IP address.

gw *<default gateway address>*
Shows routes to a default gateway.

type indirect|direct|local|broadcast|martian|multicast
Shows routes of a single type. For a description of IP routing types, see Table 4-5 on page 58

tag fixed|static|snmp|addr|rip|bgp|icmp|broadcast|martian|multicast|vip
Shows routes of a single tag. For a description of IP routing tags, see Table 4-6 on page 59

if *<interface number (1-256)>*
Shows routes on a single 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 57.

/maint/uudmp
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 analyzing the information.

If you want to capture dump information to a file, set your communication software on your workstation to capture session data prior to issuing the `uudmp` command. This will ensure that you do not lose any information. Once entered, the `uudmp` command will cause approximately 1460 lines of data to be displayed on your screen and copied into the file.

Using the `uudmp` command, dump information can be read multiple times. The command does not cause the information to be updated or cleared from flash memory.

NOTE – Dump information is not cleared automatically. In order for any subsequent dump information to be written to flash memory, you must manually clear the dump region. For more information on clearing the dump region, see page 308.

To access dump information, at the `Maintenance#` prompt, enter:

```
Maintenance# uudmp
```

The dump information is displayed on your screen and, if you have configured your communication software to do so, captured to a file. If the dump region is empty, the following appears:

```
No FLASH dump available.
```

/maint/ptdmp <server> <filename>
TFTP System Dump Put

Use this command to `put` (save) the system dump to a TFTP server.

NOTE – If the TFTP server is running SunOS or the Solaris operating system, the specified `ptdmp` file must exist *prior* to executing the `ptdmp` 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 `Maintenance#` prompt, enter:

```
Maintenance# ptdmp server filename
```

Where *server* is the TFTP server IP address or hostname, and *filename* is the target dump file.

/maint/cldmp
Clearing Dump Information

To clear dump information from flash memory, at the `Maintenance#` prompt, enter:

```
Maintenance# cldmp
```

The switch clears the dump region of flash memory and displays the following message:

```
FLASH dump region cleared.
```

If the flash dump region is already clear, the switch displays the following message:

```
FLASH dump region is already clear.
```

/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 Wednesday January 24, 2001...

Boot version 1.0.1

Alteon 184

Rebooted because of console PANIC command.

Booting complete 11:55:01 Wednesday January 24, 2001:
```

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 Thursday Jan 25, 2001. Use /maint/uudmp to
      extract the dump for analysis and /maint/cldmp to
      clear the FLASH region. The region must be cleared
      before another dump can be taken.
```

APPENDIX A

Web OS Syslog Messages

The following syntax is used when outputting syslog messages:

<*Time stamp*><*Log Label*>`Web OS`<*Thread ID*>`:`<*Message*>

where

- <*Timestamp*>

  The time of the message event is displayed in month day hour:minute:second format. For example: `Aug 19 14:20:30`

- <*Log Label*>

  The following types of log messages are recorded: `LOG_EMERG`, `LOG_ALERT`, `LOG_CRIT`, `LOG_ERR`, `LOG_WARNING`, `LOG_NOTICE`, `LOG_INFO`, and `LOG_DEBUG`

- <*Thread ID*>

  This is the software thread that reports the log message. The following thread IDs are recorded: `stp`, `ip`, `slb`, `console`, `telnet`, `vrrp`, `system`, `web server`, `ssh`, and `bgp`

- <*Message*>: The log message

Following is a list of potential syslog messages. To keep this list as short as possible, only <*Thread ID*> and <*Message*> are shown. The messages are sorted by <*Log Label*>.

Where the <*Thread ID*> is listed as *mgmt*, one of the following may be shown: `console`, `telnet`, `web server`, or `ssh`.

LOG_WARNING

FILTER “filter <*filter number*> fired on port <*port number*>, <*source IP address*> -> <*destination IP address*>, [<*ICMP type*>], [<*IP protocol*>], [<*layer-4 ports*>], [<*TCP f1ags*>]”

LOG_ALERT

stp:	own BPDU received from port <port_id>
IP	cannot contact default gateway <ip_address>
vrrp:	received errored advertisement from <ip_address>
vrrp:	received incorrect password from <ip_address>
vrrp:	received incorrect addresses from <ip_address>
vrrp:	received incorrect advertisement interval <seconds> from <ip_address>
slb:	cannot contact real server <ip_address>
slb:	real server <ip_address> has reached maximum connections
gslb:	received update from <ip_address> for unknown remote server <ip_address>
gslb:	received update from <ip_address> for unknown virtual service
gslb:	received update for unknown remote server <ip_address> from <ip_address>
gslb:	received update for unknown service <ip_address:service>
slb:	cannot contact real service <ip_address:real_port>
slb:	real server failure threshold (<threshold>) has been reach for group <group_id>
slb:	real server <ip_address> disabled through configuration
slb:	Virtual Service Pool full. gSvcPool=MAX_SERVICES
bgp:	notification (<reason>) received from <BGP peer ip_address>
bgp:	session with <BGP peer ip_address> failed (<reason>)
vrrp:	Synchronization from non-configured peer <ip_address>
vrrp:	Synchronization from non-configured peer <ip_address> was blocked
dps:	hold down triggered: <ip_address> for <min> minutes
dps:	manual hold down: <ip_address>
syn_atk	SYN attack detected: <count> new half-open sessions per se cond
tcplim	hold down triggered: <ip_address> for <min> minutes

