Web OS Switch Software

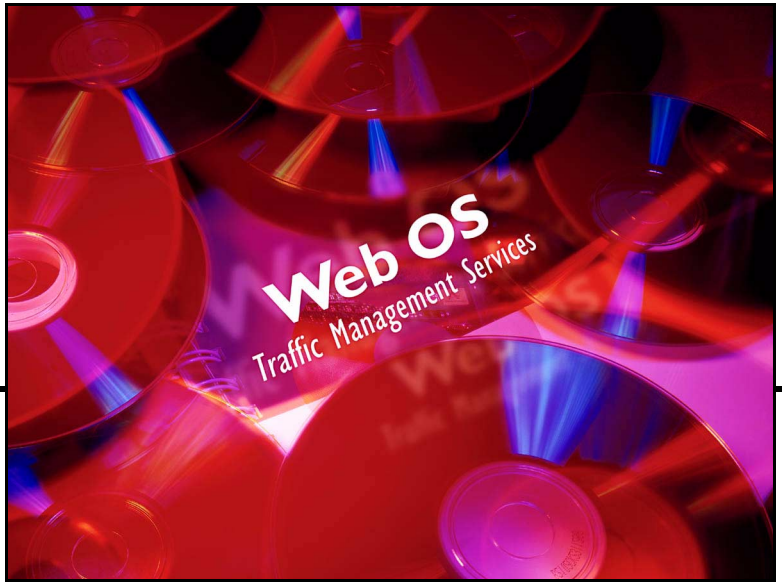

7.0 Command Reference

Part Number: 050066, Revision B, July 2000

Alteon WebSystems
Web Speed for e-Business

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

Contents

Preface

This manual describes how to configure and use the Web OS Release 7.0 software included in the Alteon WebSystems family of switches.

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

Who Should Use This Book

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

How This Book Is Organized

Chapter 1, "Web OS Software Features," provides an overview of the major features included in this release of the switch software.

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

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

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

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

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

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

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

Chapter 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.

Chapter 11, "Troubleshooting," describes switch configuration troubleshooting techniques.

Typographic Conventions

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

Table 1 Typographic Conventions

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

Contacting Alteon WebSystems

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

- URL for Alteon WebSystems Online:

  `http://www.alteonwebsystems.com`

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

- E-mail access:

  `support@alteonwebsystems.com`

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

- Telephone access to Alteon WebSystems Customer Support:

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

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

- Telephone access to Alteon WebSystems Sales:

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

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

CHAPTER 1

Web OS Software Features

The capabilities provided by Web OS Release 7.0 software include those documented with Release 5.2 for Alteon 180 and ACEdirector switches, plus many new features implemented for the new Alteon 700 series switch product family.

The capabilities of the Alteon 700 series switch are based on a distributed processing architecture, enabling the switch to deliver performance far beyond the bounds of traditional, centralized processor switch architectures. Traffic control services are distributed across multiple processors with substantial hardware assist. Background tasks such as switch management, routing updates, and running the user interface are implemented in separate management processors.

This chapter provides an overview of Web OS 7.0 features.

Layer 1 Features

Port Link Characteristics

- 100 Mbps ports support half-and full-duplex operations and 802.3u 10/100 auto-negotiation
- 1000 Mbps ports support 802.3z compliant full-duplex operations with asymmetric flow control

Layer 2 Features

- Fast Ethernet and Gigabit Ethernet ports support the same feature set
- Architectural support for up to 64,000 MAC addresses

Port Trunk Groups

Ports in a trunk group combine their bandwidth to create a single, larger virtual link. Web OS 7 supports EtherChannel-compatible trunk groups, enabling link-level redundancy and load sharing with other EtherChannel-compatible devices.

Release 7 support enables the following port trunking capabilities:

- Up to four trunk groups can be configured per switch
- Up to four ports can be trunked together to form a single virtual link with bandwidth at 100 Mb per second
- IP Session ID hashing for IP addresses
- MAC SA/DA hashing for non-IP traffic
- Trunk groups are inherently fault tolerant: the trunk is active as long as any of its ports are available
- Traffic on the trunk is statistically load balanced between the ports in the link
- Trunk connections support third-party devices such as Cisco routers and switches with EtherChannel technology, and Sun's Quad Fast Ethernet adapter

VLANs

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

Release 7.0 support enables the following VLAN capabilities:

- IEEE 802.1Q tagging (4K external, 4K internal) allows multiple VLANs per port and provides standards-based VLAN support for Ethernet systems
- Port-based VLAN PVIDs: Up to 8 PVIDs can be configured per port, enabling the network administrator to create separate VLANs for different packet types, e.g., IP, IPX, etc.
- Up to 1024 unique VLANs can be configured per switch. (VLANs can be assigned numbers from 1-4094)

Spanning Tree Support

When multiple paths exist, Spanning Tree configures the network so that a switch uses only the most efficient path. In previous releases of the Web OS software, a single Spanning Tree was allowed per switch. Running Web OS 7.0, you can configure up to 256 Spanning Trees on an Alteon 700 series switch. The default Spanning Tree (1) can support up to 1024 VLANs. All other Spanning Trees can support only one VLAN. A VLAN may belong only to one Spanning Tree.

RMON Lite Support

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

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

Layer 3 Features

IP Routing

IP Routing allows the network administrator to seamlessly connect server IP subnets to the rest of the backbone network, using a combination of configurable IP switch interfaces and IP routing options. The IP Routing feature enhances Alteon WebSystems' server switching solution in the following ways:

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

Virtual Router Redundancy Protocol

Virtual Router Redundancy Protocol (VRRP) support on Alteon WebSystems' 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.

NOTE – Alteon WebSystems has extended VRRP to include virtual servers as well, allowing for full active/active redundancy between its Layer 4 switches.

Layer 4 Features

Running Web OS software, the Alteon 700 series offers local and global server load balancing, application redirection, non-server (such as firewall, router) load balancing, active-active high availability configurations, bandwidth management, and server security services.

Server Load Balancing

With Server Load Balancing, your Alteon 700 series switch is aware of the shared services provided by your server farm. The switch can then balance user session traffic among the available servers. For even greater control, traffic is distributed according to a variety of user-selectable metrics.

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

- TCP and UDP load balancing
- Bi-directional session ID substitution
- MaxCons, back-up, and overflow server support
- Round Robin and connection-based load balancing
- Server static weighting
- Hash and Min-Misses load balancing

NOTE – URL-based server load balancing and Web cache redirection will not be supported in the initial Web OS 7.0 release.

Health Checks

The switch can perform health checks at various levels. This includes checking the Layer 3 connectivity using ICMP Ping. Layer 4 connectivity is checked by sending a TCP connection request to the server. The next level of health check supported is checking the retrieval of the actual content from various applications. Content-intelligent health checks are performed for DNS, FTP, HTTP, NNTP, POP3, IMAP, SMTP, and RADIUS services. If any server in a server pool fails, the remaining servers continue to provide access to vital applications and data. The failed server can be brought back up without interrupting access to services. As users are added and the server pool's capabilities are saturated, new servers can be added to the pool transparently.

Filtering

Alteon 700 series switches support line rate filtering for up to 1,024 filters per port, giving network administrators a powerful tool to protect their server networks. Switch-wide filtering rules can be defined on each Alteon 700 series switch, with any or all rules applied to each port. Each filter can forward, drop, or redirect packets and can optionally log results, based on any combination of the following user-specified criteria:

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

Application Redirection

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

Other Features

Switch Components

- Line-card module (LCM) configuration “memory” is provide on a per-slot basis, via parameters stored on the management processor module (MPM).
  - When LCMs are plugged into a slot, their type is recognized.
  - Each LCM type has a default configuration stored for it. The default can be modified on a slot-by-slot basis.
  - Operation can be set so that all new cards immediately take on the default configuration of the latest known modified configuration for the LCM type. Taking on the latest configuration is the default behavior.
- Switch-level configuration “moves” with MPMs. If an MPM in switch A is removed and inserted in the switch “B” chassis, switch “B” will take on the same configuration parameters that box “A” previously had, provided the configuration parameters have been saved.

Comprehensive Network Management

Network managers can configure and monitor all Alteon 700 series functions via the Web OS BBI (Browser-Based Interface), SNMP applications, and a command-line interface (CLI) accessed from the console port or via Telnet. Four levels of password protection are provided, to allow switch configuration changes and to view switch information and statistics.

- Command-Line Interface (CLI) enhancements include a Setup facility, command-line retrieval and editing capability, and tab completion function for commands and options. Aliases for real servers and real server groups are also supported, making it easier to identify them on information and statistics screens.
- Web OS Browser-based Interface (BBI) provides direct browser-to-switch interaction for switch configuration and monitoring.

The Alteon 700 series supports a private MIB and four groups of RMON on every port. The Alteon 700 series management interface is integrated with HP OpenView 5.0 under UNIX (HPUX, Solaris) and Windows NT.

SNMP MIB Support

The SNMP agent for Alteon WebSystems' switches supports the following standard Management Interface Bases (MIBs): RFC 1213 MIB-II, RFC 1493 Bridge MIB, RFC 1643 Ethernet-like MIB, RFC 1573 Interface Extensions MIB, RFC 1724 RIP2 MIB, RFC 1757 RMON (Groups 1-4) MIB, and RFC 2037 Entity MIB.

Security is provided through SNMP community strings that can be modified only through the Command Line Interface (CLI). The default community strings are "public" for SNMP GET operations and "private" for SNMP SET operations.

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

RFC 1573 Interface Extension MIB Compliance

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

Web OS supports the RFC 1573 MIB. This IF Extensions MIB allows for higher speed networking environments, providing 64-bit counters on many MIB-II statistics, plus roll-over counters for 32-bit counters.

Server Dual Homing

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

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

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

CHAPTER 2

The Command-Line Interface

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

Your Alteon WebSystems 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.

Connecting to the Switch

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

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

Establishing a Console Connection

Requirements

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

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

  Table 2 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 (see page 27).

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 33)

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 61).

Running Telnet

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

telnet *IP-address*

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

If you have trouble making a Telnet connection to the switch, refer to Chapter 11, "Troubleshooting."

Entering Passwords

Once you are connected to the switch via local console or Telnet, you are prompted to enter a password. There are four levels of access to the switch:

- User
- Administrator
- Layer 4 administrator
- Layer 4 operator

Each level has a different password and is granted different access privileges.

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 43.

The User Account

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

The Administrator Account

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

Layer 4 Administrator Account

The Layer 4 administrator has limited control of the switch. He or she 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`.

Layer 4 Operator Account

The Layer 4 Operator has limited control of the switch. He or she can view all switch information and statistics, and make temporary "operational" changes (`/oper/slb`) to the Server Load Balancing menus. The default password for the Layer 4 administrator account is `l4oper`.

Accessing the Switch

To enable better switch management and user accountability, three levels or *classes* of user access have been implemented on the switch. Levels of access to CLI and 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; that is they cannot change anything on the switch. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.
- Operators can only effect temporary changes on the switch; these changes will be lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the switch, operators cannot severely impact switch operation.
- Administrators are the only ones that may make permanent changes to the switch configuration; that is, changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the switch. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

Access to switch functions is controlled through the use of unique surnames and passwords. Once you are connected to the switch via local console or Telnet, you are prompted to enter a password. The default usernames/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 43.

Table 3 User Access Levels

User Account	Description and Tasks Performed	Password
User	The user has no direct responsibility for switch management. He or she can view switch status information and statistics, but cannot make any configuration changes.	`user`
Administrator	The superuser administrator has complete access to all menus, information, and configuration commands on the switch, including the ability to change both the user and administrator passwords. Access includes "`oper`" and "`l4admin`" functions.	`admin`
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`
Layer 4 Administrator	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.	`l4admin`

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 vs. Setup

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

The following figure shows the Main Menu with administrator privileges.

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

Figure 1 Administrator Main Menu

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 Chapter 4, “Menu Basics.”

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 151.

CHAPTER 3

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

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

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

```
Enter Password:
```

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

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

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

NOTE – If the default `admin` login is unsuccessful, or if the administrator Main Menu 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 245.

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, IP address and subnet mask, Spanning Tree,
 Port and Link characteristics, and VLANs.
 [Type Ctrl-C to abort "Set Up"]
------------------------------------------------------------------

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

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

   If you decide not to configure VLANs during this session, you can configure them later using the configuration menus, or by restarting the Setup facility. For more information on VLANs issues, see "Setup Part 3: VLANs" on page 38.

   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 [5]:
   ```

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

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

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

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

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

```
Enter year [00]:
```

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

The system displays the date and time settings:

```
System clock set to 13:56:52 Fri May 5, 2000.
```

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

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

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

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

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

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

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

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

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

The system displays the date and time settings:

```
System clock set to 13:56:52 Fri May 5, 2000.
```

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

```
BootP Option [must DISABLE to configure IP]:
Current BOOTP: disabled
Enter new BOOTP [d/e]:
```

If available on your network, a BOOTP server can supply the switch with IP parameters so that you do not have to enter them manually. BOOTP must be disabled however, before the system will prompt for IP parameters.

Enter **d** to disable the use of BOOTP, or enter **e** to enable the use of BOOTP. To keep the current setting, press <Enter>.

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

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

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

Setup Part 2: Port Configuration

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

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

```
Port Configuration:

Enter port (e.g. b12): a1
```

If you wish to change settings for individual ports, enter the number of the port you wish to configure. To skip port configuration, press <Enter> without specifying any port and go to "Setup Part 3: VLANs" on page 38.

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

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

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

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

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

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

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

4. **If appropriate, configure the link mode.**

   The system prompts for the link mode, which can be full duplex or half duplex:

```
Link Mode:
Current Port A1 mode setting: none
Pending new mode setting: any
Enter new mode [full|half|any]: 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>.

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

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

```
Link Autonegotiation:
Current Port A1 autonegotiation:  off
Pending new autonegotiation: on
Enter new value [on|off]: on
```

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

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

   To keep the current setting, press <Enter>.

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

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

   To keep the current setting, press <Enter>.

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

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

   To keep the current setting, press <Enter>.

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

9. **Select whether to forward or discard tagged frames:**

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

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

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

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

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

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

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

Setup Part 3: VLANs

If you chose to skip VLANs configuration back in Part 1, skip to "Setup Part 4: IP Configuration" on page 39.

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

```
VLAN Config:
Enter VLAN number from 1 to 4094, NULL at end:
```

If you wish to change settings for individual VLANs, enter the number of the VLAN you wish to configure. To skip VLAN configuration, press <Enter> without typing a VLAN number and go to "Setup Part 4: IP Configuration" on page 39.

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

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

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

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

```
Define Ports in VLAN:
Current VLAN 1:  A1-16 B1-15 C1-16 D1-16 T1-T8
Enter ports one per line, NULL at end:
```

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

```
>
```

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

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

   The default spanning tree group index is 1. All VLANs must belong to a spanning tree group.

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

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

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

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

Setup Part 4: IP Configuration

If BOOTP was enabled back in Part 1, skip to "Setup Part 5: Final Steps" on page 42. Otherwise, if you disabled BOOTP, the system prompts for IP parameters.

IP Interfaces

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

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

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

```
IP Config:

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

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 41.

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

   To keep the current setting, press <Enter>.

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

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

   To keep the current setting, press <Enter>.

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

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

   To keep the current setting, press <Enter>.

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

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

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

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

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

6. **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: [1-1024]
```

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

Default Gateways

1. **At the prompt, pick a default gateway to configure, or skip default gateway configuration:**

```
IP default gateways:
Enter default gateway number: [1-4]
```

Enter the number for the default gateway to be configured. To skip default gateway configuration, press <Enter> without typing a gateway number and go to "IP Routing" on page 41.

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

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

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

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

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

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

```
Enter default gateway number: [1-4]
```

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

IP Routing

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

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

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

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

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

Setup Part 5: Final Steps

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Setting Passwords

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

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

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

Changing the Default Administrator Password

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

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

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

```
Main# cfg
```

The Configuration Menu is displayed.

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

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

```
>> Configuration# sys
```

The System Menu is displayed.

```
[System Menu]
      date    - Set system date
      time    - Set system time
      idle    - Set timeout for idle CLI sessions
      bootp   - Enable/disable use of BOOTP
      snmp    - Enable/disable SNMP management access
      web     - Enable/disable Web management access
      wport   - Set Web server port number
      mnet    - Set management network
      mmask   - Set management netmask
      banner  - Set login banner
      usrpw   - Set user password
      admpw   - Set administrator password
      l4apw   - Set L4 administrator password
      l4opw   - Set L4 operator password
      current - Display current system-wide configuration
```

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

```
System# 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 by entering `usrpw` at the `System#` prompt.**

```
System# usrpw
```

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

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

```
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 (see "SLB Configuration" on page 185)

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
```

The System Menu is displayed.

```
[System Menu]
      date    - Set system date
      time    - Set system time
      idle    - Set timeout for idle CLI sessions
      bootp   - Enable/disable use of BOOTP
      snmp    - Enable/disable SNMP management access
      web     - Enable/disable Web management access
      wport   - Set Web server port number
      mnet    - Set management network
      mmask   - Set management netmask
      banner  - Set login banner
      usrpw   - Set user password
      admpw   - Set administrator password
      l4apw   - Set L4 administrator password
      l4opw   - Set L4 operator password
      current - Display current system-wide configuration
```

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
```

Changing the Default Layer 4 Operator Password

The Layer 4 operator has limited control of the switch. Through a Layer 4 operator account, you can view all switch information and statistics, but can make only temporary, "operational" changes to the Server Load Balancing menus (see "Operations-Level SLB Options" on page 236.)

The default password for the Layer 4 operator account is `l4oper`. 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
```

   The System Menu is displayed.

```
[System Menu]
      date    - Set system date
      time    - Set system time
      idle    - Set timeout for idle CLI sessions
      bootp   - Enable/disable use of BOOTP
      snmp    - Enable/disable SNMP management access
      web     - Enable/disable Web management access
      wport   - Set Web server port number
      mnet    - Set management network
      mmask   - Set management netmask
      banner  - Set login banner
      usrpw   - Set user password
      admpw   - Set administrator password
      l4apw   - Set L4 administrator password
      l4opw   - Set L4 operator password
      current - Display current system-wide configuration
```

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

```
System# l4opw
```

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

```
Changing L4 OPERATOR 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 operator password at the prompt:**

```
Enter new L4 operator password:
```

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

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

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

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

CLI vs. Setup

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

The following figure shows the Main Menu with administrator privileges.