LOG_CRIT

SYSTEM: temperature at sensor <sensor_id> exceeded threshhold

SYSTEM: internal power supply failed

SYSTEM: redundant power supply failed

SSH can't allocate memory in load_MP_INT

LOG_ERR

mgmt: PANIC at <file>:<line> in thread <thread id>

mgmt: VERIFY at <file>:<line> in thread <thread id>

mgmt: ASSERT at <file>:<line> in thread <thread id>

ntp: cannot contact NTP server <ip_address>

ntp: unable to listen to NTP port

isd: unable to listen to BOOTP_SERVER_PORT port

stp: Error: Error writing STG config to FLASH

stp: Error: Error writing config to FLASH

mgmt: Apply not done

mgmt: Save not done

mgmt: " <""apply""|""save""> is issued by another user. Try later"

cli: Error: Error writing %s config to FLASH

cli: New Path Cost for Port <port_id> is invalid

cli: VLAN <vlan_id> has a member port that can not support jumbo frame

cli: PVID <vlan_id> for port <port_id> is not created

cli: RADIUS secret must be 1-32 characters long

cli: Please configure primary RADIUS server address

cli: STP changes can't be applied since STP is OFF

cli: Switch reset is required to turn STP on/off

cli: Trunk group <trunk_id> contains ports with different PVIDs

LOG_ERR (Continued)

cli: Trunk group <trunk_id> has more than <max_trunk_ports> ports

cli: Trunk group <trunk_id> contains no ports but is enabled

cli: Not all ports in trunk group <trunk_id> are in VLAN <vlan_id>

cli: Trunk groups <trunk_id> and <trunk_id> can not share the same port

port_mirr: Port Mirroring changes are not applied

cli: Broadcast address for IP interface <interface_id> is invalid

cli: IP Interfaces <interface_id> and <interface_id> are on the same subnet

cli: Multiple static routes have same destination

cli: Virtual router <vr_id> must have sharing disabled when hotstandby is enabled

cli: Virtual router group must be enabled when hotstandby is enabled

cli: At least one virtual router must be enabled when group is enabled

cli: Virtual router group must have sharing disabled when hotstandby is enabled

cli: Virtual router group must have preemption enabled when hotstandby is enabled

cli: Virtual router <vr_id> must have an IP address

cli: Virtual router <vr_id> cannot have same VRID and VLAN as <vlan_id>

cli: Virtual router <vr_id> cannot have same IP address as <ip_address>

cli: Virtual router <vr_id> corresponding virtual server <server_id> is not enabled

cli: Hot-standby must be enabled when a virtual router has a PIP address

cli: Virtual router <vr_id> IP interface should be <interface_id>

cli: Enabled real server <server_id> has no IP address

cli: Real server <server_id> has same IP address as IP interface <interface_id>

cli: Real server <server_id> has same IP address as switch

cli: Real server <server_id> (Backup for <server_id>) is not enabled

cli: Real server <server_id> has same IP address as virtual server <server_id>

cli: Real server <server_id> has same IP address as real server <server_id>

cli: Real server group <group_id> cannot backup itself

cli: Real server <server_id> cannot be added to same group

LOG_ERR (Continued)

cli: Enabled virtual server <server_id> has no IP address

cli: Virtual server <server_id> has same IP address as IP interface <interface_id>

cli: Virtual server <server_id> has same IP address as switch

cli: Virtual servers <server_id> and <server_id> with same IP address must support same layr3 configuration

cli: Real server <server_id> cannot be backup server for both real server <server_id> and group <group_id>

cli: Virtual server <server_id> has same IP address and vport as virtual server <server_id>

cli: RS <server_id> can't exist for VS <server_id> vport <virtual_port>

cli: Switch port <port_id> has same proxy IP address as port <port_id>

cli: Switch port <port_id> has same IP address as IP interface <interface_id>

cli: A hot-standby port cannot also be an inter-switch port

cli: There must be at least one inter-switch port if any hot-standby port exist

cli: "With VMA, ports 1-8 must all have a PIP if any one does"

cli: Client bindings are not supported with proxy IP addresses

cli: DAM must be turned on or a PIP must be enabled for port <port_id> in order for virtual server to support FTP parsing

cli: Real server <server_id> and group %u cannot both have backups configured

cli: Virtual server <server_id> : port mapping but layer3 bindings

cli: Extracting length has to set to 8 or 16 for cookie rewrite mode

cli: DAM must be turned on or a PIP must be enabled for port <port_id> in order for virtural server <server_id> to support URL parsing

cli: Port filtering must be disabled on port <port_id> in order to support cookie based persistence for virtual server <server_id>

cli: Virtual server <server_id>: port mapping but Direct Access Mode

cli: Virtual server %lu: support nonat IP but not layer 3 bindings

cli: Virtual servers: all that support IP must use same group

cli: Virtual servers <server_id> and <server_id> that include the same real server <server_id> cannot map the same real port or balance UDP

cli: Virtual server <server_id>: UDP service <virtual_port> with out-of-range port number

LOG_ERR (Continued)

cli: Switch cannot support more than <MAX_VIRT_SERVICES> virtual services

cli: Switch cannot support more than <MAX_SMT> real services

cli: Trunk group (<trunk_id>) ports must have same L4 config

cli: Trunk group (<trunk_id>) ports must all have a PIP

cli: DAM must be turned on or a PIP must be enabled for ports <port_id> in order to do URL based redirection

cli: “Two services have same hostname, <host_name>.<domain_name>”

cli: Direct access mode is not supported with default gateway load balancing

cli: SLB Radius secret must be 16 characters long

cli: Dynamic NAT filter <filter_id> must be cached

cli: NAT filter <filter_id> must have same smask and dmask

cli: NAT filter <filter_id> cannot have port ranges

cli: NAT filter <filter_id> must be cached

cli: NAT filter <filter_id> dest range includes VIP <server_id>

cli: NAT filter <filter_id> dest range includes RIP <server_id>

cli: Redirection filter <filter_id> must be cached

cli: Filter with L4 ports configured <port_id> must have IP protocol configured

cli: “For Global SLB, Web server must be moved from TCP port 80”

cli: Remote site <site_id> does not have a primary IP address

cli: Primary and secondary remote site <site_id> switches must differ

cli: Remote sites <site_id> and <site_id> must use different addresses

cli: Remote site <site_id> and real server <server_id> must use different addresses

cli: Remote site <site_id> and virtual server <server_id> must use different addresses

cli: Only <MAX_SLB_SITES> remote servers are allowed per group

cli: Only <MAX_SLB_SERVICES> remote services are supported

cli: Enabled external lookup IP address has no IP address

cli: domain name must be configured

LOG_ERR (Continued)

cli:	Network <static_network_id> has no VIP address
cli:	duplicate default entry
cli:	BGP peer <bgp_peer_id> must have an IP address
cli:	BGP peers <bgp_peer_id> and <bgp_peer_id> have same address
cli:	BGP peer <bgp_peer_id> have same address as IP interface <ip_interface_id>
cli:	BGP peer <bgp_peer_id> IP interface <ip_interface_id> is not enabled
cli:	Filter with ICMP types configured (<icmp_type>) must have IP protocol configure to ICMP
cli:	"Two services have same hostname, <host_name>.<domain_name>"
cli:	Loadbalance string must be added to real server <server_id> in order to enable exclusionary string matching
cli:	intrval input value must be in the range [0-24]
mgmt:	unapplied changes reverted
mgmt:	unsaved changes reverted
mgmt:	Attempting to redirect a previously redirected output
vrrp:	Attempting to redirect a previously redirected output
vrrp:	cfg_sync_tx_putsn: ABORTED
vrrp:	Synchronization TX Error
vrrp:	Synchronization TX connection RESET
vrrp:	Synchronization TX connection TIMEOUT
vrrp:	Synchronization TX connection UNREACEABLE
vrrp:	Synchronization TX connection UNKNOWN CLOSE
vrrp:	Synchronization RX connection RESET
vrrp:	Synchronization RX connection TIMEOUT
vrrp:	Synchronization RX connection UNREACEABLE
vrrp:	Synchronization RX connection UNKNOWN CLOSE
vrrp:	Synchronization connection RCLOSE by peer
vrrp:	Synchronization connection RCLOSE before RX

LOG_ERR (Continued)

vrrp: Synchronization connection early RCLOSE in RX

vrrp: Synchronization connection Wait-For-Close Timeout

vrrp: Synchronization connection Transmit Timeout

vrrp: Synchronization Receive Timeout

vrrp: Synchronization Receive UNKNOWN Timeout

vrrp: Sync transmit in progress … cannot start Sync

vrrp: Sync receive in progress … cannot start Sync

vrrp: Sync already in progress … cannot start Sync

vrrp: Config Sync route find error

vrrp: Config Sync tcp_open error

vrrp: Config Synchronization Timeout - Resuming Console thread

vrrp: “<""apply""|""save""> is issued by another user. Try later”

vrrp: new configuration did not validate (rc =)

vrrp: new configuration did not apply (rc =)

vrrp: new configuration did not save (rc =)

vrrp: Sync config apply error

vrrp: Restoring Current Config

vrrp: Sync rx tcp open error

vrrp: Sync Version/Password Failed-No Version/Password Line

vrrp: Sync Version Failed - peer:%s config:%s

vrrp: Sync Password Failed-Bad Password

vrrp: Sync receive already in progress … cannot start Sync receive

vrrp: Sync transmit in progress … cannot start Sync receive

LOG_NOTICE

system:	internal power supply ok
system:	redundant power supply present and ok
system:	rebooted <last_reset_information>
system:	rebooted <last_reset_information> administrator logged in
mgmt:	boot config block changed
mgmt:	boot image changed
mgmt:	switch reset from CLI
mgmt:	syslog host changed to <ip_address>
mgmt:	syslog host changed to this host
mgmt:	second syslog host changed to <ip_address>
mgmt:	second syslog host changed to this host
mgmt:	Next boot will use active config block
mgmt:	user password changed
mgmt:	SLB operator password changed
mgmt:	L4 operator password changed
mgmt:	operator password changed
mgmt:	SLB administrator password changed
mgmt:	L4 administrator password changed
mgmt:	administrator password changed
ssh:	scp <login_level> login
ssh:	“scp <login_level> <""connection closed""|""idle timeout""|""logout""> ”
mgmt:	RADIUS server timeouts
mgmt:	Failed login attempt via TELNET from host %s
mgmt:	PASSWORD FIX-UP MODE IN USE
mgmt:	<login_level> login on Console
mgmt:	“<login_level> <""idle timeout""|""logout""> from Console”
mgmt:	PANIC command from CLI

LOG_NOTICE (Continued)

port_mirr:	“port mirroring is <""enabled""|""disabled""> ”
vlan:	Default VLAN can not be deleted
mgmt:	<login_level> login from host <ip_address>
mgmt:	“<login_level> <""connection closed""|""idle timeout""|""logout""> from”
IP	“default gateway <ip_address> <""enabled""|""disabled""> ”
IP	default gateway <ip_address> operational
vrrp:	virtual router <ip_address> is now master
vrrp:	virtual router <ip_address> is now backup
slb:	“backup server <ip_address> <""enabled""|""diabled""> for real server <server_id>”
slb:	“backup server <ip_address> <""enabled""|""disabled""> for real server group <group_id>”
slb:	“backup group server <ip_address> <""enabled""|""disabled""> for real server group group_id>”
slb:	“overflow server <ip_address> <""enabled""|""disabled""> for real server <server_id>”
slb:	“overflow server <ip_address> <""enabled""|""disabled""> for real server group <group_id>”
slb:	“overflow group server <ip_address> <""enabled""|""disabled""> for real server group <group_id>”
slb:	real server <ip_address> operational
slb:	real service <ip_address:real_port> operational
slb:	No services are available for Virtual Server <virtual_server>
slb:	Services are available for Virtual Server <virtual_server>
bgp:	session established with <BGP_peer_ip_address>