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

Figure 2 Administrator Main Menu

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

CHAPTER 4

Menu Basics

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

To make the CLI easy to use, the various commands have been logically grouped into a series of menus and 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. Figure 3 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]
```

Figure 3 Administrator Main Menu

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 FLASH memory.

- **Operations Command Menu**

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

- **Boot Options Menu**

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

- **Maintenance Menu**

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

Global Commands

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

Table 4 Global Commands

Command	Action
? *command*	Provides more information about a specific command on the current menu. When used without the *command* parameter, a summary of the global commands is displayed.
.	Display the current menu.
..	Go up one level in the menu structure.
/	If placed at the beginning of a command, go to the Main Menu. Otherwise, this is used to separate multiple commands placed on the same line.
diff	Show any pending configuration changes.
apply	Apply pending configuration changes.
save	Write configuration changes to non-volatile flash memory.
exit	Exit from the command-line interface and log out.
ping	Use this command to verify station-to-station connectivity across the network. The format is as follows: **ping** *address* [*tries* [*delay*]] Where *address* is the hostname or IP address of the device, *tries* (optional) is the number of attempts (1-32), and *delay* (optional) is the number of milliseconds between attempts. The DNS parameters must be configured if specifying hostnames (see "Domain Name System Configuration" on page 183).
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.

Table 4 Global Commands

Command	Action
pwd	Display the command path used to reach the current menu.
lines *n*	Set the number of lines (*n*) that display on the screen at one time; the default is 24 lines. When used without a value, the current setting is displayed.
verbose *n*	Sets the level of information displayed on the screen: **0** = Quiet: Nothing appears 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 5 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.

Table 5 Command-Line History and Editing Options

Option	Description
<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 5

The Information Menu

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

/info

Information Menu

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

The information provided by each menu option is briefly described in Table 6 on page 58, followed by examples of each information screen.

Table 6 Information Menu Options (/info)

Command Syntax and Usage
sys Displays system information, including ■ System date and time ■ Switch model name and number ■ Switch name and location ■ Time of last boot ■ Base MAC address of the switch management processor ■ IP address of IP interface #1 ■ Hardware version and part number ■ Software image file and version number ■ Configuration name ■ Log-in banner, if one is configured For more information, see page 61.
env Displays the switch temperature and fan status. For more information, see page 62.
log Displays 10 most recent syslog messages. For more information, see page 63.
link Displays configuration information about each port, including: ■ Slot letter (A-D on Alteon 708; A-H on Alteon 714) ■ 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 more information, see page 64.
port Displays port status information, including : ■ Slot letter (A-D on Alteon 708; A-H on Alteon 714) ■ Port number ■ Whether the port forwards or discards tagged frames ■ Whether the port forwards or discards untagged frames ■ Whether the port forwards or discards priority-tagged frames ■ Whether RMON is enabled (e) or disabled (d) ■ Port VLAN ID (PVID) ■ Port name For more information, see page 65.

Table 6 Information Menu Options (/info)

Command Syntax and Usage
slot Displays the slot status for LCM and SFM slots. For more information, see page 66.
vlan Displays VLAN configuration information, including: ■ VLAN Number ■ VLAN Name ■ Status ■ Port membership of the VLAN For more information, see page 67.
stp Displays Spanning Tree Protocol status, and the following STP bridge information: ■ Priority ■ Hello interval ■ Maximum age value ■ Forwarding delay ■ Aging time Also displays the following port-specific STP information: ■ Port number and priority ■ Cost ■ State For more information, see page 68.
trunk When trunk groups are configured, you can view the state of each port in the various trunk groups. For more information, see page 70.
arp Displays the Address Resolution Protocol (ARP) Information Menu. For more information, see page 71.
fdb Displays the Forwarding Database Information Menu. For more information, see page 73.

Table 6 Information Menu Options (/info)

Command Syntax and Usage
ip Displays the IP Information Menu. IP information includes: ■ For each IP interface: Interface number, IP address, subnet mask, broadcast address, VLAN number, class of service, and operational status ■ For each Default gateway: 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 For more information, see page 75.
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 more information, see page 76.
slb Displays the Server Load Balancing Information Menu. For more information, see page 79.
vrrp Displays VRRP information. For more information, see page 82.
swkey Displays a list of all the optional software packages which have been activated or installed on your switch. For more information, see page 83.
dump Dumps all switch information available from the Information Menu (10K or more, depending on your configuration). If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump commands. For more information, see page 83.

/info/sys
System Information

```
System Information at 14:40:35 Wed May 17, 2000

Alteon 708
sysName:     Main-switch
sysLocation: IT Lab
Last boot: 10:21:44 Wed May 17, 2000 (reset from console)

Switch base MAC address: 00:22:33:44:55:00
Interface 1 MAC address: 00:22:33:44:55:05  IP address: 123.255.224.10
Hardware Revision: A
Hardware Part No: 210019P1
Software Version 7.0.30(FLASH image2), active configuration.
```

System information includes:

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

/info/env
Environment Information

```
Thermal MP:     36C

Thermal Fan:    20C

Thermal SFM  1: 26C
             2: 24C

Thermal LCM  1: 33.5C
             2: 37C

Power 1: Present
      2: Not Present

Fan: Present @ approx. 2050 RPM
```

Environment information includes:

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

/info/log
Show Last 10 Syslog Messages

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

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

/info/link
Link Status Information

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

Use this command to display link status information about each port on an Alteon 700 Series switch slot, including:

- `Slot`: LCM Slot (A through D on Alteon 708; A through H on Alteon 714)
- `Port`: port number (1-4 for Gigabit-SX; 1-16 for Fast Ethernet, 1-4 for 1000Base-T)
- `Speed`: port speed (10, 100, 10/100, or 10/100/1000)
- `Duplex`: duplex mode (half, full, or auto)
- `Flow Ctrl`: flow control for transmit (Tx) and receive (Rx) (no, yes, or auto)
- `Link`: link status (up or down)

/info/port
Port Information

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

Port information includes:

- `Slot`: indicates LCM slot letter (A through D for Alteon 708; A through H for Alteon 714)
- `Port number` (1-4 for Gigabit-SX; 1-16 for Fast Ethernet, 1-4 for 1000Base-T)
- `Tag`: shows whether tagged frames are discarded (disc), or forwarded (frwd)
- `Untag`: shows whether untagged frames are discarded (disc), or forwarded (frwd)
- `PriTag`: Priority tagged frames, which are frames with a tag on the priority part of a tag, not the VLAN part. This shows whether priority tagged frames are discarded (disc), or forwarded (frwd).
- `RMON` status: `d`=disabled, `e`=enabled
- `VLAN`: indicates VLAN number(s) to which the port belongs

/info/slot
Slot Status Information

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

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

Slot status shows what type of devices are inserted in the LCM slots, and the usage of the SFM slots. Slot status information includes:

- `Slot`: LCM slots, A–D (for Alteon 708) or A–H (for Alteon 714)
- `State`: Shows the status of the LCM
- `LCM type` : Gigabit-SX, Fast Ethernet, 1000Base-T
- `SP States`: shows the states of the switch processors (SPs) that are in use. For information about SP index numbers, see page 67.
- `SFM slots`: shows the SFM slot number(s)
- `Usage`: shows the type of LCM that is using the SFM. SFM types are `GE` (Gigabit-SX or 1000Base-T), or `FE` (Fast Ethernet).
- `LCM slots`: shows the slot letter of the LCM that are using the SFM

Interpreting SP Index Numbers

Switch processors are numbered by SP Index, x0 through x15 on the Alteon 708 (x0 through x31on the Alteon 714).

On the Alteon 708, slot A is assigned the sp index numbers x0 through x3, slot B=x4 through x7, slot C=x8 through x11, and slot D=x12 through x15.

Each LCM slot is assigned four SP index numbers, regardless of the actual number of SP's present on the *installed* LCM:

- A 4-port Gigabit-SX or 1000Base-T LCM requires four SPs; therefore, each SP index number corresponds to a port number. The /info/slot command will show four SP indexes for each LCM slot.
- An 8-port Fast Ethernet LCM requires only two SPs; each SP controls two ports. Therefore, even though there are four possible SP indexes assigned to each LCM slot, the **`/info/slot`** command will show only two SP index numbers for a Fast Ethernet LCM slot.

`/info/vlan`
VLAN Information

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

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

VLAN information includes:

- VLAN Number
- VLAN Name
- Status- enabled (ena), disabled (dis)
- Port membership of the VLAN. Port membership is represented as *slot* (letter) and *port* (number) together, like b12.

/info/stp
Spanning Tree Information

```
Spanning Tree Group 1: On

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

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

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

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

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

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

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

The following table describes the STP parameters.

Table 7 Spanning Tree Parameter Descriptions

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

/info/trunk
Trunk Group Information

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

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

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

/info/arp
Address Resolution Protocol Information

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

The ARP information includes the following:

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

Table 8 ARP Information Menu Options (/info/arp)

Command Syntax and Usage
find *<IP address (e.g., 192.4.17.101)>* Displays a single ARP entry by IP address.
port *<port (must be slot and port together, like b12)>* Displays the ARP entries on a single port.
refpt *<port (must be slot and port together, like b12)>* Displays the ARP entries referenced by a single port.
vlan *<VLAN number 1-4094>* Displays the ARP entries on a single VLAN.
dump Displays all ARP entries, including: ■ IP address and MAC address of each entry ■ Address status flag (see below) ■ The VLAN and port to which the address belongs ■ The ports that have referenced the address (empty if no port has routed traffic to the IP address shown) For more information, see page 72.

Table 8 ARP Information Menu Options (/info/arp)

Command Syntax and Usage	
addr	Displays the ARP address list, which is the list of IP addresses that receive responses to ARP requests. For more information, see page 72.

/info/arp/dump

Show All ARP Entry Information

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

The `Flag` field is interpreted as follows:

Table 9 ARP Flag Parameter

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

/info/arp/addr

ARP Address List Information

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

/info/fdb
FDB Information

```
[Forwarding Database Menu]
      find    - Show a single FDB entry by MAC address
      port    - Show FDB entries on a single port
      vlan    - Show FDB entries on a single VLAN
      dump    - Show all FDB entries
      count   - Show count of FDB entries
```

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

Table 10 FDB Information Menu Options (/info/fdb)

Command Syntax and Usage
find <*MAC-addr*> [<*VLAN*>] Displays a single database entry by its MAC address, and optionally, its VLAN. You are prompted to enter the MAC address of the device. Enter the MAC address using the hex format, `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 as slot and port, like b12*> Displays all FDB entries for a particular port.
vlan <*VLAN number 1-4094*> Displays all FDB entries on a single VLAN.
dump Displays all entries in the Forwarding Database. For more information, see page 74.
count Displays the number of Forwarding Database entries.

/info/fdb/dump

Show All FDB Information

```
   MAC Address      VLAN    Port    State
00:11:22:33:44:45     1      A1      FWD
00:11:22:33:44:7a     1      A1      FWD
00:11:22:33:44:3d     1      A1      FWD
00:11:22:33:44:a4     1      A1      FWD
00:11:22:33:44:24     1      A1      FWD
00:11:22:33:44:4c     1      A1      FWD
00:11:22:33:44:25     1      A1      FWD
00:11:22:33:44:b0     1      A1      FWD
00:11:22:33:44:40     1      A1      FWD
00:11:22:33:44:00     1      A1      FWD
00:11:22:33:44:05     1      A1      FWD
00:11:22:33:44:3f     1      A1      FWD
00:11:22:33:44:ee     1      A1      FWD
00:11:22:33:44:10     1      A1      FWD
00:11:22:33:44:28     1      A1      FWD
00:11:22:33:44:77     1      A1      FWD
00:11:22:33:44:e5     1      A1      FWD
00:11:22:33:44:64     1      A1      FWD
00:11:22:33:44:c1     1      B6      FWD
00:11:22:33:44:15     1      A1      FWD
00:11:22:33:44:5b     1      A1      FWD
00:11:22:33:44:0b     1      A1      FWD
                                      :
                                      :
```

- An address that is in the *forwarding* (FWD) state, means that it has been learned by the switch.
- If the state for the FDB entry 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 FDB entry 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; that is, a virtual router with the same IP address as a virtual server.

Clearing Entries from the Forwarding Database

To delete a MAC address from the FDB or to clear the entire FDB, see page 253.

/info/ip

IP Information

```
IP information:

Interface information:
   1: 172.28.1.200,    255.255.0.0,     172.28.255.255,  vlan 1, up
  10: 192.94.8.3,      255.255.255.0,   192.94.8.255,    vlan 1, up

Default gateway information: metric strict
  1: 172.28.1.1,      up

Current forwarding setting: ON, dirbr disabled

Current port settings: all ports ON
```

IP information includes:

- Router ID, IP address, AS number
- IP interface information: Interface number, IP address, subnet mask, broadcast address, VLAN number, and operational status
- Default gateway information: Metric for selecting which configured gateway to use, gateway number, IP address, and health status
- Current forwarding status
- Current port settings

/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 11 Route Information Menu Options (/info/route)

Command Syntax and Usage
find *<IP address (e.g., 192.4.17.101)>* Displays a single route by destination IP address.
gw *<default gateway address (e.g., 192.4.17.44)>* Displays routes to a single gateway.
type <[indirect|direct|local|broadcast|martian|multicast]> Displays routes of a single type. For more information, see Table 12 on page 77.
tag <**[fixed|static|snmp|addr|rip|icmp|broadcast|martian|multicast|dynamic]**> Displays routes of a single tag. For more information, see Table 13 on page 78.
if *<interface number [1-1024]>* Displays routes on a single interface.
dump Displays all routes known to the switch. For more information, see page 77.

/info/route/dump

Show All IP Route Information

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

Table 12 describes the routing `Type` parameters.

Table 12 IP Routing Type Parameters

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

Table 13 describes the routing `Tag` parameters.

Table 13 IP Routing Tag Parameters

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

/info/slb
SLB Information Menu

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

Layer 4 information includes the following:.

Table 14 Layer 4 Information Menu Options (/info/slb)

Command Syntax and Usage
sess Displays the Session Table Information Menu. To view menu options, see page 80.
real *<real server number [1-4096]>* Real server number, real IP address, MAC address, VLAN, physical switch port, layer where health check is performed, and health check result.
virt *<virtual server number [1-1024]>* ■ Displays Virtual Server State: Virtual server number, virtual IP address, virtual MAC address ■ Virtual Port State: Virtual service or port, server port mapping, real server group, group backup server
filt *<filter ID [1-1024]>* Displays the filter number, destination port, real server port, real server group, health check layer, group backup server, and real server group, IP address, backup server, and status.
port *<port as slot and port, like b12>* Displays the physical port number, proxy IP address, filter status, a list of applied filters, and client and/or server Layer 4 activity.
gslb Displays remote switch number, remote switch IP address, IP subnet mask, and health status.
dump Displays all Layer 4 information for the switch. For details, see page 81.

/info/slb/sess

Show Session Table Information

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

Table 15 Session Information Menu Options (/info/slb/sess)

Command Syntax and Usage
find *<IP address>* Displays all session entries with source IP address.
sp *<SP as port (e.g. b12) or index (e.g. x5)>* Displays all session entries on the switch processor.
dump Displays all session entries.

/info/slb/dump

Show All Layer 4 Information

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

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

Redirect filter state:

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

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

/info/vrrp
VRRP Information

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

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

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

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

/info/swkey
Software Enabled Keys

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

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

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

info/dump
Information Dump

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

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

This data is useful for tuning and debugging switch performance.

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

CHAPTER 6

The Statistics Menu

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

/stats

Statistics Menu

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

Table 16 Statistics Menu Options (/stats)

Command Syntax and Usage
`port` *<port as slot and port, like b12>* Displays the Port Statistics Menu for the specified port. Use this command to display traffic statistics on a port-by-port basis. Traffic statistics include SNMP Management Information Base (MIB) objects. To view menu options, see page 86.
`slb` Displays the SLB (server load balancing) Statistics Menu. To view menu options, see page 101.
`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 116.

Table 16 Statistics Menu Options (/stats)

Command Syntax and Usage
sp Displays the Switch Processor Statistics Menu. To view menu options, see page 140.
dump After you enter **y** to confirm, this command dumps all switch statistics. Use this command to gather data for tuning and debugging switch performance. If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump command. To view menu options, see page 145.

/stats/port <slot-number 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
      link    - Link statistics
      rmon    - RMON statistics
      clear   - Clear statistics
```

NOTE – The `ip` and `maint` commands have been removed from this menu.

Table 17 Port Statistics Menu Options (/stats/port)

Command Syntax and Usage
brg Displays bridging IEEE 802.1 (“dot1”) statistics for the port. For more information, see page 87.
ether Displays Ethernet IEEE 802.3 (“dot3”) statistics for the port. For more information, see page 89.
if Displays interface statistics for the port. For more information, see page 93.

Table 17 Port Statistics Menu Options (/stats/port)

Command Syntax and Usage
link Displays link statistics for the port. For more information, see page 95.
rmon Displays RMON statistics for the port. For more information, see page 96
clear Clears all switch statistics for the port.

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

Bridging ("dot1") statistics

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

Table 18 Bridging "dot1" Statistics* (/stats/port/brg)

Statistic	Description
dot1PortInFrames	The number of frames that have been received by this port from its segment. Note that a frame received on the interface corresponding to this port is only counted by this object if and only if it is for a protocol being processed by the local bridging function, including bridge management frames
dot1PortOutFrames	The number of frames that have been transmitted by this port to its segment. Note that a frame transmitted on the interface corresponding to this port is only counted by this object if and only if it is for a protocol being processed by the local bridging function, including bridge management frames.
dot1PortInDiscards	The count of valid frames received which were discarded (i.e., filtered) by the Forwarding Process.

*** Extract from RFC-1493**

Table 18 Bridging “dot1” Statistics* (/stats/port/brg)

Statistic	Description
dot1TpLearnedEntryDiscards	The total number of Forwarding Database entries, which have been or would have been learnt, but have been discarded due to a lack of storage space in the Forwarding Database. If this counter is increasing, it indicates that the Forwarding Database is regularly becoming full (a condition which has unpleasant performance effects on the subnetwork). If this counter has a significant value but is not presently increasing, it indicates that the problem has been occurring but is not persistent.
dot1BasePortDelayExceededDiscards	The number of frames discarded by this port due to excessive transit delay through the bridge. It is incremented by both transparent and source route bridges.
dot1BasePortMtuExceededDiscards	The number of frames discarded by this port due to an excessive size. It is incremented by both transparent and source route bridges.
dot1StpPortForwardTransitions	The number of times this port has transitioned from the Learning state to the Forwarding state.

*** Extract from RFC-1493**

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

Ethernet (“dot3”) statistics

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

Table 19 Ethernet (“dot3”) statistics* (/stats/port/ether)

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

*** Extract from RFC-1643**

Table 19 Ethernet (“dot3”) statistics* (/stats/port/ether)

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

*** Extract from RFC-1643**

Table 19 Ethernet (“dot3”) statistics* (/stats/port/ether)

Statistic	Description
dot3StatsLateCollisions	The number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet. Five hundred and twelve bit-times corresponds to 51.2 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 purposes of other collision-related statistics.
dot3StatsExcessiveCollisions	A count of frames for which transmission on a particular interface fails due to excessive collisions.
dot3StatsInternalMacTransmitErrors	A count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object. The precise meaning of the count represented by an instance of this object is implementation- specific. In particular, an instance of this object may represent a count of transmission errors on a particular interface that are not otherwise counted.
dot3StatsCarrierSenseErrors	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.
dot3StatsFrameTooLongs	A count of frames received on a particular interface that exceed the maximum permitted frame size. The count represented by an instance of this object is 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 obtain are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.

*** Extract from RFC-1643**

Table 19 Ethernet (“dot3”) statistics* (/stats/port/ether)

Statistic	Description
dot3StatsInternalMacReceiveErrors	A count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsFrameTooLongs object, the dot3StatsAlignmentErrors object, or the dot3StatsFCSErrors object. The precise meaning of the count represented by an instance of this object is implementation- specific. In particular, an instance of this object may represent a count of receive errors on a particular interface that are not otherwise counted.
dot3CollFrequencies [1-15]	A count of individual MAC frames for which the transmission (successful or otherwise) on a particular interface occurs after the frame has experienced exactly the number of collisions in the associated dot3CollCount object. For example, a frame which is transmitted on interface 77 after experiencing exactly 4 collisions would be indicated by incrementing only dot3CollFrequencies.77.4. No other instance of dot3CollFrequencies would be incremented in this example.

*** Extract from RFC-1643**

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

Interface (“if”) statistics

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

Table 20 Interface “if” statistics* (/stats/port/if)

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

*** Extract from RFC-1213**

Table 20 Interface “if” statistics* (/stats/port/if)

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

*** Extract from RFC-1213**

Table 20 Interface “if” statistics* (/stats/port/if)

Statistic	Description
ifOutErrors	For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.

*** Extract from RFC-1213**

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

Link Statistics

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

Table 21 Link Statistics (/stat/sport/link)

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

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

RMON Statistics

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

Table 22 RMON Statistics* (/stats/port/rmon)

Statistic	Description
etherStatsDropEvents	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.

* Extract from RFC-1757

Table 22 RMON Statistics* (/stats/port/rmon)

Statistic	Description
etherStatsOctets	Total number of data octets (including those in bad packets) received on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval. The differences in the sampled values are Pkts and Octets, respectively, and the number of seconds in the interval is Interval. These values are used to calculate the Utilization as follows: $$\text{Utilization} = \frac{\text{Pkts} * (9.6 + 6.4) + (\text{Octets} * .8)}{\text{Interval} * 10{,}000}$$ The result of this equation is the value "Utilization", which is the percent utilization of the ethernet segment on a scale of 0 to 100 percent.
etherStatsPkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
etherStatsBroadcastPkts	The total number of good packets received that were directed to the broadcast address. Note that this does not include multicast packets.
etherStatsMulticastPkts	The total number of good packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
etherStatsCRCAlignErrors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
etherStatsUndersizePkts	The total number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
etherStatsOversizePkts	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.

* Extract from RFC-1757

Table 22 RMON Statistics* (/stats/port/rmon)

Statistic	Description
etherStatsFragments	The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that it is entirely normal for etherStatsFragments to increment. This is because it counts both runts (which are normal occurrences due to collisions) and noise hits.
etherStatsJabbers	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.

* Extract from RFC-1757

Table 22 RMON Statistics* (/stats/port/rmon)

Statistic	Description
etherStatsCollisions	The best estimate of the total number of collisions on this Ethernet segment. The value returned will depend on the location of the RMON probe. Section 8.2.1.3 (10BASE-5) and section 10.3.1.3 (10BASE-2) of IEEE standard 802.3 states that a station must detect a collision, in the receive mode, if three or more stations are transmitting simultaneously. A repeater port must detect a collision when two or more stations are transmitting simultaneously. Thus a probe placed on a repeater port could record more collisions than a probe connected to a station on the same segment would. Probe location plays a much smaller role when considering 10BASE-T. 14.2.1.4 (10BASE-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10BASE-T station can only detect collisions when it is transmitting. Thus probes placed on a station and a repeater, should report the same number of collisions. Note also that an RMON probe inside a repeater should ideally report collisions between the repeater and one or more other hosts (transmit collisions as defined by IEEE 802.3k) plus receiver collisions observed on any coax segments to which the repeater is connected.
etherStatsPkts64Octets	The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including FCS octets).
etherStatsPkts65to127Octets	The total number of packets (including bad packets) received that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
etherStatsPkts128to255Octets	The total number of packets (including bad packets) received that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
etherStatsPkts256to511Octets	The total number of packets (including bad packets) received that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).

* Extract from RFC-1757

Table 22 RMON Statistics* (/stats/port/rmon)

Statistic	Description
etherStatsPkts512to1023Octets	The total number of packets (including bad packets) received that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
etherStatsPkts1024to1518Octets	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).

* Extract from RFC-1757

/stats/slb
Server Load Balancing Statistics