LOG_INFO

SYSTEM:	bootp response from <ip_address>
mgmt:	new configuration applied
mgmt:	new configuration saved
mgmt:	unsaved changes reverted
mgmt:	Could not revert unsaved changes
mgmt:	" <image1|image2> downloaded from host <ip_address>, file <file_name> <software_version>"
mgmt:	serial EEPROM downloaded from host <ip_address> file <file_name>
ssh:	scp <login_level> login
ssh:	" scp <login_level> <""connection closed""|""idle timeout""|""logout"">"
mgmt:	<login_level> login on Console
mgmt:	" <login_level> <""idle timeout""|""logout""> from Console"
mgmt:	<login_level> login from host <ip_address>
mgmt:	" <login_level> <""connection closed""|""idle timeout""|""logout""> from Telnet/SSH."
ssh:	server key autogen starts
ssh:	server key autogen completes
ssh:	server key autogen timer timeouts
vrrp:	new synch configuration applied
vrrp:	new synch configuration saved
vrrp:	Synchronizing from <host_name>
vrrp:	Synchronizing to <host_name>
vrrp:	Config Synchronization Transmit Successful
vrrp:	Config Synchronization Receive Successful
vrrp:	new configuration VALIDATED

APPENDIX B

Web OS SNMP Agent

The Web OS SNMP agent supports SNMP Version 1. Security is provided through SNMP community strings. The default community strings are “public” for `SNMP GET` operation and “private” for `SNMP SET` operation. The community string can be modified only through the Command Line Interface (CLI). Alteon WebSystems is registered as Vendor 1872. Detailed SNMP MIBs and trap definitions of the Web OS SNMP agent can be found in the following Alteon WebSystems enterprise MIB documents:

- `Altroot.mib` - Alteon product registrations which are returned as sysObjectID.
- `Altswitch.mib` - Alteon enterprise MIB definitions.
- `Alttrap.mib` - Alteon enterprise trap definitions.

Users may specify up to two trap hosts for receiving SNMP Traps. The agent will send the SNMP Trap to the specified hosts when appropriate. Traps will not be sent if there is no host specified.

Web OS SNMP agent supports the following standard MIBs:

- RFC 1213 - MIB II (System, Interface, Address Translation, IP, ICMP, TCP, UDP, SNMP Groups)
- RFC 1573 - MIB II Extension (IFX table)
- RFC 1643 - EtherLike MIB
- RFC 1493 - Bridge MIB
- RFC 1757 - RMON MIB (Statistics, History, Alarm, Event Groups)

Web OS SNMP agent supports the following generic traps as defined in RFC 1215:

- ColdStart
- WarmStart
- LinkDown
- LinkUp
- AuthenticationFailure

The SNMP agent also supports two Spanning Tree traps as defined in RFC 1493:

- NewRoot
- TopologyChange

The following are the enterprise SNMP traps supported in Web OS:

Table 10-6 Web OS-Supported Enterprise SNMP Traps

Trap Name	Description
altSwDefGwUp	Signifies that the default gateway defined is alive.
altSwDefGwDown	Signifies that the default gateway defined is down
altSwDefGwInService	Signifies that the default gateway is up and in service
altSwDefGwNotInService	Signifies that the default gateway is alive but not in service
altSwSlbRealServerUp	Signifies that the real server is up and operational
altSwSlbRealServerDown	Signifies that the real server is down and out of service
altSwSlbRealServerMaxConnReached	Signifies that the real server has reached maximum connections
altSwSlbBkupRealServerAct	Signifies that the backup real server is activated due to availablity of the primary real server
altSwSlbBkupRealServerDeact	Signifies that the backup real server is deactivated due to the primary real server is available
altSwSlbBkupRealServerActOverflow	Signifies that the backup real server is deactivated due to the primary real server is overflowed
altSwSlbBkupRealServerDeactOverflow	Signifies that the backup real server is deactivated due to the primary real server is out from overflow situation
altSwfltFilterFired	Signifies that the packet received on a switch port matches the filter rule
altSwSlbRealServerServiceUp	Signifies that the service port of the real server is up and operational
altSwSlbRealServerServiceDown	Signifies that the service port of the real server is down and out of service
altSwVrrpNewMaster	The newMaster trap indicates that the sending agent has transitioned to 'Master' state.
altSwVrrpNewBackup	The newBackup trap indicates that the sending agent has transitioned to 'Backup' state.
altSwVrrpAuthFailure	A vrrpAuthFailure trap signifies that a packet has been received from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. Implementation of this trap is optional.

Table 10-6 Web OS-Supported Enterprise SNMP Traps

Trap Name	Description
altSwLoginFailure	A altSwLoginFailure trap signifies that someone failed to enter a valid username/password combination.
altSwSlbSynAttack	A altSwSlbSynAttack trap signifies that a SYN attack has been detected.
altSwTcpHoldDown	A altSwTcpHoldDown trap signifies that new TCP connection requests from a particular client will be blocked for a pre-determined amount of time since the rate of new TCP connections from that client has reached a pre-determined threshold.
altSwTempExceedThreshold	A altSwTempExceedThreshold trap signifies that the switch temperature has exceeded maximum safety limits.

APPENDIX C

Performing a Serial Download

You can perform a serial download of the new Web switch software if you are upgrading Web OS directly from any image. Note that a serial download will reset the switch configuration back to its factory defaults.

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)

Use the following procedure to perform a serial upgrade.