```
[Server Load Balancing Statistics Menu]
      real    - Real server statistics
      group   - Real server group statistics
      virtual - Virtual server statistics
      filter  - Filter statistics
      sp      - SLB switch processor statistics
      gslb    - Global SLB statistics
      maint   - Maintenance statistics
      clear   - Clear all SLB statistics
      dump    - Dump all SLB statistics
```

NOTE – The `port` command, previously in Release 6, has been removed from this menu.

Table 23 SLB Statistics Menu Options (/stats/slb)

Command Syntax and Usage
`real` *<real server number [1-4096]>* Displays the following real server statistics: ■ Number of times the real server has failed its health checks. ■ Number of sessions currently open on the real server. ■ Total sessions on the real server. ■ Highest number of simultaneous sessions recorded for each real server. ■ Real server transmit/receive octets. To view an example of what is displayed on-screen, see page 103.
`group` *<real server group number [1-1024]>* Displays the following real server group statistics: ■ Current and total sessions for each real server in the real server group. ■ Current and total sessions for all real servers associated with the real server group. ■ Highest number of simultaneous sessions recorded for each real server. ■ Real server transmit/receive octets. For more information, see page 103.

Table 23 SLB Statistics Menu Options (/stats/slb)

Command Syntax and Usage
virtual *<virtual server number [1-1024]>* Displays the following virtual server statistics: ■ Current and total sessions for each real server associated with the virtual server. ■ Current and total sessions for all real servers associated with the virtual server. ■ Highest number of simultaneous sessions recorded for each real server. Real server transmit/receive octets. For more information, see page 104.
filter *<filter ID [1-1024]>* Displays the total number of times the specified filter has been used. For more information, see page 104.
sp *<Enter SP as port (e.g. b12) or SP index (e.g. x5)>* Displays the Switch Processor SLB Statistics Menu. Enter the sp as port or index. For more information, see page 105.
gslb Displays the Global SLB Statistics Menu. For more information, see page 110.
maint Displays the SLB maintenance statistics. To see an example of what is displayed on-screen, see page 113.
clear Clears all switch SLB statistics.
dump Dumps all switch SLB statistics. Use this command to gather data for tuning and debugging switch performance. If you want to capture dump data to a file, set your communication software on your workstation to capture session data before you enter the **dump** command.

/stats/slb/real <real server number [1-4096]>

Real Server SLB Statistics

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

/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
---- --------------- -------- ---------- --------  ---------------
   1 123.255.13.12          0        745       15                0
   2 123.255.13.21          0        739       17                0
---- --------------- -------- ---------- --------  ---------------
                            0       1484       32                0
```

This menu displays the statistics of “Current sessions”, “Total sessions” “Highest Sessions”, and “Octets” for all real servers that belong to a particular real server group on a switch port.

/stats/slb/virt <virtual-server-number [1-1024]>

Virtual Server SLB Statistics

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

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

This menu displays the statistics of “Current sessions”, “Total sessions” “Highest Sessions, and “Octets” for all virtual servers that belong to a particular real server group.

/stats/slb/filt <filter-number>

Filter SLB Statistics

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

Total firings displays the number of filter hits on the switch port for a particular filter.

/stats/slb/sp <sp-number>

Switch Processor SLB Statistics

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

Table 24 Switch Processor SLB Statistics Menu Options (/stats/slb/sp)

Command Syntax and Usage
real *<real server number [1-4096]>* Displays real server statistics for the selected Switch Processor (SP). For more information, see page 106.
group *<real server group number [1-1024]>* Displays real server group statistics for the selected SP. For more information, see page 106.
virt *<virtual server number [1-1024]>* Displays virtual server statistics for the selected SP. For more information, see page 106.
filt *<filter ID [1-1024]>* Displays filter statistics for the selected SP. For more information, see page 107.
maint Displays maintenance statistics for the selected SP. For more information, see page 107.
clear Clears all switch processor SLB statistics for this switch processor, resetting them to zero.

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

SP Real Server SLB Statistics

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

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

SP Real Server Group SLB Statistics

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

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

SP Virtual Server SLB Statistics

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

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

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

SP Filter SLB Statistics

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

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

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

SP Maintenance SLB Statistics

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

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

Table 25 SP Server Load Balancing Maintenance Statistics (/stats/slb/sp/maint)

Statistic	Description
Current sessions	Number of session bindings currently in use for the specified SP.
Allocation failures	The number of times a session entry could not be created because the table for the specified SP was full.
L4 disabled	The number of frames addressed to a virtual server for which client processing is disabled for the specified SP.
Overflows	The number of hash chain overflows for the specified SP.
Not ready	The number of times the a session entry for the specified SP was not ready for second or subsequent frames.

Table 25 SP Server Load Balancing Maintenance Statistics (/stats/slb/sp/maint)

Statistic	Description
Lookup failures	The number of frames for the specified SP for which a Layer 3 (MAC address) or Layer 4 (IP address) lookup failed.
TCP fragments	The number of TCP fragments that were dropped for the specified Switch Processor.
UDP datagrams	The number of UDP datagrams that were dropped because the virtual service for the specified SP did not support UDP.
IP options	The number of frames for the specified SP that could not be processed completely because they contain IP options.
Non TCP/IP frames	The number of non-IP based frames for the specified SP that were received by the virtual server.
Incorrect VIPs	The number of times the a frame was dropped for the specified SP because it was addressed to a VMAC but an unknown VIP. The switch has received a Layer 4 request for a virtual server that was not configured.
Incorrect Vports	The number of frames for the specified SP that were dropped to a valid VIP but not a valid VPORT (destination port). The virtual server has received frames for TCP/UDP services that have not been configured. Normally this indicates a mis-configuration on the virtual server or the client, but it may be an indication of a potential security probing application like SATAN.
No Server available	The number of times a frame for the specified SP was dropped because there were no real servers available to accept the session. All real servers are either out of service or at their `mcon` limit.
Filtered (denied) Frames	The number of frames that where dropped because they matched an active filter with the "`deny`" action set.
Client slowpath frames	The number of client frames processed by the control processor(s)
Server slowpath frames	The number of server frames processed by the control processor(s)
ICMP slowpath frames	Number of client ICMP frames processed by the control processor(s)
DAM slowpath frames	The number of DAM frames processed by the control processor(s)

Table 25 SP Server Load Balancing Maintenance Statistics (/stats/slb/sp/maint)

Statistic	Description
DAM sessions created	The number of DAM sessions added to the session table by the control processor(s)
Slowpath sessions created	The number of sessions added to session table by the control processor(s)
SW alloc failures	The number of instances where the switch ran out of available bindings for a port.

/stats/slb/gslb

Global SLB Statistics

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

Table 26 Global SLB Statistics Menu Options (/stats/slb/gslb)

Command Syntax and Usage
real *<real server number [1-4096]>* Where the real server number represents the real server ID on this switch, under which the remote server is configured. For more information, see page 111.
group *<real server group number [1-1024]>* Displays real server group global statistics: For more information, see page 111.
virt *<virtual server number [1-1024]>* Displays virtual server group global statistics. For more information, see page 112.
maint Displays global SLB maintenance statistics. For more information, see page 113.
clear Clears all global SLB statistics for the switch.

/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 to the remote server
- Number of HTTP redirects to the remote server

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

Real Server Group Global SLB Statistics

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

Real server group global statistics include the following:

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

/stats/slb/gslb/virt <virtual-server-number>

Virtual Server Global SLB Statistics

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

Virtual server global statistics include the following:

- Service: type of service running on the virtual server
- Server: type of server configuration and server ID number
  - **v**# represents a local virtual server number
  - **r**# represents a remote site. Since each remote sites is configured on its peers as if it were a real server (with certain special properties), the number represents the real server ID on this switch, under which the remote server is configured.
- IP address of the server
- Response time: the average time (present weighted) that each service takes to respond to information exchanges with its peers. The time is specified in ticks of 65 milliseconds.
- 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 received from remote sites.
- The number of bad DSSP updates received from remote sites. Bad updates usually indicate that there is a GSLB switch configuration problem. If bad updates occur, check your `syslog` for configuration error messages.

/stats/slb/maint

SLB Maintenance Statistics

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

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

Table 27 Server Load Balancing Maintenance Statistics (/stats/slb/maint)

Statistic	Description
Current sessions	The number of session bindings currently in use.
Highest sessions	The highest number of session bindings.

Table 27 Server Load Balancing Maintenance Statistics (/stats/slb/maint)

Statistic	Description
Allocation failures	The number of times a session entry could not be created because the table was full.
L4 disabled	The number of frames addressed to a virtual server for which client processing is disabled.
Overflows	The number of hash chain overflows.
Not ready	The number of times the a session entry was not ready for second or subsequent frames.
Lookup failures	The number of frames for which a Layer 3 (MAC address) or Layer 4 (IP address) lookup failed.
Non TCP/IP frames	The number of non-IP based frames received by the virtual server.
IP options	The number of frames that could not be processed completely because they contain IP options.
TCP fragments	The number of TCP fragments that were dropped.
UDP datagrams	The number of UDP datagrams that were dropped because the virtual service did not support UDP.
Filtered (denied) Frames	The number of frames that where dropped because they matched an active filter with the "`deny`" action set.
Incorrect VIPs	The number of times the a frame was dropped because it was addressed to a VMAC but an unknown VIP. The switch has received a Layer 4 request for a virtual server that was not configured.
Incorrect Vports	The number of frames dropped to a valid VIP but not a valid VPORT (destination port). The virtual server has received frames for TCP/UDP services that have not been configured. Normally this indicates a mis-configuration on the virtual server or the client, but it may be an indication of a potential security probing application like SATAN.
No available real server	The number of times a frame was dropped because there were no real servers available to accept the session. All real servers are either out of service or at their `mcon` limit.
Timed-out sessions	The number of sessions that timed out.

Table 27 Server Load Balancing Maintenance Statistics (/stats/slb/maint)

Statistic	Description
Backup server activated	The number of times a real server failure has occurred and caused a backup server to be brought online.
Overflo server activated	The number of times a real server has reached the `mcon` limit and caused an overflow server to be brought online.
Client slowpath frames	The number of client frames processed by the control processor(s)
Client sessions created	The number of client frames created by the control processor(s)
Server slowpath frames	The number of server frames processed by the control processor(s)
ICMP slowpath frames	Number of client ICMP frames processed by the control processor(s)
DAM slowpath frames	The number of DAM frames processed by the control processor(s)
DAM sessions created	The number of DAM sessions added to the session table by the control processor(s)
Slowpath alloc failures	The number of allocation failures in the control processor(s)

/stats/mp
Management Processor Statistics

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

Table 28 Management Processor Statistics Menu Options (/stats/mp)

Command Syntax and Usage
if *<interface number [1-1024]>* Displays IP interface statistics for the management processors.
ip Displays IP statistics for the management processors.
icmp Display ICMP statistics for the management processors.
tcp Displays TCP statistics for the management processors.
udp Displays UDP statistics for the management processors.
snmp Displays SNMP statistics for the management processors.
arp Displays ARP (Address Resolution Protocol) statistics for the management processors.

Table 28 Management Processor Statistics Menu Options (/stats/mp)

Command Syntax and Usage
route Displays route statistics for the management processors.
dns Displays DNS (Domain Name Server) statistics for the management processors.
vrrp When virtual routers are configured for the management processors, you can display the following protocol statistics for VRRP: ■ Advertisements received (`vrrpInAdvers`) ■ Advertisements transmitted (`vrrpOutAdvers`) ■ Advertisements received, but ignored (`vrrpBadAdvers`)
maint Displays maintenance statistics for the management processor. This sub-menu contains information that may be requested by an Alteon WebSystems support engineer for trouble shooting purposes.
clear Clears all management processor statistics for the switch.

/stats/mp/if <interface-number>

IP Interface Statistics

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

Table 29 IP Interface Statistics* (/stats/mp/if)

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

*** extract from RFC-1213**

Table 29 IP Interface Statistics* (/stats/mp/if)

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

*** extract from RFC-1213**

/stats/mp/ip

IP Statistics

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

Table 30 IP Statistics (/stats/mp/ip)

Statistic	Description
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInHdrErrors	The number of input datagrams discarded 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, etc.
ipInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., 0.0.0.0) and addresses of unsupported Classes (e.g., 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 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, 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.

*** Extract from RFC-1213**

Table 30 IP Statistics (/stats/mp/ip)

Statistic	Description
ipInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., 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 (e.g., 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.
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, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.
ipFragOKs	The number of IP datagrams that have been successfully fragmented at this entity.

*** Extract from RFC-1213**

Table 30 IP Statistics (/stats/mp/ip)

Statistic	Description
ipFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, e.g., 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.
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 field of the IP header of datagrams originated at this entity, 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.

*** Extract from RFC-1213**

/stats/mp/icmp

ICMP Statistics

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

Table 31 ICMP Statistics* (/stats/mp/icmp)

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

*** Extract from RFC-1213**

Table 31 ICMP Statistics* (/stats/mp/icmp)

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

*** Extract from RFC-1213**

/stats/mp/tcp

TCP Statistics

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

Table 32 TCP Statistics (/stats/mp/tcp)

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

*** Extract from RFC-1213**

Table 32 TCP Statistics (/stats/mp/tcp)

Statistic	Description
tcpAttemptFails	The number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
tcpEstabResets	The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
tcpInSegs	The total number of segments received, including those received in error. This count includes segments received on currently established connections.
tcpOutSegs	The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets.
tcpRetransSegs	The total number of segments retransmitted - that is, the number of TCP segments transmitted containing one or more previously transmitted octets.
tcpInErrs	The total number of segments received in error (e.g., 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.

*** Extract from RFC-1213**

/stats/mp/udp

UDP Statistics

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

Table 33 UDP Statistics (/stats/mp/udp)

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

/stats/mp/snmp

SNMP Statistics

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

Table 34 SNMP Statistics (/stats/mp/snmp)

Statistic	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 protocol entity which used a SNMP community name not known to said entity.
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 received SNMP Messages.

Table 34 SNMP Statistics (/stats/mp/snmp)

Statistic	Description
snmpEnableAuthTraps	Indicates whether the SNMP agent process is permitted to generate authentication-failure traps. The value of this object overrides any configuration information; as such, it provides a means whereby all authentication-failure traps may be disabled. Note that it is strongly recommended that this object be stored in non-volatile memory so that it remains constant between re-initializations of the network management system.
snmpOutPkts	The total number of SNMP Messages 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 PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is ‘tooBig’.
snmpInNoSuchNames	The total number of SNMP 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 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 PDUs which were delivered to the SNMP protocol entity and for which the value of the error-status field is ‘readOnly’. It should be noted that it is a protocol error to generate an SNMP PDU which contains the value `readOnly' in the error-status field, as such this object is provided as a means of detecting incorrect implementations of the SNMP.
snmpInGenErrs	The total number of SNMP PDUs 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 the result of receiving valid SNMP Get-Request and Get-Next PDUs.
snmpInTotalSetVars	The total number of MIB objects which have been altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.

Table 34 SNMP Statistics (/stats/mp/snmp)

Statistic	Description
snmpInGetRequests	The total number of SNMP Get-Request PDUs which have been accepted and processed by the SNMP protocol entity.
snmpInGetNexts	The total number of SNMP Get-Next PDUs which have been accepted and processed by the SNMP protocol entity.
snmpInSetRequests	The total number of SNMP Set-Request PDUs which have been accepted and processed by the SNMP protocol entity.
snmpInGetResponses	The total number of SNMP Get-Response PDUs which have been accepted and processed by the SNMP protocol entity.
snmpInTraps	The total number of SNMP Trap PDUs which have been accepted and processed by the SNMP protocol entity.
snmpOutTooBigs	The total number of SNMP PDUs which were generated by the SNMP protocol entity and for which the value of the error-status field is ‘tooBig.’
snmpOutNoSuchNames	The total number of SNMP 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 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 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 PDUs which have been generated by the SNMP protocol entity.
snmpOutGetNexts	The total number of SNMP Get-Next PDUs which have been generated by the SNMP protocol entity.
snmpOutSetRequests	The total number of SNMP Set-Request PDUs which have been generated by the SNMP protocol entity.
snmpOutGetResponses	The total number of SNMP Get-Response PDUs which have been generated by the SNMP protocol entity.

Table 34 SNMP Statistics (/stats/mp/snmp)

Statistic	Description
snmpOutTraps	The total number of SNMP Trap PDUs which have been generated by the SNMP protocol entity.

/stats/mp/route

Route statistics

```
Route statistics:
ripInPkts:                    0   ripOutPkts:                    0
ripInDiscards:                0   ripRoutesAgedOut:             0
ipRoutesCur:                  8   ipRoutesHighWater:            8
ipRoutesMax:              65536
```

Table 35 Route Statistics (/stats/mp/route)

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

/stats/mp/arp

Address Resolution Protocol Statistics

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

Table 36 ARP Statistics (/stats/mp/arp)

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

/stats/mp/dns

DNS Statistics

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

Table 37 DNS Statistics (/stats/mp/dns)

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

/stats/mp/vrrp

VRRP Statistics

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

Virtual Router Redundancy Protocol (VRRP) support on Alteon WebSystems' 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:

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

/stats/mp/maint

MP Maintenance Statistics Menu

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

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

/stats/mp/maint/ltr

Letter Statistics

```
Letter statistics for MP:

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

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

/stats/mp/maint/msg

STEM Message Statistics for MP

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

/stats/mp/maint/mem

STEM Memory Statistics for Management Processor

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

/stats/mp/maint/amem

All Stem Memory Allocated Blocks

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

/stats/mp/maint/thr

STEM Thread Statistics for MP

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

/stats/mp/maint/pkt

Packet Counts

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

/stats/mp/maint/tcb

All TCB Allocated Control Blocks

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

This menu displays statistics on the TCP Control Block (TCB). For example:

```
0019bf94:  0.0.0.0            0 <=> 0.0.0.0           80  listen
 |          |                 |       |                |    |-- state
 |          |                 |       |                |-- source port
 |          |                 |       |-- source IP
 |          |                 |-- destination port
 |          |-- destination IP
 |-- memory location of TCB
```

/stats/mp/maint/uart

UART Statistics

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

/stats/sp
Switch Processor Statistics

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

Table 38 Switch Processor Statistics Menu Options (/stats/sp)

Command Syntax and Usage
ip *<interface number [1-1024]>* Displays IP interface statistics for the specified switch processor. For more information, see page 120.
maint Displays maintenance statistics for the switch processor. This sub-menu contains information that may be requested by an Alteon WebSystems support engineer for trouble shooting purposes.
clear Clears all switch processor statistics.

/stats/sp/ip <interface-number [1-1024]>

IP Statistics

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

Table 39 IP Statistics (/stats/sp/ip)

Statistic	Description
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInHdrErrors	The number of input datagrams discarded 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, etc.
ipInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
ipForwDatagrams	The number of input datagrams for which this entity 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, 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 (e.g., for lack of buffer space). Note that this counter does not include any data-grams discarded while awaiting re-assembly.
ipFragOKs	The number of IP datagrams that have been successfully frag-mented at this entity.

Table 39 IP Statistics (/stats/sp/ip)

Statistic	Description
ipFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
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 but could not be, e.g., because their “Don’t Fragment” flag was set.
ipOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this ‘no-route’ criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.

/stats/sp/maint

SP Maintenance Statistics

```
[SP index 0 (port A1) Maintenance Statistics Menu]
      ltr     - Letter counts
      msg     - STEM message statistics
      mem     - STEM memory statistics
      thr     - STEM thread statistics
      hdlr    - Event handler statistics
```

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

/stats/sp/maint/ltr

Letter Counts for SP

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

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

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

/stats/sp/maint/msg

STEM Message Statistics for SP

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

/stats/sp/maint/mem

STEM Memory Statistics for SP

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

/stats/sp/maint/thr

STEM thread statistics for SP

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

/stats/dump
Statistics Dump

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

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

CHAPTER 7

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 differences are called out in the text.

NOTE – After configuring the switch, make sure to use to **`apply`** command to make the configurations operational.

/cfg Configuration Menu

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

Table 40 Configuration Menu Options (/cfg)

Command Syntax and Usage
sys Displays the System Configuration Menu.To view menu options, see page 151.
udef Displays the User-Defined Defaults Menu. Use this menu to set user-defined default link configurations and user-defined Ethertypes. To view menu options, see page 153.
port *<port as slot and port, like b12>* Displays the Port Configuration Menu. To view menu options, see page 156.
vlan *<VLAN number [1-4094]>* Displays the VLAN Configuration Menu. To view menu options, see page 161.
stp *<group number [1-8]>* Displays the Spanning Tree Configuration Menu. To view menu options, see page 163.
snmp Displays the SNMP Configuration Menu. To view menu options, see page 167.
trunk *<group number [1-8]>* Displays the Trunk Group Configuration Menu. Alteon 700 switches support eight trunk groups and up to four wire ports per trunk group.To view menu options, see page 170.
ip Displays the IP Configuration Menu. To view menu options, see page 173.
slb Displays the Server Load Balancing Configuration Menu. To view menu options, see page 185.
vrrp Displays the Virtual Router Redundancy Protocol Configuration Menu. To view menu options, see page 222.
setup Walks you through the configuration of System Date and Time, BOOTP, IP address and subnet mask, Spanning Tree, Port and Link characteristics, and VLANs. For more information, see page 31.
dump Dumps current configuration to a script file.

Table 40 Configuration Menu Options (/cfg)

Command Syntax and Usage
putcfg *<host name or IP address> <filename on host>* Backs up current configuration to TFTP server.
getcfg *<host name or IP address> <filename on host>* Restores the backed-up configuration from TFTP server. Reads the configuration file into the *new* (pending) configuration block, not the current block. Also, you must apply the changes. See "Viewing, Applying, and Saving Changes" on page 149.

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
- Revert, delete all pending changes

Viewing Pending Changes

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

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

Applying Pending Changes

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

```
# apply
```

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

NOTE – All configuration changes take effect immediately when applied, except for starting Spanning-Tree Protocol. To turn STP on or off, you must apply the changes, save them (see below), and then reset the switch (see "Resetting the Switch" on page 245).

Saving the Configuration

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

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

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

```
# save
```

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

```
# save noback
```

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

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

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

For instructions on selecting the configuration to run at the next system reset, see "Selecting a Configuration Block" on page 245."

/cfg/sys
System Configuration

```
[System Menu]
      date    - Set system date
      time    - Set system time
      idle    - Set timeout for idle CLI sessions
      bootp   - Enable/disable use of BOOTP
      telnet  - Enable/disable Telnet access
      snmp    - Enable/disable SNMP management access
      web     - Enable/disable Web management access
      wport   - Set Web server port number
      mnet    - Set management network
      mmask   - Set management netmask
      banner  - Set login banner
      usrpw   - Set user password
      admpw   - Set administrator password
      l4apw   - Set L4 administrator password
      l4opw   - Set L4 operator password
      current - Display current system-wide configuration
```

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

NOTE – In previous releases of switch software, the `web` command was referred to as `http`.

Table 41 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
`date` Prompts the user for the system date.
`time` Prompts the user for the system time using a 24-hour clock format.
`idle` *<idle timeout in minutes; affects both console and Telnet>* Prompts the user for the idle timeout for command-line interface sessions; the range is 1 to 60 minutes. The default is 5 minutes.

Table 41 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
bootp disable|enable Enables or disables the use of BOOTP; if you enable BOOTP, the switch will query its BOOTP server for switch IP parameters such as addr, default gateway, etc.
snmp [d|e] Enables or disables SNMP-based network management. If disabled, you cannot configure or manage the switch using an SNMP network management station.
web disable|enable Enables or disables HTTP access to the web-based interface.
wport *<TCP port number [1-65535]>* Sets the switch port to be used for serving switch web content. The default is HTTP port 80. If Global Server Load Balancing is to be used, change this parameter to use a different port (such as 8080).
mnet *<IP subnet (e.g., 192.4.17.0)>* Sets the base source IP subnet allowed to access switch management through Telnet, SNMP, RIP, or the WebOS web interface. A range of IP addresses is produced when used with mmask (below). Specify an IP address in dotted-decimal notation.
mmask *<IP subnet mask (e.g., 255.255.0.0)>* This IP address mask is used with mnet to set a range of source IP addresses allowed to access switch management functions. Specify the mask in dotted-decimal notation.
banner *<string, maximum 80 characters>* Configures a login banner of up to 80 characters. When a user or administrator logs into the switch via Telnet, the login banner is displayed; it is also displayed as part of the output from the /info/sys command.
usrpw Configures the user password; the user password can have a maximum of 15 characters.
admpw Configures the administrator password; the administrator password can have a maximum of 15 characters.
l4apw Configures the Layer 4 administrator password; the L4 administrator password can have a maximum of 15 characters. The Layer 4 administrator can view all switch information and statistics, but can configure changes only on the Server Load Balancing menus.

Table 41 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
l4opw Configures the Layer 4 operator password; the Layer 4 operator password can have a maximum of 15 characters.
current Displays the values for most system parameters (excluding passwords).

/cfg/udef
User-Defined Defaults

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

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

Table 42 User-Defined Defaults Configuration Menu Options (/cfg/udef)

Command Syntax and Usage
link *<link type: [FE|GE-SX|1000Base-T]>* Displays the Link Defaults Menu. This menu enables you to create user-defined default link configurations, for Fast Ethernet, 1000Base-SX Gigabit Ethernet (GE-SX), or 1000Base-T Gigabit Ethernet over Copper. To view menu options, see "User-Defined Link Defaults" on page 154.
etype *<Ethertype index [1-10]>* Displays the User-defined Ethertype Menu. This menu enables you to configure up to 10 user-defined Ethertypes per switch, to be used for PVID configuration. To view menu options, see page 155.

/cfg/udef/link

User-Defined Link Defaults

```
[1000sx Default Menu]
      speed   - Set link speed
      mode    - Set full or half duplex mode
      fctl    - Set flow control
      auto    - Control autonegotiation
      current - Display current link configuration
```

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

However, you may need to make an exception for one or more ports. For example, if connecting port b12 to a laptop that requires a 10 Mbps link, you would configure the link speed for port b12 using the menu for "Port Link Configuration" on page 157.

Table 43 User-Defined Link Defaults Menu Options (/cfg/udef/link)

Command Syntax and Usage
speed 10|100|1000|any (not all options are valid on all ports) Sets the link speed; the choices include: ■ "Any," for automatic detection (default) ■ 10 Mbps ■ 100 Mbps
mode full|half|any Sets the operating mode; the choices include: ■ "Any," for autonegotiation (default) ■ Full-duplex ■ Half-duplex
fctl rx|tx|both|none Sets the flow control for the user-defined link; the choices include: ■ Autonegotiation (default) ■ Receive flow control only (will honor PAUSE frames but not send them) ■ Transmit flow control (will send PAUSE frames but not honor them) ■ Both receive and transmit flow control ■ No flow control

Table 43 User-Defined Link Defaults Menu Options (/cfg/udef/link)

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

/cfg/udef/etype

Ethertype Defaults

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

This menu enables you to configure up to 10 user-defined Ethertypes per switch, to be used for PVID configuration.

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

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

Table 44 Ethertype Defaults Menu Options (/cfg/udef/etype)

Command Syntax and Usage
ether *<Ethertype value [0800-ffff, 0 to delete]>* Used for specifying the Ethertype value in the frame. The values 0800-ffff identify the range of Ethertype values.
current Displays the current Ethertype value.
dump Displays all user-defined Ethertypes.

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

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

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

Table 45 Port Configuration Menu Options (/cfg/port)

Command Syntax and Usage
link *<link type* **[none|FE|GE-SX|1000Base-T]** Displays the Link Menu for the specified port. Enter the new link type **[none|FE|GE-SX|1000Base-T]** to display the link menu for that type of link. To view menu options, see page 157.
pvid Displays the Port VLAN ID Menu, which is used to configure default PVIDs for each port. The PVID is the default VLAN number that will be used to forward frames that are not VLAN-tagged. Alteon 700 Series switches allows users to configure PVIDs for each different Ethernet type. There can be up to eight PVIDs per port – IP, IPX, and five user-defined Ethernet types. Also, PVIDs can be configured for each port to classify untagged frames that do not match the other Ethertypes. The factory default values for these will be “1”. To view menu options, see page 159.
tag forward|discard (or just **f|d**). Forwards/discards all VLAN tagged frames.
utag forward|discard (or just **f|d**) Forwards/discards all untagged frames.
ptag forward|discard (or just **f|d**) Forwards/discards all priority tagged frames.

Table 45 Port Configuration Menu Options (/cfg/port)

Command Syntax and Usage
rmon disable|**enable** (or just **d**|**e**) Enables or disables RMON support on the port.
enable Enables the port.
disable Disables the port.
current Displays current port configuration.

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

Port Link Configuration

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

The port link configuration menu is used to configure the link speed, mode, flow control, and autonegotiation for a port.

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

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

Table 46 Port Link Configuration Menu Options (/cfg/port/link)

Command Syntax and Usage
speed 10|100|1000|any (not all options are valid on all ports) Sets the link speed; the choices include: ■ "Any," for automatic detection (default) ■ 10 Mbps ■ 100 Mbps ■ 1000 Mbps
mode full|half|any Sets the operating mode; the choices include: ■ "Any," for autonegotiation (default) ■ Full-duplex ■ Half-duplex
fctl rx|tx|both|none Sets flow control for the port; the choices include: ■ Autonegotiation (default) ■ Receive flow control only ■ Transmit flow control only ■ Both receive and transmit flow control ■ No flow control
auto on|off Enable or disable autonegotiation for the port.
current Displays current port parameters.

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

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

Use this menu to configure default Port [Default] VLAN IDs (PVIDs) for ports. The Alteon 700 Series switches allow you to configure PVIDs for each different Ethernet type. There can be up to eight PVIDs per port: IP, IPX, and five user-defined Ethernet types, as well as the PVID that can be configured for each port to classify untagged frames that do not match the other Ethertypes. The factory default values for the PVIDs is "1."

Table 47 Port VLAN ID Configuration Menu Options (/cfg/port/pvid)

Command Syntax and Usage
ip *<VLAN number [1-4094]>* Sets the PVID for IP/ARP/RARP.
ipx *<VLAN number [1-4094]>* Sets the PVID for IPX.
user *<user PVID index number [1-5]>* This menu is for setting the PVID for one user-specified Ethertype. To view menu options, see page 160.
other *<VLAN number [1-4094]>* Sets the PVID for all other Ethertypes.
current Displays the current PVIDs for the port.

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

User PVID Port Configuration

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

Table 48 User PVID Port Configuration Menu (/cfg/port/pvid/user)

Command Syntax and Usage

etype *<Ethertype index [1-10]>*

Specifies the Ethertype index.

vlan *<VLAN number [1-4094]>*

Specifies the VLAN ID.

undef

Undefines this user PVID.

current

Displays the current Ethertype and VLAN for this user PVID.

/cfg/vlan <vlan-number [1-4094]>
VLAN Configuration

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

The commands in this menu configure VLAN attributes, change the status of the VLAN, delete the VLAN, and change the port membership of the VLAN. For more information on configuring VLANs, see "Setup Part 3: VLANs" on page 38.

Table 49 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.

stg *<Spanning Tree Group index [1-256]>*
Assigns this VLAN to a Spanning Tree Group.

add
Adds port(s) or trunk group(s) to the VLAN membership.

remove
Removes port(s) or trunk group(s) from this VLAN.

clear
Removes all ports and trunk groups from this VLAN.

enable
Enables this VLAN.

disable
Disables this VLAN without removing it from the configuration.

Table 49 VLAN Configuration Menu Options (/cfg/vlan)

Command Syntax and Usage
delete Deletes this VLAN.
current Displays all currently configured VLANs.

NOTE – You cannot add a port to more than one VLAN unless the port has VLAN tagging turned on (see "Port Configuration" on page 156).

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.

/cfg/stp <Spanning Tree Group Index [1-256]>
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
      current - Display current Spanning Tree configuration
```

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

NOTE – When VRRP is used for active/active redundancy, STP must be enabled.

Table 50 Spanning Tree Configuration Menu (/cfg/stp)

Command Syntax and Usage
brg Displays the Bridge Spanning Tree Menu. To view menu options, see page 164.
port *<port as slot and port, like b12>* Displays the Spanning Tree Port Menu. To view menu options, see page 166.
add *<VLAN number [1-4094]>* Adds a VLAN to this Spanning Tree Group.
remove *<VLAN number [1-4094]>* Removes a VLAN from a Spanning Tree Group.
clear Removes all VLANs from a Spanning Tree Group.
on Globally enables STP.
off Globally disables STP.

Table 50 Spanning Tree Configuration Menu (/cfg/stp)

Command Syntax and Usage
default Displays the default Spanning Tree and Member parameters
cur Displays current STP parameters.

/cfg/stp <index[1-256]>/brg
Bridge Spanning Tree Configuration

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

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 51 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.
maxage *<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 any bridge port has to wait before it changes from learning state to forwarding state. The range is 4 to 30 seconds, and the default is 15 seconds.
aging *<new bridge Aging Time [10-65535 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 10 to 65535 seconds, and the default is 300 seconds. To disable aging, set this parameter to 0.
current Displays the current bridge STP parameters.

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

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

/cfg/stp <index[1-256]>/port < port (e.g. b12) or trunk group (e.g. t2):>

Spanning Tree Port Configuration

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

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

- Port priority
- Port path cost

Table 52 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 cost will be computed for an autonegotiated link speed.
on Enables STP on the port.
off Disables STP on the port.
current Displays the current STP port parameters.

/cfg/snmp
SNMP Configuration

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

NOTE – The `trap1`, `trap2`, `t1comm`, and `t2comm` commands in Release 6.0 have been replaced with commands in the trap - SNMP Trap Host Configuration Menu To view menu options, see page 169.

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 53 SNMP Configuration Menu Options (/cfg/snmp))

Command Syntax and Usage
name *<new string, maximum 64 characters>* Configures the name for the system. The name can have a maximum of 64 characters.
locn *<new string, maximum 64 characters>* Configures the system location information (for example, “Bldg. 1”. The system location can have a maximum of 64 characters.
contact *<new string, maximum 64 characters>* Configures the system contact information (for example, a phone number). The system contact can have a maximum of 64 characters.
rcomm *<new SNMP read community string, maximum 32 characters>* Configures the SNMP read community string. The read community string controls SNMP “get” access to the switch. It can have a maximum of 32 characters.
wcomm *<new SNMP write community string, maximum 32 characters>* Configures the SNMP write community string. The write community string controls SNMP “set” and “get” access to the switch. It can have a maximum of 32 characters.
trap *<trap host index:* **1**|**2**>* Displays the SNMP Trap Host Configuration Menu. The SNMP trap host is the device that receives SNMP trap messages from the switch. To view menu options, see page 169.
auth **disable**|**enable** (or just **d**|**e**) 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 for the given port.
current Displays the current STP port parameters.

/cfg/snmp/trap <*trap-host-number*>

Trap Host SNMP Configuration

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

Table 54 Trap Host SNMP Configuration Menu Options (/cfg/snmp/trap)

Command Syntax and Usage
addr <*new SNMP trap host IP address (e.g., 192.4.17.101)*> Configures the IP address of this SNMP trap host using dotted decimal notation. The SNMP trap host is the device that receives SNMP trap messages from the switch.
comm <*new trap host community string, maximum 32 characters*> Configures the community string for this trap host.

/cfg/trunk <group number [1-4]>
Trunk Configuration

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

Trunk groups can provide super-bandwidth connections between Alteon WebSystems 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 switch. The following restrictions apply:

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

Table 55 Trunk Configuration Menu Options (/cfg/trunk)

Command Syntax and Usage
add *<port as slot and port, like b12>* Adds a physical port to the current trunk group.
remove *<port as slot and port, like b12>* Removes a physical port from the current trunk group.
pvid Displays the Default Port VLAN ID configuration menu. To view menu options, see page 171.
tag forward|discard (or just **f|d**) Forwards/discards all VLAN tagged frames.

Table 55 Trunk Configuration Menu Options (/cfg/trunk)

Command Syntax and Usage
utag forward|discard (or just **f|d**) Forwards/discards all untagged frames.
ptag forward|discard (or just **f|d**) Forwards/discards all priority tagged frames
enable Enables the current trunk group.
disable Disables the current trunk group.
current Displays current trunk group parameters.

/cfg/trunk <trunk-group-number> /pvid

Port VLAN ID Trunk Configuration

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

Table 56 PVID Trunk Configuration Menu Options (/cfg/trunk/pvid)

Command Syntax and Usage
ip *<VLAN number [1-4094]>* Sets the PVID for IP/ARP/RARP.
ipx *<VLAN number [1-4094]>* Sets the PVID for IPX.
user *<user PVID index number [1-5]>* Displays the PVID configuration menu for user defined Ethertypes. To view menu options, see page 172.

Table 56 PVID Trunk Configuration Menu Options (/cfg/trunk/pvid)

Command Syntax and Usage
other *<VLAN number [1-4094]>* Sets the PVID for all other Ethertypes.
current Displays current PVID configuration.

/cfg/trunk <trunk-group-number> /pvid/user <user PVID index number [1-5]>

User PVID Port Configuration Menu

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

Table 57 PVID Trunk Configuration Menu Options (/cfg/trunk/pvid)

Command Syntax and Usage
etype *<Ethertype config index [1-10]>* Specifies the Ethertype index for this user PVID. See page 155.
vlan *<VLAN number [1-4094]>* Specifies the VLAN ID for this user PVID.
undef Undefines this user PVID.
current Displays current configuration for this user PVID.

/cfg/ip
IP Configuration

```
[IP Menu]
      if      - Interface menu
      gw      - Default gateway menu
      route   - Static route menu
      frwd    - Forwarding menu
      rip     - Routing Information Protocol (RIP) menu
      port    - Port menu
      dns     - Domain Name System menu
      metric  - Set default gateway metric
      rearp   - Set re-ARP period in minutes
      log     - Set IP address of syslog host
      log2    - Set IP address of second syslog host
      logfac  - Set facility of syslog host
      log2fac - Set facility of second syslog host
      current - Display current IP configuration
```

Table 58 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage

if *<interface number [1-1024]>*

Displays the IP Interface Menu. To view menu options, see page 175.

gw *<default gateway number [1-4]>*

Displays the IP Default Gateway Menu. To view menu options, see page 176.

route

Displays the IP Static Route Menu. To view menu options, see page 177.

frwd

Displays the IP Forwarding Menu. To view menu options, see page 178.

rip

Displays the Routing Interface Protocol Menu. To view menu options, see page 179.

port *<port as slot and port, like b12>*

Displays the IP Port Menu. To view menu options, see page 182.

dns

Displays the IP Domain Name System Menu. To view menu options, see page 183.

Table 58 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage
metric s|r Sets the default gateway metric to strict or roundrobin.
rearp *<2-120 minutes>* Sets the re-ARP period in minutes. To view menu options, see page 184
log *<syslog host IP address (e.g., 192.4.17.223)>* Sets the IP address of syslog host.
log2 *<second syslog host IP address (e.g., 192.4.17.223)>* Sets the IP address of second syslog host.
logfac *<syslog host number [0-7]>* Sets the facility of the syslog host
log2fac *<second syslog host number [0-7]>* Sets the facility of the second syslog host
current Displays the current IP configuration.

/cfg/ip/if <interface-number [1-1024]>

IP Interface Configuration

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

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

Table 59 IP Interface Menu Options (/cfg/ip/if)

Command Syntax and Usage
addr *<IP address (e.g., 192.4.17.101)>* Configures the IP address of the switch interface using dotted decimal notation.
mask *<IP subnet mask (e.g., 255.255.255.0)>* Configures the IP subnet address mask for the interface using dotted decimal notation.
bcast *<broadcast address (e.g., 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 only one VLAN, though any VLAN can have multiple IP interfaces in it.
enable Enables this interface.
disable Disables this interface.
delete Deletes this interface.
current Displays the current interface settings.

/cfg/ip/gw <gateway-number [1-4]>

Default IP Gateway Configuration

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

The switch can be configured with up to four default IP gateways.

Table 60 Default Gateway Options (/cfg/ip/gw)

Command Syntax and Usage
addr *<default gateway address (e.g., 192.4.17.44)>* Configures the IP address of the default IP gateway using dotted decimal notation.
inter *<value [0-60 seconds]>* The switch pings the default gateway to verify that the gateway is up. The **inter** option lets you choose the time between health checks. The range is from 0 to 60 seconds. The default interval is 2 seconds. Setting the interval to 0 disables gateway pinging, and the interface is always assumed to be up.
retry *<attempts [1-120]>* Set the number of failed health check attempts required before declaring this default gateway inoperative. The range is from 1 to 120 attempts. The default is 8 attempts.
arp disable|enable (or just **d|e**) Enables or disables ARP-only health checks.
enable Enables the gateway for use.
disable Disables the gateway.