1. **Using the serial cable, connect the computer to the switch Console port.**
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:**

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

4. **Turn on the switch power and press <Shift-F> while the switch is first attempting to boot.**

   When performed correctly, the following message appears:

```
Xmodem flash download 1.0.5
To download to flash use xmodem at 57600 baud
Power cycle to end xmodem.
```

5. **Reconfigure your terminal emulation software for the following parameters:**

Parameter	Value
Baud Rate	57,800
Data Bits	8
Parity	None
Stop Bits	1

6. **Set the file transfer mode to Xmodem.**

7. **Transfer the binary firmware image file to the switch.**

   This process can take three or four minutes to complete. When finished, the message “done” will appear on your terminal.

8. **Disconnect the terminal emulation session and reconfigure your terminal emulation software for normal switch connection parameters:**

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

9. **Reconnect the terminal session to the switch.**

10. **Turn the switch power off, and then back on again.**

    The switch should now boot normally.

Glossary

DIP (Destination IP Address)	The destination IP address of a frame.
Dport (Destination Port)	The destination port (application socket: for example, http-80/https-443/DNS-53)
NAT (Network Address Translation)	Any time an IP address is changed from one source IP or destination IP address to another address, network address translation can be said to have taken place. In general, half NAT is when the destination IP or source IP address is changed from one address to another. Full NAT is when both addresses are changed from one address to another. No NAT is when neither source nor destination IP addresses are translated. Virtual server-based load balancing uses half NAT by design, because it translates the destination IP address from the Virtual Server IP address, to that of one of the real servers.
Preemption	In VRRP, preemption will cause a Virtual Router that has a lower priority to go into backup should a peer Virtual Router start advertising with a higher priority.
Priority	In VRRP, the value given to a Virtual Router to determine its ranking with its peer(s). Minimum value is 1 and maximum value is 254. Default is 100. A higher number will win out for master designation.
Proto (Protocol)	The protocol of a frame. Can be any value represented by a 8-bit value in the IP header adherent to the IP specification (for example, TCP, UDP, OSPF, ICMP, and so on.)
Real Server Group	A group of real servers that are associated with a Virtual Server IP address, or a filter.

Redirection or Filter-Based Load Balancing

A type of load balancing that operates differently from virtual server-based load balancing. With this type of load balancing, requests are transparently intercepted and "redirected" to a server group. "Transparently" means that requests are not specifically destined for a Virtual Server IP address that the switch owns. Instead, a filter is configured in the switch. This filter intercepts traffic based on certain IP header criteria and load balances it. Filters can be configured to filter on the SIP/Range (via netmask), DIP/Range (via netmask), Protocol, SPort/Range or DPort/Range. The action on a filter can be Allow, Deny, Redirect to a Server Group, or NAT (translation of either the source IP or destination IP address). In redirection-based load balancing, the destination IP address is not translated to that of one of the real servers. Therefore, redirection-based load balancing is designed to load balance devices that normally operate transparently in your network—such as a firewall, spam filter, or transparent Web cache.

RIP (Real Server)

Real Server IP Address. An IP addresses that the switch load balances to when requests are made to a Virtual Server IP address (VIP).

SIP (Source IP Address)

The source IP address of a frame.

SPort (Source Port)

The source port (application socket: for example, HTTP-80/HTTPS-443/DNS-53).

Tracking

In VRRP, a method to increase the priority of a virtual router and thus master designation (with preemption enabled). Tracking can be very valuable in an active/active configuration.

You can track the following:

- Vrs: Virtual Routers in Master Mode (increments priority by 2 for each)
- Ifs: Active IP interfaces on the Web switch (increments priority by 2 for each)
- Ports: Active ports on the same VLAN (increments priority by 2 for each)
- l4pts: Active Layer 4 Ports, client or server designation (increments priority by 2 for each
- reals: healthy real servers (increments by 2 for each healthy real server)
- hsrp: HSRP announcements heard on a client designated port (increments by 10 for each)

VIP (Virtual Server IP Address)

An IP address that the switch owns and uses to load balance particular service requests (like HTTP) to other servers.

VIR (Virtual Interface Router)

A VRRP address that is an IP interface address shared between two or more virtual routers.

Virtual Router

A shared address between two devices utilizing VRRP, as defined in RFC 2338. One virtual router is associated with an IP interface. This is one of the IP interfaces that the switch is assigned. All IP interfaces on the Alteon Web switches must be in a VLAN. If there is more than one VLAN defined on the Web switch, then the VRRP broadcasts will only be sent out on the VLAN of which the associated IP interface is a member.

Virtual Server Load Balancing

Classic load balancing. Requests destined for a Virtual Server IP address (VIP), which is owned by the switch, are load balanced to a real server contained in the group associated with the VIP. Network address translation is done back and forth, by the switch, as requests come and go.
Frames come to the switch destined for the VIP. The switch then replaces the VIP and with one of the real server IP addresses (RIP's), updates the relevant checksums, and forwards the frame to the server for which it is now destined. This process of replacing the destination IP (VIP) with one of the real server addresses is called half NAT. If the frames were not half NAT'ed to the address of one of the RIPs, a server would receive the frame that was destined for it's MAC address, forcing the packet up to Layer 3. The server would then drop the frame, since the packet would have the DIP of the VIP and not that of the server (RIP).

VRID (Virtual Router Identifier)

In VRRP, a value between 1 and 255 that is used by each virtual router to create its MAC address and identify its peer for which it is sharing this VRRP address. The VRRP MAC address as defined in the RFC is 00-00-5E-00-01-{VRID}. If you have a VRRP address that two switches are sharing, then the VRID number needs to be identical on both switches so each virtual router on each switch knows whom to share with.

VRRP (Virtual Router Redundancy Protocol)

A protocol that acts very similarly to Cisco's proprietary HSRP address sharing protocol. The reason for both of these protocols is so devices have a next hop or default gateway that is always available. Two or more devices sharing an IP interface are either advertising or listening for advertisements. These advertisements are sent via a broadcast message to an address such as 224.0.0.18.
With VRRP, one switch is considered the master and the other the backup. The master is always advertising via the broadcasts. The backup switch is always listening for the broadcasts. Should the master stop advertising, the backup will take over ownership of the VRRP IP and MAC addresses as defined by the specification. The switch announces this change in ownership to the devices around it by way of a Gratuitous ARP, and advertisements. If the backup switch didn't do the Gratuitous ARP the Layer 2 devices attached to the switch would not know that the MAC address had moved in the network. For a more detailed description, refer to RFC 2338.

VSR (Virtual Server Router)

A VRRP address that is a shared Virtual Server IP address. VSR is Alteon WebSystems' proprietary extension to the VRRP specification. The switches must be able to share Virtual Server IP addresses, as well as IP interfaces. If they didn't, the two switches would fight for ownership of the Virtual Server IP address, and the ARP tables in the devices around them would have two ARP entries with the same IP address but different MAC addresses.

Index

Symbols

/ command 45
[ ] 13

A

abbreviating commands (CLI) 48
actio (SLB filtering option) 257
activating optional software 293
active configuration block 168, 299
Active FTP SLB Parsing and Filter Statistics 120
active switch, saving and loading configuration 233
add
    ARP entry 305
    SLB port option 267
    SLB virtual server option 249
addr
    ARP entries 305
    IP route tag 59
Address Resolution Protocol (ARP)
    add, delete entries 305
    address list 305
    interval 182
    statistics 82
administrator account 20, 23
admpw (system option) 176
aging
    STP bridge option 208
    STP information 73
application health checking 243
application redirection 237, 257
    filter states 53
    filters 237
    within real server groups 242
apply (global command) 167
applying configuration changes 167
ARP. *See* Address Resolution Protocol.
ASCII terminal 16
authentication, application health checking 281
autoconfiguration
    duplex mode 28
    link 28, 29
    port speed 28
auto-negotiation 28
    configuring flow control 179
    enable/disable on port 180
    setup 28, 29

B

backup
    SLB real server group option 244
    SLB real server option 239
backup configuration block 168, 299
backup connector (back), Port Menu option 177
backup server activations (SLB statistics) 128
banner (system option) 170
baud rate
    console connection 16, 327
    serial download 328
BGP
    configuration 197
    filters 200
    IP route tag 59
BGP Peers 198
binary 327
binary firmware image 327, 328
binding failure 127
binding table 250
BLOCKING (port state) 73
Boot Options Menu 295
BOOTP 17
    setup (enable/disable) 27
    system option 170

Border Gateway Protocol 59
configuration.....197
BPDU. *See* Bridge Protocol Data Unit.
bridge parameter menu, for STP 206
bridge priority 73
Bridge Protocol Data Unit (BPDU).....73
STP transmission frequency.....208
Bridge Spanning-Tree parameters.....208
broadcast
IP route tag 59
IP route type 58
broadcast IP address 33

C

capture dump information to a file 307
Cisco EtherChannel.....214
clear
ARP entries 305
dump information 308
FDB entry.....303
routing table.....306
Clearing the SLB Statistics.....128
client traffic processing.....266
command (help) 45
Command-Line Interface (CLI) ..... 15 to 21, 23, 43
commands
abbreviations.....48
conventions used in this manual 13
shortcuts.....48
stacking.....48
tab completion.....48
configuration
administrator password 176
apply changes..... 167
default gateway interval, for health checks 184
default gateway IP address 184
dump command 232
effect on Spanning-Tree Protocol 167
Fast Ethernet 177
flow control..... 179
Gigabit Ethernet..... 177
IP broadcast address 183
IP static route 185
Layer 4 administrator password 176
operating mode..... 179
port link speed 179
route cache..... 186
save changes 168
setup command..... 232
switch IP address 183
user password..... 175
view changes..... 167
VLAN default (PVID)..... 178
VLAN IP interface..... 183
VLAN tagging..... 178
configuration block
active 299
backup..... 299
factory 299
selection 299
Configuration Menu..... 165
connecting
via console 16
via Telnet..... 17
connection timeout (Real Server Menu option)..... 250
console port
communication settings 16, 327
connecting 16
serial download settings 328
contacting Alteon WebSystems..... 14
cost
STP information 73
STP port option 209
counters, No Server Available (dropped frames)... 128
cur (system option) 173, 174
current bindings..... 127
customer support 14

D

date
- setup 25
- system option 169

debugging 301
default gateway
- information 52
- interval, for health checks 184
- metrics 203
- round robin, load balancing for 203

default password 20
delete
- ARP entry 305
- FDB entry 303

deny (filtering) 128
designated port. 61
diff (global) command, viewing changes 167
dip (destination IP address for filtering) 258
direct (IP route type) 58
DISABLED (port state) 73
disconnect idle timeout 21
Distributed Site State Protocol (DSSP), setting update interval 269
dmask, destination mask for filtering 258
Domain Name System (DNS)
- health checks 243
- peer site handoffs 269

downloading software 296
dropped frames (No Server Available) counter 128
dump
- configuration command 232
- maintenance 301
- state information 309

duplex mode 28
- link status 51, 71
- setup 28

dynamic routes 306

E

EMS,Alteon EMS 36
emulation software 327, 328
EtherChannel, as used with port trunking 214

F

factory configuration block 299
factory default configuration 21, 23, 24
Fast Ethernet, configuring ports for 177
File Transfer Protocol 119
filter statistics 110
filtered (denied) frames 128
filters, IP address ranges 258
Final Steps 35
first-time configuration ..... 21, 23 to 40
fixed, IP route tag 59
flag field 61
flow control 51, 71
- configuring 179
- setup 28, 29

forwarding database (FDB) 301
- delete entry 303

Forwarding Database Information Menu 67
Forwarding Database Menu 303
forwarding state (FWD) 68, 73, 78
FTP server health checks 243
FTP SLB Maintenance Statistics 121
FTP SLB Statistics Dump 121
full-duplex 28
fwd (STP bridge option) 208
FwdDel (forward delay), bridge port 73