Table 60 Default Gateway Options (/cfg/ip/gw)

Command Syntax and Usage
delete Delete this gateway from the configuration.
current Displays the current gateway settings.

/cfg/ip/route

IP Route Configuration

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

Table 61 IP Static Route Menu (/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, gateway address, and interface number. Enter all addresses using dotted decimal notation.
remove *<destination> <mask>* Removes a static route. Specify the destination address and mask of the route to remove, using dotted decimal notation.
current Displays the current IP static routes.

/cfg/ip/frwd

IP Forwarding Configuration

```
[IP Forwarding Menu]
      dirbr   - Enable/disable forwarding directed broadcasts
      on      - Globally turn Forwarding ON
      off     - Globally turn Forwarding OFF
      current - Display current IP Forwarding configuration
```

NOTE – The `lnet` and `lmask` commands in Release 6.0 have been removed from this menu.

The IP Forwarding Menu is used for setting the local network address and netmask for the route cache, and to turn IP forwarding (routing) on or off.

Table 62 IP Forwarding Options (/cfg/ip/frwd)

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

/cfg/ip/rip

Routing Information Protocol Configuration

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

Table 63 Routing Information Protocol Menu (/cfg/ip/rip)

Command Syntax and Usage

if *<interface number [1-1024]>*

Displays the RIP Interface Menu. To view menu options, see page 180.

update *<update period [1-120 seconds]>*

Sets the RIP update period in seconds. The default is 30 seconds.

on

Globally turns RIP ON.

off

Globally turns RIP OFF.

current

Displays the current RIP configuration.

/cfg/ip/rip/if <interface-number[1-1024]>

RIP Interface Configuration

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

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

NOTE – Do not configure RIP2 parameters if your routing equipment uses RIP version 1.

Table 64 RIP Interface Menu Options (/cfg/ip/rip/if)

Command Syntax and Usage
version 1|2 Sets the version of Routing Information Protocol.
supply disable|enable (or just **d|e**) When enabled, the switch supplies routes to other routers.
listen disable|enable (or just **d|e**) When enabled, the switch learns routes from other routers.
default none|listen|supply|both ■ When **listen** is enabled, the switch accepts default routes from other routers and gives them priority over configured default gateways. ■ When **supply** is enabled, the switch advertises RIP default routes to other routers. ■ When **both** is enabled, the switch accepts and advertises RIP default routes. ■ When **none** is enabled, the switch rejects RIP default routes.
poison disable|enable (or just **d|e**) When enabled, the switch uses split horizon with poisoned reverse. When disabled, the switch uses only split horizon.

Table 64 RIP Interface Menu Options (/cfg/ip/rip/if)

Command Syntax and Usage
trigg disable|**enable** (or just **d**|**e**) Enables or disables triggered updates.
mcast disable|**enable** (or just **d**|**e**) Enables or disables multicast updates. This parameter can be configured only in RIP version 2.
metric <*value [1-15]*> Sets the metric.
auth none|**password** Sets the authentication type. This parameter can be configured only in RIP version 2.
key <*key*> Sets the authentication key. This parameter can be configured only in RIP version 2.
current Displays the current RIP settings.

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

IP Port Configuration

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

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

Table 65 IP Forwarding Port Options (/cfg/ip/port)

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

/cfg/ip/dns

Domain Name System Configuration

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

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

Table 66 Domain Name Service Menu Options (/cfg/ip/dns)

Command Syntax and Usage
`prima` *<IP address (e.g., 192.4.17.101)>* Used for setting the IP address for your primary DNS server. Use dotted decimal notation.
`secon` *<IP address (e.g., 192.4.17.101)>* Used for setting the IP address for your secondary DNS server. If the primary DNS server fails, the configured secondary will be used instead. Use dotted decimal notation.
`dname` *<dotted DNS notation>*|**`none`** Used for setting the default domain name used by the switch. For example: `mycompany.com`
`current` Displays the current Domain Name System settings.

/cfg/ip/metric <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; each is described in the table below:

Table 67 Default Gateway Metrics (/cfg/ip/metric)

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

/cfg/ip/rearp <re-arp-interval [2-120]>

Re-ARP Interval Configuration

The switch periodically sends ARP (Address Resolution Protocol) requests to refresh its address cache. This command is used for setting the interval between ARP refresh requests.

re-arp-interval is the number of minutes (from 2 to 120) between refreshes of the next IP address in the cache.

/cfg/slb
SLB Configuration

```
[Layer 4 Menu]
      real    - Real server menu
      group   - Real server group menu
      virt    - Virtual server menu
      filt    - Filtering menu
      port    - Layer 4 port menu
      gslb    - Global SLB menu
      sync    - Config synch menu
      adv     - Layer 4 advanced menu
      on      - Globally turn Layer 4 processing ON
      off     - Globally turn Layer 4 processing OFF
      current - Display current Layer 4 configuration
```

Table 68 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage

real *<real server number [1-4096]>*

Displays the menu for configuring real servers. To view menu options, see page 187.

group *<real server group number [1-1024]>*

Displays the menu for placing real servers into real server groups. To view menu options, see page 191.

virt *<virtual server number [1-1024]>*

Displays the menu for defining virtual servers. To view menu options, see page 196.

filt *<filter ID [1-1024]>*

Displays the menu for Filtering and Application Redirection. To view menu options, see page 203.

port *<port as slot and port, like b12>*

Displays the menu for setting physical switch port states for Layer 4 activity. To view menu options, see page 210.

gslb

Displays the menu for configuring Global Server Load Balancing. To view menu options, see page 212.

Table 68 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage
synch Displays the configuration synchronization menu. To view menu options, see page 217.
adv Displays the Layer 4 Advanced Menu. To view menu options, see page 219.
on Globally turns on Layer 4 software services for Server Load Balancing and Application Redirection. This option can be performed only once the optional Layer 4 software is enabled (see "Activating Optional Software on page 238). Enabling Layer 4 services is not necessary for using filters only to allow, deny, or NAT traffic (see "Filtering and Layer 4" on page 186).
off Globally disables Layer 4 services. All configuration information will remain in place (if `applied` or `saved`), but the software processes will no longer be active in the switch.
current Displays the current Server Load Balancing configuration.

Filtering and Layer 4

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

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

/cfg/slb/real <server-number [1-4096]>

Real Server SLB Configuration

NOTE – The real-server-number (1 to 4096) represents a real server that you wish to configure.

```
[Real server 1 Menu]
      rip     - Set IP addr of real server
      name    - Set server name
      weight  - Set server weight
      maxcon  - Set maximum number of connections
      backup  - Set backup real server
      inter   - Set interval between health checks
      retry   - Set number of failed attempts to declare server DOWN
      restr   - Set number of successful attempts to declare server UP
      remote  - Enable/disable remote site operation
      proxy   - Enable/disable client proxy operation
      enable  - Enable real server
      disable - Disable real server
      delete  - Delete real server
      current - Display current real server configuration
```

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

- Real server IP address
- Enabling the real server

Table 69 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
rip *<server IP address>* Sets the IP address of the real server in dotted decimal format. When this command is used, the address entered is PINGed to determine if the server is up, and the administrator will be warned if the server does not respond.
name *<string, maximum 15 characters>* Defines a 15-character alias for each Real Server. This will enable the network administrator to quickly identify the server by a natural language keyword value.

Table 69 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
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 194).
maxcon *<maximum connections [0-65535]>* Sets the maximum number of connections that this server should simultaneously support. This option sets a threshold as an artificial barrier, such that new connections will not be issued to this server if the `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 dropped by the virtual server.
backup *<real server number [1-4096]>*|**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.

Table 69 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
inter *<number of seconds between health checks [0-60]>* Sets the interval between real server health verification attempts. Determining the health of each real server is a necessary function for Layer 4 switching. For TCP services, the switch verifies that real servers and their corresponding services are operational by opening a TCP connection to each service, using the defined service ports configured as part of each virtual service. For UDP services, the switch pings servers to determine their status. The `intr` option lets you choose the time between health checks. The range is from 0 to 60 seconds. The default interval is 2 seconds. Setting the interval to 0 disables health checking. The real server is assumed to be always available.
retry *<number of consecutive health checks [1-63]>* Sets the number of failed health check attempts required before declaring this real server inoperative. The range is from 1 to 63 attempts. The default is 4 attempts.
restr *<number of consecutive health checks [1-63]>* Sets the number of successful health check attempts required before declaring a UDP service operational. The range is from 1 to 63 attempts. The default is 8 attempts.
remote disable|**enable** (or just **d**|**e**) Enables or disables remote site operation for this server. This should be enabled when the real IP address supplied above represents a remote server (real or virtual) this switch will access as part of its Global Server Load Balancing network.
proxy disable|**enable** (or just **d**|**e**) Enables or disables proxy IP address translation. With this option enabled (default), a client request from any application can be proxied using a load-balancing Proxy IP address (PIP).
enable You *must* perform this command to enable this real server for Layer 4 service. When enabled, the real server can process virtual server requests associated with its real server group. This option, when the `apply` and `save` commands are used, enables this real server for operation until explicitly disabled. See `/oper/slb/enable` on page 236 for an operations-level command.

Table 69 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage
disable Disables this real server from Layer 4 service. Any disabled server will no longer process virtual server requests as part of the real server group to which it is assigned. This option, when the `apply` and `save` commands are used, disables this real server until it is explicitly re-enabled. This option *does not* perform a graceful server shutdown. See `/oper/slb/disable` on page 236 for an operations-level command.
delete Deletes this real server from the Layer 4 switching software configuration. This removes the real server from operation within its real server groups. Use this command with caution, as it will delete any configuration options that have been set for this real server. This option *does not* perform a graceful server shutdown.
current Displays the current configuration information for this real server.

/cfg/slb/group <group-number [1-1024]>

Real Server Group SLB Configuration

NOTE – The *real-server-group-number* (1 to 1024) represents the number of the real server group that you wish to configure.

```
[Real server group 1 Menu]
      add     - Add real server
      remove  - Remove real server
      metric  - Set metric used to select next server in group
      content - Set health check content
      health  - Set health check type
      backup  - Set backup real server or group
      name    - Set real server group name
      realthr - Set real server failure threshold
      delete  - Delete real server group
      current - Display current group configuration
```

This menu is used for combining real servers into real server groups. Each real server 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 70 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
add *<real server number [1-4096]>* Adds a real server (1-4096) o this real server group.
rem *<real server number [1-4096]>* 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.
metric leastconns|roundrobin|minmisses|hash Set the load balancing metric used for determining which real server in the group will be the target of the next client request. See "Server Load Balancing Metrics" on page 194.
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 70 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage

`health icmp|tcp|http|dns|pop3|smtp|nntp|ftp|imap|radius`

Sets the type of health checking performed. The options are as follows:

- `icmp` For Layer 3 health checking, `ping` the server.
- `tcp` For TCP service, open and close a TCP/IP connection to the server.
- `http` For HTTP service, uses HTTP 1.1 `GETS` when a `HOST:` header is required to check that the URL content specified in `content` is accessible on the server. Otherwise, an `HTTP/1.0 GET` occurs.
- `dns` For Domain Name Service, check that the domain name specified in `content` can be resolved by the server.
- `pop3` For user mail service, check that the *user*:*password* account specified in `content` exists on the server.
- `smtp` For mail-server to mail-server services, check that the user specified in `content` is accessible on the server.
- `nntp` For newsgroup services, check that the newsgroup name specified in `content` is accessible on the server.
- `ftp` For FTP services, check that the filename specified in `content` is accessible on the server through anonymous login.
- `imap` For user mail service, check that the *user:password* value specified in `content` exists on the serve
- `radius` For remote access (RADIUS) server authentication, check that the *user:password* value specified in `content` exists on the switch and the server. To perform application health checking to a RADIUS server, the network administrator must also configure the `/cfg/slb/adv/secret` parameter. The `secret` value is a field of 16 alphanumeric characters that is used by the switch to encrypt a password during the RSA Message Digest Algorithm (MD5) and by the RADIUS server to decrypt the password during verification.

`backup` < **`r`**<*real server number*> **`|none`**

Sets the real server used as the backup/overflow server for this real server.

To prevent loss of service if the entire real server fails, use this option to assign a backup real server number. Then, if the real server group becomes inoperative, the switch will activate the backup real server until one of the original real servers becomes operative again.

The backup server is also used in overflow situations. If all the servers in the real server group reach their `maxcon` (maximum connections) limit, the backup comes online to provide additional processing power until one of the original servers becomes desaturated.

The same backup/overflow server may be assigned to more than one real server at the same time.

Table 70 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
name *<string, maximum 15 characters>* Defines a 15-character alias for each Real Server Group. This will enable the network administrator to quickly identify the server group by a natural language keyword value.
realthr *<real-server failure threshold [0-255], 0 for disabled>* Sets real server failure threshold. Once the configured number of servers have failed, a syslog message is generated.
delete Deletes this real server group from the Layer 4 software configuration. This removes the group from operation under all virtual servers it is assigned to. Use this command with caution: if you remove the only group assigned to a virtual server, the virtual server will become inoperative.
current Displays the current configuration parameters for this real server group.

Server Load Balancing Metrics

Using the *metric* command, you can set a number of metrics for selecting which real server in a group gets the next client request. These metrics are described in the following table:

Table 71 Real Server Group Metrics

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

Table 71 Real Server Group Metrics

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

NOTE – Under the `leastconns` and `roundrobin` metrics, when real servers are configured with weights (see the `weight` option on page 187), a higher proportion of connections is given to servers with higher weights. This can improve load balancing among servers of different performance levels. Weights are not applied when using `hash` or `minmisses`.

/cfg/slb/virt <server-number>

Virtual Server SLB Configuration

NOTE – The *virtual-server-number* (1 to 1024) represents the number of the virtual server that you wish to configure.

```
[Virtual server 1 Menu]
      vip     - Set IP addr of virtual server
      dname   - Set domain name of virtual server
      layer3  - Enable/disable layer 3 only balancing
      service - Virtual service menu
      enable  - Enable virtual server
      disable - Disable virtual server
      delete  - Delete virtual server
      current - Display current virtual configuration
```

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

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

Table 72 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
vip *<server IP address>* Sets the IP address of the virtual server using dotted decimal notation. The virtual server created within the switch will respond to ARPs and PINGs from network ports as if it was a normal server. Client requests directed to the virtual server's IP address will be balanced among the real servers available to it through real server group assignments.
dname *<domain name>*|**none** Sets the domain name for this virtual server. The domain name typically includes the name of the company or organization, and the Internet group code (.com, .edu, .gov, .org, etcetera). An example would be foocorp.com. It does not include the hostname portion (www, www2, ftp, and so on). To define the hostname, see hname below. To clear the dname, specify the name as **none**.

Table 72 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
`layer3 disable|enable` (or just **`d|e`**) Normally, the client IP address is used with the client Layer 4 port number to produce a session identifier. When the `layr3` option is used, the switch uses only the client IP address as the session identifier, associating all the connections from the same client with the same real server while any connection exists between them. This is necessary for some server applications where state information about the client system is divided across different simultaneous connections, and also in applications where TCP fragments are generated. If the real server that the client is assigned to becomes unavailable, the Layer 4 software will allow the client to connect to a different server.
`service` *<virtual port or name, [2 - 65535]>* Displays the Virtual Services Menu. The virtual port name can be a well-known port name, such as http, ftp, and so on. To view a list of well-known ports, see page 205. To view menu options, see page 198.
`enable` Enables this virtual server and its services. This option activates the virtual server within the switch so that it can service client requests sent to its defined IP address.
`disable` This option disables the virtual server so that it no longer services client requests.
`delete` This command removes this virtual server from operation within the switch and deletes it from the Layer 4 switching software configuration. Use this command with caution, as it will delete the options that have been set for this virtual server.
`current` Displays the current parameters for this virtual server.

/cfg/slb/virt <server-number>/service <virtual port or name>

Virtual Server Service Configuration

```
[Virtual Server 1 2 Service Menu]
      group   - Set real server group number
      rport   - Set real port
      hname   - Set hostname
      udp     - Enable/disable UDP balancing
      frag    - Enable/disable remapping UDP server fragments
      delete  - Delete virtual service
      cur     - Display current virtual service configuration
```

This menu is used for configuring services assigned to a virtual server.

Table 73 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage
group *<real server group number [1-1024]>* Sets a real server group for this service. You will be prompted to enter the number (1 to 1024) of the real server group to add to this service.
rport *<real port [2-65535]>* Defines the real server TCP or UDP port assigned to this service. By default, this is the same as the virtual port (service virtual port). If `rport` is configured to be different than the virtual port defined in **/cfg/slb/virt/service** *<virtual port>*, the switch will map the virtual port to this real port.
hname *<hostname>*|**none** Sets the hostname for a service added. This is used in conjunction with dname (page 196) to create a full host/domain name for individual services. The format for this command is as follows: # **hname** *<hostname>* For example, to add a hostname for Web services, you could specify “www” as the host-name. If a `dname` of “foocorp.com” was defined (page 196), “www.foocorp.com” would be the full host/domain name for the service. To clear the `hname`, specify the name as **none**.

Table 73 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage
udp disable|**enable**|**stateless** (or just **d**|**e**|**s**) Enables, disables, or makes stateless UDP balancing for a virtual port. You can configure this option if the service(s) to be load balanced include UDP and TCP. (For example, DNS uses UDP and TCP.) In those environments, you must activate UDP balancing for the particular virtual servers that clients will communicate with using UDP.
frag disable|enable (or just **d|e**) Enables or disables substitution of IP addresses in server response fragments with virtual addresses, when a virtual server is load-balancing a real server.
delete This command removes this virtual service from operation within the switch and deletes it from the Layer 4 switching software configuration. Use this command with caution, as it will delete the options that have been set for this virtual service.
current Displays the current configuration of services on the specified virtual server.

Direct Client Access to Real Servers

Some clients may need direct access to the real servers, to, for example, monitor a real server from a management workstation. This access can be provided in a number of ways, listed below and described in this section:

- Direct Access Mode
- Multiple IP addresses on the server
- Proxy IP addresses
- Port mapping
- Management network

Direct Access Mode

When Direct Access Mode (`/cfg/slb/adv/direct`) is enabled on a switch, any client can communicate with any real server to its load-balanced service. Also, in Direct Access Mode, any number of virtual services can be configured to load balance a real service.

Traffic sent directly to real server IP addresses is excluded from load balancing decisions. The same clients may also communicate to the virtual server IP address and have their requests load balanced.

Multiple IP Addresses on the Server

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

Proxy IP Addresses

Proxy IP addresses are used primarily to eliminate Server Load Balancing topology restrictions in complex networks. Proxy IP addresses can also provide direct access to real servers.

If the switch port to the client is configured with a proxy IP address, the client can access each real server directly using the real server's IP address. This requires that the switch port connected to the real server has server and client processing disabled (see the `server` and `client` options under `/cfg/slb/port` on page 210).

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

Port Mapping

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

NOTE – When Direct Access Mode is enabled on a server, Layer 4 port mapping and default gateway load balancing are not supported.

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

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

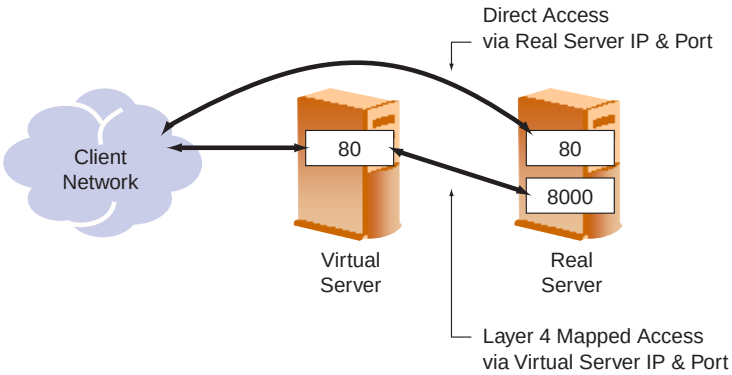

Figure 4 Mapped and Non-Mapped server access

Management Network

Typically, the management network is used by network administrators to monitor real servers and services. By configuring the `mnet` and `mmask` options of the SLB Configuration Menu (`cfg/slb`) you can access the real services being load balanced.

NOTE – Clients on the management network do not have access to Layer 4 services and cannot access the virtual services being load balanced.

The `mnet` and `mmask` options are described below:

- `mnet`: If defined, management traffic with this source IP address will be allowed direct (non-Layer 4) access to the real servers. Specify an IP address in dotted decimal notation. A range of IP addresses is produced when used with the `mmask` option
- `mmask`: This IP address mask is used with the `mnet` to select management traffic which is allowed direct real server access.