G

gig (Port Menu option) 177
Gigabit Ethernet, configuration 177
Greenwich Mean Time (GMT) 174
gtcfg (TFTP load command) 233

H

half-duplex 28
hash metric 245
health checking (SLB real server group option) 243
health checks 239
- default gateway interval, retries 184
- layer information 53
- parameters for most protocols 243
- redirection (rport) 257

hello, STP information 73
help 45

Hot Standby Router Protocol (HSRP)
priority increment value for L4 client ports 226
use with VRRP.....220, 224
VRRP priority increment value 226, 227
HP-OpenView 15
HSRP. *See* Hot Standby Router Protocol.
HTTP
application health checks.....243
redirects (Global SLB option) 270
system option.....171

I

ICMP
IP route tag 59
Layer 3 health checks.....243
idle timeout 21
IEEE standards
802.1d Spanning-Tree Protocol 72, 206
image
downloading 296
software, selecting 297
IMAP server health checks.....243
imask (IP address mask) 281
incorrect VIPs (statistic).....127
incorrect Vports (dropped frames counter).....128
indirect (IP route type).....58
Information Menu 49
intr, SLB real server option 239
IP address.....32
ARP information 60
BOOTP.....17
configuring default gateway.....184
filter ranges.....258
IP interface 32
local route cache ranges.....187
Telnet.....17
IP address mask (mmask).....170
IP address mask for SLB.....281
IP configuration via setup 32
IP forwarding.....201
IP forwarding information.....52
IP Forwarding Menu.....186
IP Information Menu 52, 76
IP interface 32, 58, 183
broadcast address (broad)..... 183
configuring address 183
configuring VLANs 183
information 52
IP route tag 59
priority increment value (ifs) for VRRP 226
IP Port Menu..... 201
IP Route Manipulation Menu..... 306
IP routing..... 32
tag parameters 59
IP Static Route Menu 185
IP subnet mask 32
IP, switch processor statistics for..... 82

J

jumbo frames, setup..... 30

L

l4apw (L4 administrator system option)..... 176
Layer 4, administrator account..... 20
LEARNING (port state) 73
least connections (SLB Real Server metric)..... 245
licence certificate..... 293
license password..... 293
lines (display option) 46
Link 51
speed, configuring..... 179
Link status 51
command..... 71
duplex mode..... 51, 71
port speed 51, 71
Link Status Information 71
linkt (SNMP option) 211
LISTENING (port state)..... 73
lmask (routing option)..... 52
lnet (routing option)..... 52
local (IP route type) 58
local route cache, IP address ranges for..... 187

M

MAC (media access control) address.. 50, 60, 67, 69, 293, 303
  switch location 17
Main Menu 43
  Command-Line Interface (CLI) 21
  summary..... 44
Maintenance Menu 301
Management Processor (MP)..... 304
  display MAC address 50, 69
manual style conventions 13
martian
  IP route tag (filtered)..... 59
  IP route type (filtered out)..... 58
mask, IP interface subnet address..... 183
MaxAge (STP information)..... 73
mcon (maximum connections)..... 128, 244
  SLB real server option..... 238
media access control. *See* MAC address.
metrc (SLB real server group option) 242
metrics, SLB..... 245
minimum misses (SLB real server metric) 245
Miscellaneous Debug Menu 304
mmask
  IP address mask for SLB..... 281
  system option 170
mnet
  management traffic IP address for SLB..... 281
  system option 170
monitor port..... 212
MP. *See* Management Processor.
multicast
  IP route tag 59
  IP route type..... 58
mxage (STP bridge option)..... 208

N

Network Address Translation (NAT)
  filter action 257
network management..... 15
non TCP/IP frames 127
notice 170

O

online help 45
operating mode, configuring 179
Operations Menu 287
Operations-Level Port Options..... 289
operations-level SLB options 290
operations-level VRRP options 291, 292
optional software 52, 79
  activating 293
  removing..... 294
OSPF Database Information 64
OSPF general 62
OSPF General Information 63
OSPF Host Entry Configuration Menu..... 196
OSPF Information..... 62
OSPF Information Route Codes 66
OSPF Interface Configuration Menu..... 193
OSPF Summary Range Configuration Menu 192
OSPF Virtual Link Configuration Menu 195
ospf/ip/aindex/host..... 190
overflow server activations 128
overflow servers 239

P

panic
  command 309
  switch (and Maintenance Menu option)..... 301
parameters
  tag..... 59
  type 58
Passive FTP SLB Parsing Statistics 120
password
  administrator account..... 20
  default 20
  L4 administrator account..... 20
  user account 19
  VRRP authentication 225
passwords 19
persistent bindings
  real server 250
ping..... 45, 238
poisoned reverse, as used with split horizon 188
POP3, server health checks..... 243
port configuration 177
port flow control. *See* flow control.

Port Menu
configuration options 177
configuring Fast Ethernet....177
configuring Gigabit Ethernet (gig) 177
Port number....71
port speed....51, 71
auto-sense....28
setup....28
port states
UNK (unknown) 68
port trunking
description....214
ports
configuration....27
disabling (temporarily) 180
information....75
IP status....52
membership of the VLAN 51, 74
priority....73
SLB state information 53
STP port priority....209
VLAN ID....51, 75
preemption
assuming VRRP master routing authority 219
virtual router....218, 222
preferred connector (pref), Port Menu option 177
priority (STP port option)....209
proxies, IP address translation 240
proxy IP address (PIP)....53
ptcfg (TFTP save command) 233
PVID (port VLAN ID)....51, 75
pwd 45

Q

quiet (screen display option)....46

R

RADIUS, server authentication 243, 281
read community string (SNMP option)....211
real server groups
combining servers into 242
statistics 109
real server SLB configuration....237
real servers
backup....244
menu options....238
priority increment value (reals) for VRRP 226
SLB state information....53
statistics....108
reboot....301, 309
receive flow control 28, 29, 179
redir (SLB filtering option)....257
reference ports....68
referenced port 61
remote site servers 240
removing optional software 294
reset key combination 301
restarting switch setup....25
restr (SLB real server UDP option)....239
retry
health checks for default gateway....184
SLB real server option....239
rip, IP route tag 59
RIP. *See* Routing Information Protocol.
rmkey....294
round robin, as used in gateway load balancing 203
roundrobin, SLB Real Server metric....245
route
cache configuration....186
switch statistics for Route protocol....82
Routing Information Protocol (RIP) 59
options 188
poisoned reverse 188
split horizon 188
version 1 parameters....188
Routing Information Protocol Menu 188
RTSP SLB Statistics 122
rx flow control 28, 29

S

save (global command) 168
noback option 168
save command 299
serial cable 16
serial download 327
Server Load Balancing
client traffic processing 266
information 53
menu options 236
metrics 245
operations-level options 290
port options 267
real server group options 244
real server weights 238
server traffic processing 266
Server Load Balancing Maintenance Statistics Menu .. 118, 119, 127
Server Load Balancing Metrics 245
server port mapping 53
server traffic processing 266
Session Binding Table 239
session identifier 247
setup command, configuration 232
setup facility 21, 23
BOOTP 27
duplex mode 28
IP configuration 32
IP subnet mask 32
jumbo frames 30
port auto-negotiation mode 28, 29
port configuration 27
port flow control 28, 29
port speed 28
restarting 25
Spanning-Tree Protocol 27
starting 24
stopping 25
system date 25
system time 26
VLAN name 30
VLAN port numbers 31
VLAN tagging 29
VLANs 30
shortcuts (CLI) 48
SIP (source IP address for filtering) 258
SLB Information 53
smask, source mask for filtering 258
smtp 170
SMTP server health checks 243
snap traces, buffer 304
SNMP 15, 82
HP-OpenView 15
IP route tag 59
menu options 211
set and get access 211
SNMP Agent 323
SNMP Support 36
software
image 296
image file and version 50, 69
license 293
Spanning-Tree Protocol 78, 167
bridge aging option 208
bridge parameters 208
bridge priority 73
port cost option 209
port priority option 209
root bridge 73, 208
setup (on/off) 27
switch reset effect 299
split horizon 188
stacking commands (CLI) 48
starting switch setup 24
state (STP information) 73
state information, client system 250
static, IP route tag 59
static (IP route tag) 59
statistics
ARP 82
Statistics Menu 81
stopping switch setup 25
subnet address maskconfiguration
IP subnet address 183
subnet mask 32
subnets 32
IP interface 183
switch
name and location 50, 69
resetting 299
Switch Processor (SP) 304
display trace buffer 304
swkey 293
synchronization, VRRP switch 278, 290

system
contact (SNMP option) 211
date and time 50, 69
information 69
location (SNMP option) 211
System Maintenance Menu 302
system options
admpw (administrator password) 176
BOOTP 170
cur (current system parameters) 173, 174
date 169
HTTP access 171
l4apw (Layer 4 administrator password) 176
login banner 170
mmask 170
mnet 170
time 170
usrpw (user password) 175
wport 170
system parameters, current 173, 174

T

tab completion (CLI) 48
TCP 82
fragments 127, 247
health checking using 239
health checks 243
source and destination ports 256
Telnet 17
BOOTP 17
configuring switches using 232
Telnet support 36
terminal emulation 16
text conventions 13
TFTP 296
PUT and GET commands 233
time
setup 26
system option 170
timeouts, idle connection 21
time-to-live, DNS response (global SLB menu option) 268
tnet 170
trace buffer 304
Switch Processor 304
traceroute 45
tracking criteria 219
transmit flow control 28, 29, 179
transparent proxies, when used for NAT 257
Trunk Group Information 78
ttl (global SLB menu option) 268
tx flow control 28, 29
type parameters 58
typographic conventions, manual 13
tzone 174

U

UDP 82
datagrams 127
server status using 239
source and destination ports 256
unknown (UNK) port state 68
Unscheduled System Dump 309
upgrade, switch software 296
URL for health checks 53
User Access Control Configuration 175
user account 19
usrpw (system option) 175
Uuencode Flash Dump 307

V

verbose 46
vip 189
IP route tag 59
virtual IP address (VIP) 53
virtual port state, SLB information about 53
Virtual Router Redundancy Protocol (VRRP)
authentication parameters for IP interfaces 225
configuration menu options 215, 228
operations-level options 291, 292
password, authentication 225
priority tracking options 198, 200, 219
virtual router options 217, 222
virtual routers
description 216, 221
HSRP failover 220, 224
HSRP priority increment value 226, 227
increasing priority level of 219
master preemption (prio) 218, 222
priority increment values (vrs) for VRRP 226
virtual servers 244
SLB state information 53
statistics 110

VLAN tagging
port configuration 178
port restrictions..... 205
setup 29
VLANs..... 32
ARP entry information 60
information 74
interface 33
name 51, 74
name setup..... 30
port membership..... 51, 74
port numbers 31
setting default number (PVID)..... 178
setup 30
tagging 29, 51, 75, 205
VRID (virtual router ID) 217, 221
VRRP Information 77
VRRP Interface Configuration..... 225
VRRP Tracking Configuration..... 226

W

WAP SLB Statistics..... 124
watchdog timer..... 301
web-based management interface 15
weights
for SLB real servers..... 238, 246
setting virtual router priority values 226
wport..... 170
write community string (SNMP option) 211

X

Xmodem..... 328