Mapping Virtual Ports to Real Ports

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

Use the **`rport`** command to map a virtual server port to a real server port. For example, to map virtual server for http to real server port 8080, enter the following from within the http service menu for Virtual server 1:

```
>> Virtual Server 1 http Service# rport 8080
```

Or, you can accomplish the same thing by entering the following direct command:

```
>> Main# /cfg/slb/virt 1/service http/rport 8080
```

NOTE – This option will not work if Direct Access Mode is enabled.

/cfg/slb/filt <filter-number [1-1024]>

SLB Filter Configuration

```
[Filter 1 Menu]
      smac    - Set source MAC address
      dmac    - Set destination MAC address
      sip     - Set source IP address
      smask   - Set source IP mask
      dip     - Set destination IP address
      dmask   - Set destination IP mask
      proto   - Set IP protocol
      sport   - Set source TCP/UDP port or range
      dport   - Set destination TCP/UDP port or range
      ack     - Enable/disable TCP ack matching
      invert  - Enable/disable filter inversion
      action  - Set action
      group   - Set real server group for redirection
      rport   - Set real server port for redirection
      nat     - Set which addresses are network address translated
      name    - Set filter name
      adv     - Filter advanced menu
      enable  - Enable filter
      disable - Disable filter
      delete  - Delete filter
      current - Display current filter configuration
```

The switch supports up to 1024 traffic filters. Each filter can be configured to allow, deny, redirect or NAT traffic according to a variety of address and protocol specifications, and each physical switch port can be configured to use any combination of filters.

The required minimum of parameters to configure is as follows:

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

Table 74 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
smac Sets the source MAC address.
dmac Sets the destination MAC address
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.
smask This IP address mask is used with the `sip` to select traffic which this filter will affect. See details below for more information on producing address ranges.
dip 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.
dmask *<IP subnet mask (e.g., 255.255.255.0)>* This IP address mask is used with the `dip` to select traffic which this filter will affect. See details below for more information on producing address ranges.
proto any|*<number>*|*<name>* If defined, traffic from the specified protocol is affected by this filter. The protocol number, name, or "**any**" can be specified: ***Number*** ***Name*** 1 icmp 2 igmp 6 tcp 17 udp 112 vrrp

Table 74 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage

sport <*any*|<*name*>|<*port*>|<*port*>–<*port*>>

Sets the source TCP/UDP port or range.

The IP protocol **proto** must be configured before configuring **sport**.

If defined, traffic with the specified TCP or UDP source port will be affected by this filter. The port number, range, name, or "**any**" can be specified. The well-known ports are as follows:

Number	*Name*	*Number*	*Name*
7	echo	179	bgp
9	discard	194	irc
11	sysdat	220	imap3
13	daytime	389	ldap
15	netstat	443	imap
19	chargen	520	rip
20	ftp-data	554	rtsp
21	ftp	1812	radius
22	ssh	1985	hsrp
23	telnet		
25	smtp		
37	time		
42	name		
43	whois		
53	domain		
67	bootps		
68	bootpc		
69	tftp		
70	gopher		
79	finger		
80	http		
109	pop2		
110	pop3		
111	sunrpc		
119	nntp		
123	ntp		
143	imap		
144	news		
161	snmp		
162	snmptrap		

Table 74 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
dport <*any*|<*name*>|<*port*>|<*port*>–<*port*>>
Sets the destination TCP/UDP port or range. The IP protocol **proto** must be configured before configuring **sport**.
If defined, traffic with the specified real server TCP or UDP destination port will be affected by this filter. The port number, range, name, or "**any**" can be specified, just as with **sport** above.
ack disable|**enable** (or just **d**|**e**)
Enables or disables TCP ack matching. Filters with this option enabled match only those frames that have the TCP ACK or RST flag set. This prevents servers from beginning a TCP connection (with a TCP SYN) from a source, such as TCP port 25. The server will drop any frames that have the ACK flag "spoofed" in them and will not allocate space for a new connection.
invert disable|**enable** (or just **d**|**e**)
Inverts the filter logic. If the conditions of the filter are met, *no action is taken.* If the conditions for the filter are *not met,* the *assigned action is performed*.
action
Specifies the action this filter takes: `allow` Allows frames to pass. `deny` Discards frames that fit this filter's profile. This can be used for building basic security profiles. `redir` Redirects frames that fit this filter's profile, such as for web-cache redirection. In addition, Layer 4 processing must be activated (see the `/cfg/slb/on` command on page 185). `nat` Performs generic Network Address Translation (NAT). This can be used to map the source or destination IP address and port information of a private network scheme to/from the advertised network IP address and ports. This is used in conjunction with the `nat` option on page 207, and can also be combined with proxies.
group <*real server group number [1-1024]*>
This option applies only when `redir` is specified at the filter action. Define a real server group (1 to 1024) to which redirected traffic will be sent.

Table 74 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
rport *<real server port [0-65535]>* This option applies only when `redir` is specified at the filter action. `rport` defines the real server TCP or UDP port to which redirected traffic will be sent. For valid Layer 4 health checks, `rport` must be configured whenever TCP protocol traffic is redirected. Also, if transparent proxies are used for Network Address Translation (NAT) on the switch (see the `pip` option on page 210), `rport` must be configured for all Application Redirection filters.
nat source|dest When `nat` is set as the filter action (see "action" on page 206), this command specifies whether the source or the destination information is re-mapped. If you specify **`source`**, the frame's source IP address (`sip`) and port number (`sport`) are replaced with the destination ip address (`dip`) and destination port address (`dport`) values. If you specify **`dest`**, the frame's destination ip address (`dip`) and destination port number (`dport`) are replaced with the `sip` and `sport` values.
name *<string, maximum 15 characters>* Sets the filter name.
adv Displays the Filter Advanced Menu. To view menu options, see page 209.
enable Enables this filter.
disable Disables this filter.
delete Deletes this filter.
current Displays the current configuration for this 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 75 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[1-1024]>/adv

Advanced Filter Configuration

```
[Filter 1 Advanced   Menu]
      proxy    - Enable/disable client proxy
      cache    - Enable/disable caching sessions that match filter
      log      - Enable/disable logging
```

Table 76 Advanced Filter Menu (/cfg/slb/filt/adv)

Command Syntax and Usage

proxy disable|enable (or just **d**|**e**)

Enables or disables client proxy. This option applies only when `redir` or `nat` is specified as the filter action. Enable or disable proxy IP address translation for traffic matching the filter criteria. By default, this option is enabled. If disabled, any proxy defined for the switch port using the `pip` command (see page 210) is not performed for traffic that meets the filter criteria. This is useful when some types of traffic must retain original IP address information, or when other forms of translation (such as Application Redirection or NAT) are preferred.

cache proxy disable|enable (or just **d**|**e**)

Enables or disables caching sessions that match filter.

log disable|enable (or just **d**|**e**)

Enables or disables logging.

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

Port SLB Configuration

```
[SLB Port A1 Menu]
      client  - Enable/disable client processing for port
      server  - Enable/disable server processing for port
      pip     - Set Proxy IP address for port
      submac  - Enable/disable source MAC address substitution
      filter  - Enable/disable filtering for port
      add     - Add filter to port
      remove  - Remove filter from port
      current - Display current port configuration
```

The Web OS Switch software allows you to enable or disable processing independently for each type of Layer 4 traffic (client and server), expanding your topology options.

Table 77 Port Configuration Menu Options (/cfg/slb/port)

Command Syntax and Usage
client disable|enable (or just **d|e**) For Server Load Balancing, the port can be enabled/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 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.
server disable|enable (or just **d|e**) Ports configured to provide real server responses to client requests require real servers to be connected to the Layer 4 switch, directly or through a hub, router, or another switch. When server processing is enabled, the switch port 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.
pip *<server IP address>* Set the proxy IP address for this port using dotted decimal notation. When defined, client address information in Layer 4 requests is replaced with this proxy address. In Server Load Balancing applications, `pip` forces response traffic to return through the switch as required, rather than around it as 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 198).

Table 77 Port Configuration Menu Options (/cfg/slb/port)

Command Syntax and Usage
submac disable|**enable** (or just **d**|**e**) Enables/disables substituting the source MAC address of Server Load Balancing (SLB) and Web Cache Redirection (WCR) client frames with one of the switch's IP interface MAC addresses.
server disable|**enable** (or just **d**|**e**) Ports configured to provide real server responses to client requests require real servers to be connected to the Layer 4 switch, directly or through a hub, router, or another switch. When server processing is enabled, the switch port 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.
filter disable|**enable** (or just **d**|**e**) Enables or disables filtering on this port.
add *<filter ID [1-1024]>* Adds a filter for use on this port.
remove *<filter ID [1-1024]>* Removes a filter from use on this port.
current Displays the current configuration of this port.

NOTE – When changing the filters on a given port, it may take some time before the port session information is updated so that the filter changes take effect. To make port filter changes take effect immediately, clear the session binding table for the port (see the `clear` command in Table 93 on page 236).

/cfg/slb/gslb

Global SLB Configuration

```
[Global SLB Menu]
      site    - Remote Site menu
      dns     - Enable/disable DNS handoffs
      ttl     - Set Time To Live of DNS resource records
      local   - Enable/disable DNS responses with only local addresses
      one     - Enable/disable DNS responses with only one address
      alway   - Enable/disable DNS responses at least one address
      geo     - Enable/disable geographic awareness
      http    - Enable/disable HTTP redirects
      usern   - Enable/disable HTTP redirect to real server name
      mincon  - Set minimum number of site connections
      inter   - Set interval between remote site updates
      weight  - Set local weight
      on      - Globally turn Global SLB ON
      off     - Globally turn Global SLB OFF
      current - Display current Global SLB configuration
```

Table 78 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage

site *<remote site [1-128]>*

Displays the Remote Site Menu for one of up to 128 remote sites. To view menu options, see page 215.

dns **disable**|**enable** (or just **d**|**e**)

Enables or disables DNS hand-offs to peer sites by this switch. This should be enabled for proper GSLB operation. If disabled, whenever the switch receives a DNS request for a configured service, it will respond only with its own virtual IP address, regardless of performance or load considerations.

ttl *<time to live in seconds [0-65535]>*

Specifies the duration (from 0 to 65535 seconds) that the DNS response from the switch (indicating site of best service) will remain in the cache of DNS servers. A lower value may increase the ability of the GSLB system to adjust to sudden changes in traffic load, but will generate more DNS traffic. Higher numbers may reduce the amount of DNS traffic, but may slow GSLB's response to sudden traffic changes.

Table 78 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage
`local disable`|**`enable`** (or just **`d`**|**`e`**) Enables or disables switch responses to DNS queries with local virtual IP addresses. When enabled, the switch will always respond to DNS queries by providing a local virtual IP address, as long as the virtual IP address has healthy real servers with an aggregate of at least 1024 available connections (the total from each server's configured `maxcons` value, minus the server's current number of connections). When the real servers for the local virtual IP addresses are unavailable or saturated, the switch will respond to DNS requests using normal GSLB rules.
`one disable`|**`enable`** (or just **`d`**|**`e`**) Enables or disables DNS responses with only one address. At most one IP address is included in each DNS response.
`alway disable`|**`enable`** (or just **`d`**|**`e`**) Enables or disables DNS responses (with) at least one address. At least one IP address is included in each DNS response. Even if all remote sites cannot handle another request, the local VIP is returned in DNS response to eliminate long DNS time-outs caused by an empty response.
`geo disable`|**`enable`** (or just **`d`**|**`e`**) Enables or disables geographic awareness, such as the IANA table. If this option is disabled, all clients and sites will be assumed to exist in the same geographic region, allowing all sites to be eligible for each client.
`http disable`|**`enable`** (or just **`d`**|**`e`**) Enables or disables HTTP redirects to peer sites by this switch. When enabled, this switch will redirect client requests to peer sites if its own real servers fail or have reached their maximum connection limits. If disabled, the switch will not perform HTTP Redirects, but will instead drop requests for new connections and cause the client's browser to eventually issue a new DNS request.
`usern disable|enable` (or just **`d`**|**`e`**) Enables or disables an HTTP redirect to a real server name. When a site redirects a client to another site using an HTTP redirect, the client is redirected to the new site's IP address. If `usern` is enabled, the client will be redirected to the domain name specified by the remote real server name plus virtual server domain name: `<remote real server name>.<virtual server domain name>`

Table 78 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage
mincon *<minimum connections, 0-65535>* Sets the minimum number of available site connections. If the site's available sessions fall below this value, traffic won't be redirected to the site. A site is not eligible for more requests (such as DNS or HTTP redirects) once the number of available connections at a site drops below this threshold.
inter *<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.
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.
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 238).
off Turns GSLB off for this switch. Any active remote sites will still perform GSLB services with each other, but will not hand off requests to this switch.
current Displays current Global SLB configuration.

/cfg/slb/gslb/site <site-number [1-128]>

GSLB Remote Site Configuration

```
[Remote site 1 Menu]
      name    - Set remote site name
      prima   - Set primary switch IP address of remote site
      secon   - Set secondary switch IP address of remote site
      update  - Enable/disable remote site updates
      enable  - Enable remote site
      disable - Disable remote site
      delete  - Delete remote site
      current - Display current remote site configuration
```

Up to 128 remote sites can be configured.

Table 79 GSLB Remote Site Menu Options (/cfg/slb/gslb/site)

Command Syntax and Usage
name *<string, maximum 15 characters>* Sets the name of the remote site.
prima *<primary 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 *<secondary server IP address>* If the remote site is configured with a redundant switch, enter the IP address of the remote secondary switch here. If the remote site primary switch fails, the local switch will address the remote site secondary switch instead.
update disable|**enable** (or just **d**|**e**) Enables or disables remote site updates. If enabled, this switch will send regular Distributed Site State Protocol (DSSP) updates to its remote peers using HTTP port 80. If disabled, the switch will not send state updates. If your local firewall does not permit this traffic, disable the updates.
enable Enables this remote site for use with Global Server Load Balancing.
disable Disables this remote site. The switch will no longer use this remote site for Global Server Load Balancing.

Table 79 GSLB Remote Site Menu Options (/cfg/slb/gslb/site)

Command Syntax and Usage
`delete` Removes this remote site from operation and deletes its configuration.
`current` Displays the current remote site configuration.

NOTE – When `update` (above) is enabled, Global Server Load Balancing uses service port 80 on the IP interface for DSSP updates. By default, the WebOS web-based interface also uses port 80. Both services cannot use the same port. If both are enabled, configure the WebOS interface to use a different service port (see the `/cfg/sys` options under Table 41 on page 151).

/cfg/slb/sync

Config Synchronization Menu

```
[Config Synchronization Menu]
   peer    - Synch peer switch menu
   filt    - Enable/disable syncing filter configuration
   ports   - Enable/disable syncing port configuration
   prios   - Enable/disable syncing VRRP priorities
   pips    - Enable/disable syncing proxy IP addresses
   current - Display current Layer 4 sync configuration
```

Table 80 Synch Peer Switch Menu Options (/cfg/slb/synch)

Command Syntax and Usage
peer *<peer switch number: (1-4) >* Displays the Peer switch menu. To view a list of menu options, see page 218.
filt disable|enable (or just **d|e**) Enables or disables syncing filter configuration.
ports disable|enable (or just **d|e**) Enables or disables syncing Layer 4 port configuration.
prios disable|enable (or just **d|e**) Enables or disables syncing VRRP priorities.
pips disable|enable (or just **d|e**) Enables or disables the current synching proxy IP address.
current Displays the current configuration for this peer switch.

/cfg/slb/synch/peer <peer-switch [1-4]>

Peer Switch Menu

The peer switch menu for each peer switch (1-4) allows you to set the IP address of the peer switch, enable, disable, or delete the current peer switch, or display the current configuration.

```
[Peer switch 1 Menu]
      addr    - Set peer switch IP address
      enable  - Enable peer switch
      disable - Disable peer switch
      delete  - Delete peer switch
      current - Display current peer switch configuration
```

Table 81 Peer Switch Menu Options (/cfg/slb/sync/peer)

Command Syntax and Usage
addr *<IP-address>* Enter the IP address of the peer switch you want to set, in dotted decimal notation.
enable Enables the peer switch.
delete Deletes the peer switch.
current Displays current peer switch configuration.

/cfg/slb/adv

Advanced Layer 4 Configuration

```
[Layer 4 Advanced Menu]
      direct  - Enable/disable Direct Access Mode
      grace   - Enable/disable graceful real server failure
      octets  - Enable/disable octet counters
      imask   - Set virtual and real IP address mask
      mnet    - Set management network
      mmask   - Set management subnet mask
      pmask   - Set persistent mask
      timeout - Set minutes inactive session remains open
      fastage - Set session entries aged every 200 us
      secret  - Set RADIUS secret
      current - Display current Layer 4 advanced configuration
```

Table 82 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage

direct disable|enable (or just **d|e**)

Enable/disables Direct Access Mode to real servers/services. This option also allows any virtual server to load balance any real server. For more information, see "Direct Access Mode" on page 221.

grace *<real server number>*

Enables or disables graceful real server failure.

octets disable|enable (or just **d|e**)

Enables or disables octet counters.

imask *<IP subnet mask (e.g., 255.255.255.0)>*

Configures the real and virtual IP address mask using dotted decimal notation. For more information, see "Configuring the imask" on page 221.

mnet *<server 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 (e.g., 255.255.255.0)>*

Sets management subnet mask. This IP address mask is used with the `mnet` to select management traffic that is allowed access to direct real servers.

Table 82 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage
pmask *<IP subnet mask (e.g., 255.255.255.0)>* Sets persistent mask, which is used to increase granularity of persistence to a subnet level.
timeout *<even number of minutes [4-60]>* Sets the number of minutes an inactive session remains open (in even numbered increments). Every client-to-server session being load balanced is recorded in the switch's Session Binding Table. When a client makes a request, the session is recorded in the binding table, the data is transferred until the client ends the session, and the binding table entry is then removed. If a client application is abnormally terminated by the client's system, TCP connections will remain registered in the switch's binding table. In order to prevent table overflow, these orphaned entries must be aged out of the binding table. Using the `timeout` option, you can set the number of minutes to wait before removing orphan table entries. Settings must be specified in even numbered increments between 2 and 60 minutes. The default setting is 10. *Note:* A session entry can be removed before the timeout value *only* for filtered or TCP sessions in which a FIN/RST is seen. All other connections such as UDP and ICMP, as well as normal TCP sessions, must timeout through inactivity using the timeout value.
fastage *<power-of-two entries [1-256]>* Sets how many session entries are aged for each SP timer tick (200 us currently).
secret *<16 character secret>* To perform application health checking to a RADIUS server, the network administrator must configure two parameters in the switch: the `/cfg/slb/adv/secret` value and the `content` parameter with a *username:password* value. The `secret` value is a field of 16 alphanumeric characters that is used by the switch to encrypt a password during the RSA Message Digest Algorithm (MD5). It is also used by the RADIUS server to decrypt the password during verification.
current Displays the current Layer 4 advanced configuration.

Direct Access Mode

Some clients may need direct access to the real servers, to, for example, monitor a real server from a management workstation. When Direct Access Mode (`/cfg/slb/adv/direct`) is enabled on a switch, any client can communicate with any real server to its load-balanced service. Also, in Direct Access Mode, any number of virtual services can be configured to load balance a real service.

NOTE – When Direct Access Mode is enabled on a server, Layer 4 port mapping and default gateway load balancing is not supported.

Traffic sent directly to real server IP addresses is excluded from load balancing decisions. The same clients may also communicate to the virtual server IP address and have their requests load balanced.

Configuring the imask

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

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

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

/cfg/vrrp
VRRP Configuration

```
[Virtual Router Redundancy Protocol Menu]
      vr      - VRRP Virtual Router menu
      if      - VRRP Interface menu
      track   - VRRP Priority Tracking menu
      on      - Globally turn VRRP ON
      off     - Globally turn VRRP OFF
      current - Display current VRRP configuration
```

Virtual Router Redundancy Protocol (VRRP) support on Alteon WebSystems 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.

Alteon WebSystems has extended VRRP to include virtual servers as well, allowing for full active/active redundancy between its Layer 4 switches.

Table 83 Virtual Router Redundancy Protocol Options (/cfg/vrrp)

Command Syntax and Usage
vr *<virtual router number [1-1024]>* Displays the VRRP virtual router menu. This menu is used for configuring up to 1024 virtual routers on this switch. To view menu options, see page 223.
if *<interface number [1-1024]>* Displays the VRRP virtual router interface menu. This menu is for setting VRRP authentication parameters for up to 1024 IP interfaces. To view menu options, see page 228.
track Displays the VRRP tracking menu. This menu is used for weighting the criteria used when modifying priority levels in the master router election process.
on Globally enables VRRP on this switch.

Table 83 Virtual Router Redundancy Protocol Options (/cfg/vrrp)

Command Syntax and Usage
off Globally disables VRRP on this switch.
current Displays the current VRRP parameters.

/cfg/vrrp/vr <router-number>

Virtual Router Configuration

```
[VRRP Virtual Router 1 Menu]
      vrid    - Set virtual router ID
      addr    - Set IP address
      if      - Set interface number
      prio    - Set renter priority
      adver   - Set advertisement interval
      preem   - Enable/disable preemption
      share   - Enable/disable sharing
      track   - Priority tracking menu
      enable  - Enable virtual router
      disable - Disable virtual router
      delete  - Delete virtual router
      current - Display current VRRP virtual router configuration
```

This menu is used for configuring one of 1024 possible 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.

Table 84 VRRP Virtual Router Options (/cfg/vrrp/vr)

Command Syntax and Usage
vrid *<virtual router ID [1-255]>* Defines the virtual router ID. This is used in conjunction with `addr` (below) to define a virtual router on this switch. To create a pool of VRRP-enabled routing devices that can provide redundancy to each other, each participating VRRP device must be configured with the same virtual router: one that shares the same `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. All `vrid` values must be unique within the VLAN to which the virtual router's IP interface (see `if` below) belongs.
addr *<IP address (e.g., 192.x4.17.101)>* Defines the IP address for this virtual router using dotted decimal notation. This IP address (`addr`) is used in conjunction with the `vrid` (above) to configure the same virtual router on each participating VRRP device.
if *<interface number [1-1024]>* Selects a switch IP interface (between 1 and 1024). 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 that has assumed master routing authority. This preemption occurs even if the `preem` option below is disabled.
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 *<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.

Table 84 VRRP Virtual Router Options (/cfg/vrrp/vr)

Command Syntax and Usage
preem disable|**enable** (or just **d**|**e**) Enables or disables master preemption. When enabled, if this virtual router is in backup mode but has a higher priority than the current master, this virtual router will preempt the lower priority master and assume control. 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).
share disable|**enable** (or just **d**|**e**) Enables or disables virtual router sharing, an Alteon WebSystems proprietary extension to VRRP. When enabled, this switch will process any traffic addressed to this virtual router, even when in backup mode.
track Displays the VRRP priority tracking menu for this virtual router. Tracking is an Alteon WebSystems proprietary extension to VRRP, used for modifying the standard priority system used for electing the master router. Sharing (`share`) should be disabled in order for tracking to be used effectively. For more information, see"Virtual Router Priority Tracking Configuration" on page 226.
enable Enables this virtual router.
disable Disables this virtual router.
delete Deletes this virtual router from the switch configuration.
current Displays the current configuration information for this virtual router.

/cfg/vrrp/vr <router-number>/track

Virtual Router Priority Tracking Configuration

```
[VRRP Virtual Router 1 Priority Tracking Menu]
      vrs     - Enable/disable tracking other virtual routers
      ifs     - Enable/disable tracking other interfaces
      ports   - Enable/disable tracking VLAN switch ports
      l4pts   - Enable/disable tracking L4 switch ports
      reals   - Enable/disable tracking L4 real servers
      hsrp    - Enable/disable tracking HSRP
      current - Display current VRRP virtual router configuration
```

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 229).

Criteria are tracked dynamically, continuously updating virtual router priority levels when enabled. If the virtual router preemption option (see `preem` in Table 84 on page 224) 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 extended virtual routers, or "virtual server routers," which perform Layer 4 Server Load Balancing functions in addition to their standard VRRP operation. A virtual *server* router is defined as any virtual router whose IP address (`addr`) is the same as any configured virtual server IP address.

Table 85 VRRP Priority Tracking Options (/cfg/vrrp/vr #/track)

Command Syntax and Usage
vrs disable|enable (or just **d|e**) When enabled, the priority for this virtual router will be increased for each other 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.
ifs disable|enable (or just **d|e**) When enabled, the priority for this virtual router will be increased for each other IP interface active on this switch. An IP interface is considered active when there is at least one active port on the same VLAN. This helps elect the virtual routers with the most available routes as the master.

Table 85 VRRP Priority Tracking Options (/cfg/vrrp/vr #/track)

Command Syntax and Usage
ports disable|enable (or just **d|e**) When enabled, the priority for this virtual router will be increased for each active port on the same VLAN. This helps elect the virtual routers with the most available ports as the master.
l4pts disable|enable (or just **d|e**) When enabled for virtual server routers, the priority for this virtual router will be increased for each physical switch port which has active Layer 4 processing on this switch. This helps elect the main Layer 4 switch as the master.
reals disable|enable (or just **d|e**) When enabled for virtual server routers, the priority for this virtual router will be increased for each healthy real server behind the virtual server IP address of the same IP address as the virtual router on this switch. This helps elect the switch with the largest server pool as the master, increasing Layer 4 efficiency.
hsrp disable|enable (or just **d|e**) Hot Standby Router Protocol (HSRP) is used with some types of routers for establishing router failover. In networks were HSRP is used, enable this switch option to increase the priority of this virtual router for each Layer 4 client-only port that receives HSRP advertisements. Enabling `hsrp` helps elect the switch closest to the master HSRP router as the master, optimizing routing efficiency.
current Displays the current configuration for priority tracking for this virtual router.

/cfg/vrrp/if <interface-number [1-1024]>

VRRP Interface Configuration

NOTE – The *interface-number* (1 to 1024) represents the IP interface on which authentication parameters must be configured.

```
[VRRP Interface 1 Menu]
     auth    - Set authentication types
     passw   - Set plain-text password
     delete  - Delete interface
     current - Display current VRRP interface configuration
```

This menu is used for configuring VRRP authentication parameters for the IP interfaces used with the virtual routers.

Table 86 VRRP Interface Options (/cfg/vrrp/if)

Command Syntax and Usage
auth none|password Defines the type of authentication: `none` No authentication used. `password` Password authentication will be used.
passw <*password*> Defines a plain-text password up to eight characters long. This password will be added to each VRRP packet transmitted by this interface when password authentication is chosen (see `auth` above).
delete Clears the authentication configuration parameters for this IP interface. The IP interface itself is not deleted.
current Displays the current configuration for this IP interface's authentication parameters.

/cfg/vrrp/track

VRRP Tracking Configuration

```
[VRRP Tracking Menu]
      vrs     - Set priority increment for virtual router tracking
      ifs     - Set priority increment for IP interface tracking
      ports   - Set priority increment for VLAN switch port tracking
      l4pts   - Set priority increment for L4 switch port tracking
      reals   - Set priority increment for L4 real server tracking
      hsrp    - Set priority increment for HSRP tracking
      current - Display current VRRP Priority Tracking configuration
```

This menu is used for setting weights for the various criteria used to modify priority levels during the master router election process. Each time one of the tracking criteria is met ("Virtual Router Priority Tracking Configuration" on page 226), the priority level for the virtual router is increased by an amount defined through this menu.

Table 87 VRRP Tracking Options (/cfg/vrrp/track)

Command Syntax and Usage
vrs *<0-254>* Defines the priority increment value (0 through 254) for virtual routers in master mode detected on this switch.
ifs *<0-254>* Defines the priority increment value (0 through 254) for active IP interfaces detected on this switch.
ports *<0-254>* Defines the priority increment value (0 through 254) for active ports on the virtual router's VLAN.
l4pts *<0-254>* Defines the priority increment value (0 through 254) for physical switch ports with active Layer 4 processing.
reals *<0-254>* Defines the priority increment value (0 through 254) for healthy real servers behind the virtual server router.

Table 87 VRRP Tracking Options (/cfg/vrrp/track)

Command Syntax and Usage
hsrp *<0-254>* Defines the priority increment value (0 through 254) for switch ports with Layer 4 client-only processing that receive Hot Standby Router Protocol (HSRP) broadcasts.
current Displays the current configuration of priority tracking increment values.

NOTE – These priority tracking options only define increment values. These options do not affect the VRRP master router election process until options under the VRRP Virtual Router Priority Tracking Menu (see page 226) are enabled.

CHAPTER 8

The Operations Menu

The Operations Menu is generally used for commands that immediately affect operation of the switch, but do not alter permanent switch configurations.

For example, you can use the Operations Menu to immediately disable a port (without the need to apply or save the change), with the understanding that when the switch is reset, the port returns to its normally configured operation.

To make permanent changes to switch configurations, use commands under Chapter 7, "The Configuration Menu".

/oper
Operations Menu

```
[Operations Menu]
      port    - Operational Port menu
      sp      - Operational SP menu
      lcm     - Operational LCM menu
      sfm     - Operational SFM menu
      slb     - Operational Layer 4 menu
      vrrp    - Operational Virtual Router Redundancy menu
      swkey   - Enter key to enable software feature
      rmkey   - Enter software feature to be removed
```

The commands of the Operations Menu enable you to alter switch operational characteristics without affecting switch configuration.

Table 88 Operations Menu Options (/oper)

Usage
port *<port as number 1-16>* Displays the Operational Port menu. To view menu options, see page 233.
sp Displays the Operational Switch Processor menu. To view menu options, see page 234.
lcm Displays the Operational Line-Card Module menu. To view menu options, see page 235.
sfm Displays the Operational Switch Fabric Module menu. To view menu options, see page 235.
slb Displays the Operational Layer 4 menu. To view menu options, see page 236.
vrrp Displays the Operational Virtual Router Redundancy menu. To view menu options, see page 237.
swkey *<16-hex-digit key to enable software feature>* Enter key to enable software feature. For more information, see page 238.
rmkey *<software feature to be removed (that is, L4)>* Enter software feature to be removed. For more information, see page 239.

/oper/port <slot-number port-number> Operations-Level Port Options

```
[Operations Port 1 Menu]
      up      - Bring port up
      down    - Shut port down
      rmon    - Up/down RMON for port
      current - Current port state
```

NOTE – The `ena` and `dis` commands, previously in Release 6.0, have been renamed `up` and `down`. The `ena` and `dis` commands always apply to *configuration*, whereas `up` and `down` always apply to *operation.*

Operations-level port options are used for temporarily disabling or enabling a port, and for changing RMON status on a port.

Table 89 Operations-Level Port Menu Options (/oper/port)

Command Syntax and Usage
up Temporarily brings up the port. The port will be returned to its configured operation mode when the switch is reset.
down Temporarily brings down the port. The port will be returned to its configured operation mode when the switch is reset.
rmon Temporarily brings up/down RMON on the port. The port will be returned to its configured operation mode when the switch is reset.
current Displays the current operational status of the port.

/oper/sp
Operations-Level Switch Processor Options

```
[SP Operations Menu]
      up      - Bring up one or all SPs
      down    - Shut down one or all SPs
      post    - Run POST on one or all SPs
      dump    - Take a snapshot dump of an SP
```

NOTE – The following commands will bring up or down all ports on the switch processor.

Table 90 Operations-Level Switch Processor Menu Options (/oper/sp)

Command Syntax and Usage
up *<SP as port (e.g. b12) or index (e.g. x5)>* Temporarily brings up the specified switch processor.
down *<SP as port (e.g. b12) or index (e.g. x5)>* Temporarily brings down the specified switch processor.
post *<SP as port (e.g. b12) or index (e.g. x5)>* Runs Power On Self Test (POST) on the specified switch processor.
dump *<SP as port (e.g. b12) or index (e.g. x5)>* Takes a snapshot dump of the specified switch processor.

/oper/lcm <module-number> Operations-Level Line Card Module Options

```
[LCM Operations Menu]
      up      - Bring an LCM up
      down    - Shut an LCM down
```

Table 91 Operations-Level Switch Processor Menu Options (/oper/lcm)

Command Syntax and Usage
up *<LCM index>* Temporarily brings up the specified line-card module.
down *<LCM index>* Temporarily brings down the specified line card module.

/oper/sfm <module-number> Operations-Level Switch Fabric Module Options

```
[SFM Operations Menu]
      up      - Bring an SFM up
      down    - Shut an SFM down
```

Table 92 Operations-Level Switch Fabric Module Menu Options (/oper/sfm)

Command Syntax and Usage
up *<SFM index>* Temporarily brings up the specified switch fabric module.
down *<SFM index>* Temporarily brings down the specified switch fabric module.

/oper/slb
Operations-Level SLB Options

```
[Server Load Balancing Operations Menu]
      synch   - Synchronize SLB, FILT, and VRRP configuration on peers
      enable  - Enable real server
      disable - Disable real server
      clear   - Clear session table on port
      current - Current layer 4 operational state
```

When the optional Layer 4 software is enabled, the operations-level Server Load Balancing options are used for temporarily disabling or enabling real servers and synchronizing the configuration between the switches.

The options are described in the following table.

Table 93 Server Load Balancing Operations Menu Options (/oper/slb)

Command Syntax and Usage
synch Synchronizes the SLB, filter, and VRRP configuration on a peer switch (a switch that owns the IP address). To take effect, VRRP must be globally enabled on the peer switch.
enable *<real server number [1-4096]>* Temporarily enables a real server. The real server will be returned to its configured operation mode when the switch is reset.
disable *<real server number [1-4096]>* Temporarily disables a real server, removing it from operation within its real server group and virtual server. The real server will be returned to its configured operation mode when the switch is reset.
clear *<port as slot and port, like b12>* Clears the session table for a specific port, and allow port filter changes to take effect immediately. Note: This disrupts current Server Load Balancing and Application Redirection sessions.
current Displays the current SLB operational state.

/oper/vrrp
Operations-Level VRRP Options

```
[VRRP Operations Menu]
        back    - Set virtual router to backup
```

This menu is used to force a master virtual router to become backup router.

Table 94 Virtual Router Redundancy Operations Menu Options (/oper/vrrp)

Command Syntax and Usage
back *<virtual-router-number [1-1024]>* Forces the specified master virtual router on this switch into backup mode. This is generally used for passing master control back to a preferred switch once the preferred switch has been returned to service after a failure. When this command is executed, the current master gives up control and initiates a new election, by temporarily advertising its own priority level as 0 (lowest). After the new election, the virtual router that had been forced into backup mode by this command, will resume master control in the following cases: ■ This switch owns the virtual router (the IP addresses of the virtual router and its IP interface are the same) ■ This switch's virtual router has a higher priority and preemption is enabled. ■ There are no other virtual routers available to take master control.

/oper/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 WebOS 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 2, "The Command-Line Interface").**

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

```
Main# oper
```

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

```
Operations# swkey
```

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

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

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

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

/oper/rmkey
Removing Optional Software

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

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

```
Operations# rmkey
```

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

```
Enter Software Feature to be removed [L4]: l4
```

CHAPTER 9

The Boot Options Menu

To use the Boot Options Menu, you must be logged in to the switch as the administrator. The Boot Options Menu provides options for:

- Selecting a switch software image to be used when the switch is next reset
- Selecting a configuration block to be used when the switch is next reset
- Downloading a new software image to the switch via TFTP

To access the Boot Options Menu, at the Main Menu prompt, enter:

```
Main# boot
```

The Boot Options Menu is displayed:

```
[Boot Options Menu]
      image   - Select software image to use on next boot
      config  - Select config block to use on next boot
      tftp    - Download new software image via TFTP
      reset   - Reset switch [WARNING: Restarts Spanning Tree]
      cur     - Display current boot options
```

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

Updating the Switch Software Image

The switch software image is the executable code running on the switch. A version of the image ships with the switch, and comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch.

Upgrading the software image on your switch requires the following:

- Loading the new image from the CD-ROM onto a TFTP server on your network
- Downloading the new image from the TFTP server to your switch
- Selecting the new software image to be loaded into switch memory the next time the switch is reset

Downloading New Software to Your Switch

The switch can store up to two different software images, called `image1` and `image2`, 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 183).

When the above requirements are met, use the following procedure to download the new software to your switch.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# tftp
```

2. **Enter the name of the switch software to be replaced:**

   To minimize mistakes, you must type the full name (“image1”, “image2”, or “boot)”.

```
Enter software image to be replaced [image1|image2|boot]:
```

3. **Enter the hostname or IP address of the TFTP server.**

```
Enter hostname or IP address of TFTP server:
```

4. **Enter the name of the new software file on the server.**

```
Enter name of file on TFTP server:
```

   The exact form of the name will vary by TFTP server. However, the file location is normally relative to the TFTP directory (usually `/tftpboot`).

5. **The system prompts you to confirm your request.**

   You should next select a software image to run, as described below.

Selecting a Software Image to Run

You can select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# image
```

2. **Enter the name of the image you want the switch to use upon the next boot.**

   The system informs you 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 software "image2" on next boot.
Specify new image to use [image1|image2]:
```

   To minimize mistakes, you must type the full name (“image1”, “image2”, or “boot”).

Selecting a Configuration Block

When you make configuration changes to the switch, you must save the changes so that they are retained beyond the next time the switch is reset. When you perform the `save` command, your new configuration changes are placed in the *active* configuration block. The previous configuration is copied into the *backup* configuration block.

There is also a *factory* configuration block. This holds the default configuration set by the factory when your switch was constructed. Under certain circumstances, it may be desirable to reset the switch configuration to the default. This can be useful when a custom-configured switch is moved to a network environment where it will be reconfigured for a different purpose.

Use the following procedure to set which configuration block you want the switch to load the next time it is reset:

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# config
```

2. **Enter the name of the configuration block you want the switch to use:**

The system informs you of which configuration block is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use active configuration block on next reset.
Specify new block to use [active|backup|factory]:
```

Resetting the Switch

You can reset the switch to make your software image file and configuration block changes occur.

NOTE – Resetting the switch causes the Spanning-Tree Protocol to restart. This process can be lengthy, depending on the topology of your network.

To reset the switch, at the `Boot Options#` prompt, enter:

```
>> Boot Options# reset
```

You are prompted to confirm your request.

CHAPTER 10

The Maintenance Menu

The Maintenance Menu is used to manage dump information and 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]
      dump    - FLASH Dump Manipulation Menu
      sys     - System Maintenance Menu
      sp      - SP Maintenance Menu
      fdb     - Forwarding Database Manipulation Menu
      debug   - Debugging Menu
      diag    - Diagnostic Menu
      arp     - ARP Cache Manipulation Menu
      route   - IP Route Manipulation Menu
      panic   - Dump MP information to FLASH and reboot
      tsdump  - Tech support dump
```

Dump information contains internal switch state data that is written to flash memory on the switch after any one of the following occurs:

- The switch administrator forces a switch *panic*. The `panic` option, found in the Maintenance Menu, causes the switch to dump state information to flash memory, and then causes the switch to reboot.
- The switch administrator enters the switch reset key combination on a device attached to the console port. The switch reset key combination is <Shift-Ctrl-6>.
- The watchdog timer forces a switch reset. The purpose of the watchdog timer is to reboot the switch if the switch software freezes.
- The switch detects a hardware or software problem that requires a reboot.

Table 95 Maintenance Menu Options (/maint)

Command Syntax and Usage
dump Displays the Flash Dump Manipulation Menu.To view menu options, see page 249.
sys Displays the System Maintenance Menu. To view menu options, see page 252.
sp Displays the SP Maintenance Menu. To view menu options, see page 253.
fdb Displays the Forwarding Database Manipulation Menu. To view menu options, see page 253.
debug Displays the Debugging Menu. To view menu options, see page 255.
diag Displays the Diagnostic Menu. To view menu options, see page 257.
arp Displays the ARP Cache Manipulation Menu. To view menu options, see page 258.
route Displays the IP Route Manipulation Menu. To view menu options, see page 259.
panic Dumps MP information to FLASH and reboots. For more information, see page 260.
tsdump Dumps all switch information, statistics, and configurations.

/maint/dump
Flash Dump Manipulation Options

```
[Dump Menu]
      current - Display currently-available FLASH dumps
      uuout   - Uuencode a FLASH dump to standard out
      tftp    - TFTP a FLASH dump to tftpserver
      delete  - Delete a FLASH dump
      mksave  - Mark a FLASH dump as "saved"
```

Table 96 Flash Dump Maintenance Menu Options (/maint/dump)

Command Syntax and Usage

current
Displays currently-available FLASH dumps. Information similar to the following is shown:
`index  MP/SP  ver     date and time taken        cause     state`
`-----  -----  ---  -------------------------  ---------  --------`
`  0     MP     1   11:08:56 Thu May 25, 2000     W/DOG     unread`
`----------------------------------------------------------------------`

uuout Uuencodes a FLASH dump to standard out. For more information, see page 250.
tftp TFTPs a FLASH dump to a TFTP server. For more information, see page 251.
delete Deletes a FLASH dump. For more information, see page 251.
mksave Marks a FLASH dump as “saved.” A new dump is called UNREAD, and it will not be overlaid with a later dump until it has been read (with uuout or tftp), or deleted. Marking a dump as SAVED will prevent it from being overlaid, even after it is read.

/maint/dump/uuout

Uuencode Flash Dump

Using this command, dump information is presented in uuencoded format. This format makes it easy to capture the dump information as a file or a string of characters. You can then contact Alteon WebSystems Customer Support for help in analyzing the information.

If you want to capture dump information to a file, set your communication software on your workstation to capture session data prior to issuing the `uuout` command. This will ensure that you do not lose any information. Once entered, the `uuout` command will cause data to be displayed on your screen and copied into the file.

Using the `uuout` command, dump information can be read multiple times. The command does not cause the information to be updated or cleared from flash memory.

NOTE – Dump information is not cleared automatically. In order for any subsequent dump information to be written to flash memory, you must manually clear the dump region. For more information on clearing the dump region, see "Clearing Dump Information" on page 251.

To access dump information, at the `Maintenance#` prompt, enter:

```
Maintenance# dump/uuout
```

The dump information is displayed on your screen and, if you have configured your communication software to do so, captured to a file. If the dump region is empty, the following appears:

```
No FLASH dump available.
```

/maint/dump/tftp <server/filename>

TFTP System Dump Put

Use this command to put (save) the system dump via TFTP.

NOTE – If the TFTP server is running SunOS or the Solaris operating system, the specified target dump file must exist *prior* to executing the `tftp` command, and must be writable (set with proper permission, and not locked by any application). The contents of the specified file will be replaced with the current dump data.

To save dump information via TFTP, at the `Dump#` prompt, enter:

```
Maintenance# tftp server filename
```

Where *server* is the TFTP server IP address or hostname, and *filename* is the target dump file.

/maint/dump/delete

Clearing Dump Information

To clear dump information from FLASH memory, at the `Maintenance#` prompt, enter:

```
Maintenance# delete
```

/maint/sys
System Maintenance Options

This menu is reserved for use by Alteon WebSystems Customer Support. The options are used to perform system debugging

```
[System Maintenance Menu]
      sw      - Display software build information
      flags   - Set NVRAM flag word; "bits" to list bits
      tmask   - Set current MP trace mask
      sptmask - Set default SP trace mask
```

Table 10-1 System Maintenance Menu Options (/maint/sys)

Command Syntax and Usage

sw
Displays software build information, such as the following:
`Software build information:`
`  Software build information:`
`  FW_VERSION: #2 Thu Jun 22 14:05:57 PDT 2000`
`  FW_COMPILE_TIME: 14:05:57`
`  FW_COMPILE_BY: devrcs`
`  FW_SANDBOX: /projects/sandbox/gne_7_0_24/src`
`  FW_COMPILER: gcc version 2.7.2`

flags Sets an NVRAM flag word that can be used by Customer Support to turn on special messages or additional information used for debugging purposes.
tmask Sets the current Management Processor trace mask, which can be used by Customer Support to turn on special messages or additional information used for debugging purposes.
sptmsk Sets the default Switch Processor trace mask, which can be used by Customer Support to turn on special messages or additional information used for debugging purposes.

/maint/sp <sp-index>
Switch Processor Maintenance Options

```
[Maint SP index 2 (port A3) Menu]
      tmask   - Temporarily set SP's trace mask
```

Table 97 Switch Processor Maintenance Menu Options (/maint/sp)

Command Syntax and Usage
tmask Temporarily sets the switch processor's trace mask, which can be used by Customer Support to turn on special messages or additional information used for debugging purposes.

/maint/fdb
Forwarding Database Options

```
[FDB Manipulation Menu]
      find    - Show a single FDB entry by MAC address
      port    - Show FDB entries for a single port
      vlan    - Show FDB entries for a single VLAN
      dump    - Show all FDB entries
      count   - Show count of FDB entries
      del     - Delete an FDB entry
      clear   - Clear entire FDB
```

The Forwarding Database Manipulation Menu can be used to view forwarding database entries, find entries based on MAC address, port, VLAN, and switch processor. It can also be used to delete a MAC address from the forwarding database or clear the entire forwarding database. This is helpful in identifying problems associated with MAC address learning and packet forwarding decisions.

Table 98 FDB Manipulation Menu Options (maint/fdb)

Command Syntax and Usage
find <*MAC-addr*> [<*VLAN*>] Displays a single database entry by its MAC address. Enter the MAC address using the format, `xx:xx:xx:xx:xx:xx`. For example, `08:00:20:12:34:56`. You can also enter the MAC address using the format, `xxxxxxxxxxxx`. For example, `080020123456`.
port <*port as slot and port, like b12*> Displays all FDB entries for a particular port.
vlan <*VLAN number 1-4094*> Displays all FDB entries on a single VLAN.
dump Displays all entries in the Forwarding Database. For more information, see page 73.
count Displays the count of all FDB entries.
del <*MAC-addr*> *[*<*VLAN*>*]* Removes a single FDB entry.
clear Clears the entire Forwarding Database from switch memory.

/maint/debug
Debugging Options

```
[Miscellaneous Debug Menu]
      tbuf    - Display MP trace buffer
      xtbuf   - Display MP XTRACE buffer
      sptbuf  - Display SP trace buffer
```

The Miscellaneous Debug Menu displays trace buffer information about events that can be helpful in understanding switch operation. You can view the following information using the debug menu:

- Events traced by the Management Processor (MP)
- Events traced by a Switch Processor (SP)
- Events traced to a buffer area when a reset occurs

If the switch resets due to a crash, the MP trace buffer is saved into the MP dump, and the SP trace buffer is saved into the SP dump, both of which are written to FLASH. No trace buffers are saved if the switch is powered off.

The output from these commands can be interpreted by the Alteon WebSystems Customer Support organization.

Table 99 Miscellaneous Debug Menu Options (maint/debug

Command Syntax and Usage
tbuf Displays the Management Processor trace buffer. The MP trace buffer displays normal traces, which contain detailed operational events that may give details about what events were occurring at the moment of a crash. Header information similar to the following is shown: `MP trace buffer at 18:46:10 Thu Jun 29, 2000; mask: 0x01dff6df` The buffer information is displayed after the header.
xtbuf Displays the Management Processor xtrace buffer. The MP xtrace buffer displays more significant operational events going back further in time. Header information similar to the following is shown: `MP XTRACE buffer at 18:48:16 Thu Jun 29, 2000` The buffer information is displayed after the header
sptb *<slot-number port-number, like b12>* Displays the Switch Processor trace buffer. The SP trace buffer displays normal traces, which contain detailed operational events that may give details about what events were occuring at the moment of a switch processor crash. Header information similar to the following is shown: `SP index 0 (A1) at 18:48:26 Thu Jun 29, 2000; mask: 0x01dff6df` The buffer information is displayed after the header.

/maint/diag
Diagnostics Menu

```
[Diagnostics Menu]
      info     - Show options and status summary
      rtc      - Real Time Clock Test
      spmem    - SP Memory Menu
      spcpu    - SP CPU Menu
      phy      - Phy Test Menu (static frames)
      fpg      - Frame Pattern Generator Menu (dynamic frames)
      burnin   - System Burn-In Menu
      clear    - Clear counters, free memory
```

CAUTION—The diagnostics are tools for analyzing hardware problems and *will halt all normal switch traffic*. If you believe you have a hardware problem, contact Alteon Technical Support, or Alteon Customer Service (see “Contacting Alteon WebSystems” on page 13). You may be asked to run Diagnostics.

The diagnostics menu is available only when logged in as “admin” and connected via console. This menu is not accessible via telnet connection. To start diagnostics, the switch must be rebooted (/b/r hard or a power cycle) and the Esc key pressed on the system console during boot to enter maintenance mode.

/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
      refpt   - Show ARP entries referenced by a single port
      vlan    - Show ARP entries on a single VLAN
      add     - Add a permanent ARP entry
      delete  - Delete an ARP entry
      clear   - Clear ARP cache
      dump    - Show all ARP entries
      addr    - Show ARP address list
```

Table 100 Address Resolution Protocol Menu Options (maint/arp)

Command Syntax and Usage

find *<IP address>*

Shows a single ARP entry by IP address.

port *<slot-number port-number>*

Shows ARP entries on a single port.

refpt *<slot-number port-number>*

Shows all ARP entries referenced by a single port.

vlan *<VLAN ID>*

Shows ARP entries on a single VLAN.

add *<IP address>*

Adds a permanent ARP entry from switch memory.

delete *<IP address>*

Removes a single ARP entry from switch memory.

clear

Clears the entire ARP list from switch memory.

dump

Shows all ARP entries.

addr

Shows the list of IP addresses for which the switch will respond to 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 71.

/maint/route
IP Route Manipulation

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

Table 101 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 single gateway.
type *<type* **[indirect|direct|local|broadcast|martian|multicast]**> Shows routes of a single type.
tag *<type* **[fixed|static|snmp|addr|rip|icmp|broadcast|martian|multicast|dynamic]**> Shows routes of a single tag.
if *<interface number [1-1024]>* 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 76.

/maint/panic
Panic Command

The `panic` command causes the switch to immediately dump state information to flash memory and automatically reboot.

To select `panic`, at the `Maintenance#` prompt, enter:

```
Maintenance# panic
```

Enter **y** to confirm the command:

```
Confirm dump and reboot [y/n]: y
```

The following messages are displayed:

```
Starting system dump...done.

Reboot at 11:54:08 Thursday May 31, 2000...

Boot version 1.0.1

Alteon 708

Rebooted because of console PANIC command.

Booting complete 11:55:01 Thursday May 31, 2000:
```

/maint/tsdump
Tech Support Dump

The tsdump command is used to provide Alteon technical support with all available system information. It causes the switch to dump the following information to flash memory:

- /info/dump
- /stats/dump
- /cfg/dump

To select `tsdump`, at the `Maintenance#` prompt, enter:

```
Maintenance# tsdump
```

Enter **y** to confirm the command:

```
Confirm dumping all information, statistics, and configuration
[y|n]: y
```

All information about the switch will be displayed. Save the dumped information from your console into a text file, and then email it to Alteon Technical Support for evaluation.

Unscheduled System Dumps

If there is an unscheduled system dump to flash memory, the following message is displayed when you log on to the switch:

```
Note: A system dump exists in FLASH. The dump was saved
      at 13:43:22 Fri Jun 23, 2000. Use /maint/dump/uuout to
      extract the dump for analysis and /maint/dump/delete to
      clear the FLASH region. The region must be cleared
      before another dump can be taken.
```

CHAPTER 11

Troubleshooting

This chapter describes the most common problems that might occur with the switch, lists the probable causes for the problems, and defines possible solutions.

Definitions

- **Management Processor (MP)**

  The processor that handles management of the switch. It processes the CLI, Telnet, SNMP operation, and Spanning-Tree.

- **Switch Processor (SP)**

  The switch processor that processes both switched user frames and switched management frames.

- **Forwarding Database (FDB)**

  This is the database of learned and being-learned MAC addresses.

- **Spanning-Tree Protocol (STP)**

  The IEEE 802.1d specified loop prevention protocol widely used in Ethernet bridge networks.

- **Bridge Protocol Data Unit (BPDU)**

  Frames used to convey Spanning-Tree information to form a loop-free network topology.

System Problems

Switch Management Problems

Cannot `ping` a switch IP interface. Cannot Telnet to a switch IP interface. MIB Browser cannot discover the switch. The switch does not send SNMP traps.

Possible Causes

- Incorrect switch IP interface configuration
- Link state of the port the `ping`ing station is connected to is the in the "down" state
- Spanning-Tree port state is not in "forwarding" state
- Incorrect SNMP community strings
- Trap server is not configured
- Switch IP interface address is used by some other device in the network

Actions

- Check `/cfg/ip/current` to be sure the switch IP interface addresses, subnet masks, and default gateways are correctly configured, and that the IP interfaces are enabled.
- Check `/info/link` to be sure the management port link is in the "up" state.
- Check `/info/stp` to be sure port Spanning-Tree is in "forwarding" state.
- Check `/cfg/snmp/current` to be sure SNMP community strings are correct.
- Check `/cfg/snmp/current` to be sure the Trap server is specified.
- Check for duplicate IP address and correct if necessary.

Link Problems

Green link LED does not come on. Link state is in "down" state from the CLI (`/info/link`).

Potential Causes

- Port Configuration mismatch between the switch and the other device
- Different version of Link Negotiation used between the switch and the other device
- Bad or incorrect cable

Actions

- If ports are configured with specific values such as 100Mbs speed, then make sure the other device is configured the same way.
- Port Configuration: Make sure both the switch port and the other device are configured with the same negotiation mode. If the switch port is configured with either Speed or Duplex mode in “auto,” the other device must have the same configuration.
- Check the cabling between the switch and the other device. If the other device is a workstation, straight through cable should be used. However, if it is either another switch or a hub, a cross-over cable should be used unless there is an “uplink” enable/disable switch used instead on the switch or hub.

Table 102 Pin-outs for Crossover cable

pin 1 ---------- pin 3
pin 2 ---------- pin 6
pin 3 ---------- pin 1
pin 6 ---------- pin 2

NOTE – These pin-outs are for the 10/100 Mbps physical ports only.

- Check link status in `/info/link`. If link state is “up”, then the problem is a bad LED.

Switch Boot Failure

The switch will not boot.

Possible Causes

- Corrupted firmware
- Firmware and configuration was corrupted when rebooting with an older firmware image

Actions

Replace the corrupted firmware by performing a serial download of a new binary firmware image.

NOTE – The procedure for serial download is different from the procedure for TFTP download.

This procedure requires the following:

- A computer running terminal emulation software
- A standard serial cable with a male DB9 connector (see your switch hardware installation guide for specifics)
- A *binary* switch firmware image (*not* the `tfpt` file used for TFTP download)

Procedure

1. **Using the serial cable, connect the computer to the switch Console port (Serial Port on some models).**
2. **Make sure that the new binary firmware file is available on the computer.**
3. **Start your terminal emulation software and set the communication parameters:**

Table 103 Console Configuration Parameters

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

4. **Turn on the switch power and press <Shift-X> while the switch is first attempting to boot.**

   When performed correctly, the following message appears:

```
Xmodem flash download 1.0.5
To download to flash use xmodem at 115200 baud
Power cycle to end xmodem.
```

5. **Reconfigure your terminal emulation software for the following parameters:**

Parameter	Value
Baud Rate	115,200
Data Bits	8
Parity	None
Stop Bits	1

6. **Set the file transfer mode to Xmodem.**

7. **Transfer the binary firmware image file to the switch.**

   This process can take three or four minutes to complete. When finished, the message “done” will appear on your terminal.

8. **Disconnect the terminal emulation session and reconfigure your terminal emulation software for normal switch connection parameters:**

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

9. **Reconnect the terminal session to the switch.**

10. **Turn the switch power off, and then back on again.**

    The switch should now boot normally.

Switching Problems

This section lists the most common switching problems, their causes, and solutions.

Connectivity Problems

Client “A” on port 1 cannot connect to server “B” on port 2.

Potential Causes

- Incorrect configuration of client/server machines: the IP address is wrong.
- Ports 1 or 2 may be down (link down).
- Spanning-Tree Port State is not in “forwarding” state.
- Frames from either “A” or “B” are received with errors or not transmitted due to error conditions on outgoing port.
- MAC Address of either “A” or “B” is learned incorrectly from ports other than 1 and 2.

Actions

- Check `/info/link` to be sure link state is up.
- Check `/info/stp` to be sure Spanning-Tree Port is in “forwarding” state.
- Check port interface statistics (`/stats/port` *port-number*`/if`) to see whether `ifInErrors`, `ifInDiscards`, `ifOutErrors`, or `ifOutDiscards` are incrementing.
  - `ifInErrors`: MAC errors
  - `ifInDiscards`: STP blocking state, filtering, frame errors, PCI busy
  - `ifOutErrors`: not used
  - `ifOutDiscards`: due to backup on link
- Check port `dot3` statistics (`/stats/port 1/ether`) for Ethernet specific errors.
- Search MAC addresses for “A” and “B” from the FDB. For example, if A’s MAC address is 00:00:00:00:00:01 and B’s is 00:00:00:00:00:02, search for A’s MAC address by typing the following from the CLI: **`/info/fdb/find 00:00:00:00:00:01`**

  Output similar to the following example should be displayed.

```
   MAC Address      VLAN   Port   State
00:00:00:00:00:01     1     A3     FWD
```

Spanning-Tree Protocol Problems

The topology in the following figure is used to illustrate the STP problems in this section.

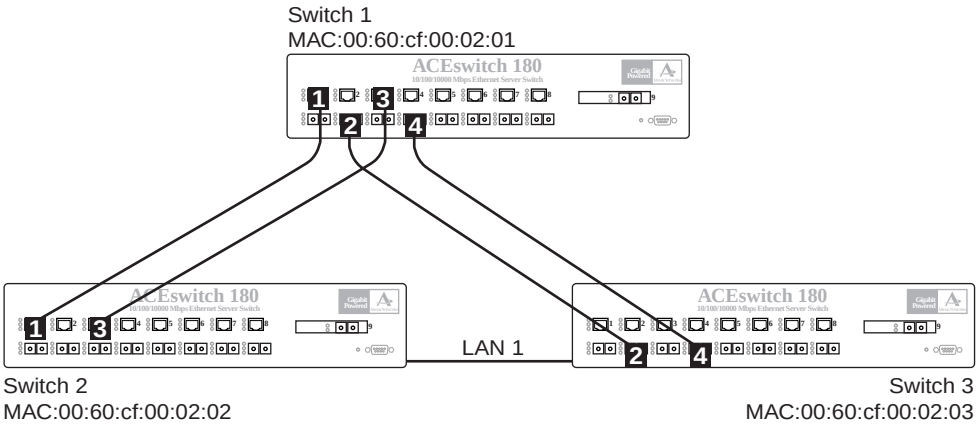

Figure 5 Spanning-Tree Topology

All switches have the default STP parameters except the following:

- Switch 1 MAC: 00:60:cf:00:02:01
- Switch 2 MAC: 00:60:cf:00:02:02, Path cost for port 1 (to Switch 1) is 10. Path cost for port 3 (to Switch 1) is 5.
- Switch 3 MAC: 00:60:cf:00:02:03, Path cost for port 2 and port 4 (to Switch 1) is 1.

Switch Receives its own Spanning-Tree BPDU Message

If the switch software receives its own bridge protocol data unit (BPDU) message, the switch port will be disabled. As an example, this could occur when the switch transmits the BPDU message out switch port 1 to a hub that has two hub ports connected together in a loop.

You must remove the loop from the port and manually re-enable the switch port. To manually re-enable the switch port, enter the following command:

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

Spanning-Tree Recalculation

The IEEE 802.1d Spanning-Tree algorithm can take up to 45 seconds from the time it detects a topology change to the time it transitions from “spanning-tree port” state to “forwarding” state. During Spanning-Tree recalculation, frame forwarding from the port will stop and interrupt normal network traffic flow. Unlike shared media environments, in a switched network environment when the end station directly connected to a switch port is rebooted, it causes the switch port link state to change, resulting in recalculation of the “spanning-tree port” state. This is seen by loss of connection upon end station reboot.

Server Load Balancing Configurations

General

The following checklist will help you resolve the most common difficulties configuring Server Load Balancing.

- Check the Server Load Balancing maintenance statistics (page 113) and the Server Load Balancing information (page 79) for anything unexpected.
- On the switch, check that the real servers, real server groups, virtual servers, etc. have been *enabled*.
- Check that the real servers are physically functioning.
- Check that all the services which the switch is expecting to find on each real server are installed, configured, and running properly.
- On the switch, make sure that you used `apply` and `save` to activate your configuration changes (see “Viewing, Applying, and Saving Changes” on page 149).
- On the switch, make sure that the real servers were added to the proper real server groups and that the real server groups are associated with a virtual server.
- Make sure that you are not violating any of the network topology restrictions, such as by connecting clients and servers to the same switch port.
- Make sure that the port state for each switch port is properly configured as `client`, `server`, or `none` (see “SLB Port Configuration” on page 210).
- Make sure that the switch is configured to accept the TCP/UDP port numbers on which each particular service is expected to run.

Service Problems

Periodic loss of a configured TCP service (such as HTTP). Real server does not come into service, or comes into service and fails periodically.

Possible Causes

- Invalid topology or port state: the real server is connected to the switch through a port configured in the “client” or “application redirection” state.
- There may be a health-check failure between the switch and the real server.
- One of the real servers of the real server group does not respond to the service request.

Actions

- Monitor the health checks. At Layer 4, there should be a 3-way TCP handshake for opening a TCP connection, followed by a 4-way TCP handshake to close a TCP connection.
- Verify that the real server has a default gateway or a route back to the client.
- Verify that the requested HTTP object is present on every real server in the real server group.

Index

Symbols

Numerics

A

B

C

D

E

F

H

I

L

M

N

O

P

Q

R

S

T

U

V

W

X