WebOS Switch Software

8.0 Command Reference

Part Number: 050081, Revision A, May 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

Chapter 3: First-Time Configuration 3-1

Chapter 4: Menu Basics 4-1

Chapter 5: The Information Menu 5-1

Preface

The *WebOS 8.0 Command Reference* describes how to configure and use the WebOS software with 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

The *WebOS 8.0 Command Reference* 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, "WebOS Feature Summary," provides brief descriptions 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, Jumbo Frames, Spanning-Tree Protocol, SNMP, Port Mirroring, IP Routing, Port Trunking, and more.

Chapter 8, "The SLB Configuration Menu," shows how to configure Server Load Balancing, Filtering, Global Server Load Balancing, and more.

Chapter 9, "The Operations Menu," shows how to use commands which affect switch performance immediately, but do not alter permanent switch configurations (such as temporarily disabling ports). Also describes how to activate or deactivate optional software features.

Chapter 10, "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 11, "The Maintenance Menu," shows how to generate and access a dump of critical switch state information, how to clear it, and how to clear part or all of the forwarding database.

Appendix A, "WebOS Syslog Messages," shows a listing of syslog messages for WebOS 8.0.

Appendix B, "WebOS SNMP Agent," lists the Management Interface Bases (MIBs) supported in WebOS 8.0.

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

Typographic Conventions

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

Table 1 Typographic Conventions

Typeface or Symbol	Meaning	Example
`AaBbCc123`	This type is used for names of commands, files, and directories used within the text.	View the `readme.txt` file.
	It also depicts on-screen computer output and prompts.	`Main#`
`AaBbCc123`	This bold type appears in command examples. It shows text that must be typed in exactly as shown.	`Main#` **`sys`**
AaBbCc123	This italicized type appears in command examples as a parameter placeholder. Replace the indicated text with the appropriate real name or value when using the command.	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

WebOS Feature Summary

This chapter provides a brief overview of features and enhancements provided in Alteon Web-System's WebOS switch software. The latest release of WebOS software extends the capabilities of the Alteon 180 and ACEdirector switches and supports features and enhancements implemented for the new Alteon 184 and AceDirector 4 switches.

WebOS extends switch application and content-intelligent traffic management services to include support for the following features:

- Intelligent traffic redirection with URL header parsing
- URL-based Server Load Balancing
- Bandwidth Management
- Session Persistence with advanced cookie parsing
- User-configurable client proximity tables to control Global SLB traffic redirection
- User-scriptable health checks and enhanced RADIUS health checking
- Advanced packet filtering
- Secure switch management
- Virtual Matrix Architecture, enabling the distribution of traffic across multiple processors and efficiently allocating memory resources to maximize switch performance

NOTE – Some features are optional and may require additional software licences on some switches. For information on activating these features on your switch (if necessary), see "Activating Optional Software" on page 9-9.

Finding the Information You Need

The following table lists the WebOS features described in this chapter, grouped according to their applicable Layer 1-7 functionality, and the page number where you'll find summary information for each feature.

NOTE – Once you have reviewed the information in this chapter, you'll find more application-specific information and configuration details for WebOS features and enhancements in the *WebOS 8.0 Application Guide*.

Layer	Feature/Function	See
Layer 5-7	**Content-Intelligent Load Balancing**	
	■ Host Header Inspection	page 1-4
	■ URL-based Server Load Balancing	page 1-5
	Content-Intelligent Web Cache Redirection	page 1-5
	■ Cachability based on domain name	
	■ Redirection based on domain name	
	Persistence-Based Load Balancing	page 1-6
	■ HTTP Cookie-Based Persistence	page 1-6
	■ Session ID-based Persistence	page 1-6
Layer 4	**Server Load Balancing**	page 1-7
	■ Health Checking	page 1-7
	■ Firewall Load Balancing	page 1-8
	■ FTP SLB	page 1-8
	■ Direct Server Return with one-arm load balancing	page 1-8
	■ Backup Server Group	
	Filtering	page 1-9
	■ Packet Filtering	
	Application Redirection	page 1-10
	■ FTP Client NAT	
	Global Server Load Balancing	page 1-12
	Bandwidth Management	page 1-13
Layer 3	**High Availability and VRRP**	page 1-14
	IP Routingl	page 1-16
	Border Gateway Protoco	page 1-16

Layer	Feature/Function	See
Layer2	**Jumbo Frames**	page 1-17
	Port Trunk Groups	page 1-17
	VLANs	page 1-18
	Spanning Tree Support	page 1-18
	Port Mirroring	page 1-18
	RMON Lite Support	page 1-18
Layer 1	**Port Link Characteristics**	page 1-19
Switch Management	**Secure Switch Administration** ■ SSH ■ SCP	page 1-19
	RADIUS Authentication	page 1-19
	Network Management ■ Command Line Interface ■ Browser-Based Interface	page 1-19 page 1-20 page 1-20
	SNMP MIB Support	page 1-21
Virtual Matrix Architecture	Distributed architecture, enabling switch resources to be shared across ports.	page 1-11
Dual Homing	Network resiliency and redundancy.	page 1-22

Layer 5-7 Features

Layer 5-7 features and support are summarized below, with pointers to where you can find topics discussed in greater detail.

NOTE – Not all features described in this section are "layer-specific."

Content-Intelligent Switching

"Content intelligent," or "content aware," switching enhancements in WebOS software are principally divided into two areas of functionality: HTTP Header Inspection and URL-based SLB. Each is briefly described below.

HTTP Host Header Inspection

Through inspection of HTTP HOST: headers, the switch supports these capabilities:

- Virtual Hosting

  Using a single virtual IP address to represent multiple domains in a hosting environment.

- Cache redirection based on domain names

  - Using the hashing algorithm, you can optimize "cache hits," redirecting client requests going to the same page of an origin server to a specific cache server.

  - Determine cachability or redirection, based on the domain name in the HOST header.

    For example, requests with URLs containing ".com" are redirected to Server Group 1 and all other requests are sent to Server Group 2.

- Defining cacheable domains

  Configure domain pages that should not be cached.

- HTTP Header-Based Server Load Balancing

  User-Agent header: direct requests to different servers, based on browser type.

NOTE – For more information about the application and configuration of content-intelligent switching, refer to the *WebOS 8.0 Application Guide*.

URL-Based Server Load Balancing

To take advantage of server memory and optimize server utilization, the switch can parse the URL information associated with incoming packets and redirect those packets based on URL sub-string matching. This enables the network administrator to send requests for static content to one server group while requests for dynamic content can be directed to another group. The switch can also inspect URLs that are split in multiple frames, up to 4500 bytes.

By inspecting the URL, the switch makes it easier to support these capabilities:

- Optimize "cache hits," redirecting client requests going to the same page of an origin server to a specific cache server.
- Exclusionary URL sub-string matching
- URL Hashing

URL-Based Web-Cache Redirection

WebOS software enables you to send requests with specific URLs or URL sub-strings to designated cache servers. The URL-based redirection option allows you to perform cache server farm tuning and offload overhead processing from the cache servers by only sending appropriate requests to the cache server farm. This feature supports both HTTP 1.0 and HTTP 1.1.

Each request is examined and handled as described below:

- If the request is a non-GET request like HEAD, POST, PUT, or HTTP with cookies, it doesn't get sent to the cache.
- If the request is an ASP or CGI request, or a dynamically generated page, it doesn't get sent to the cache.
- If the request is a cookie, it can optionally bypass the cache.

Session Persistence

Session persistence allows you to re-establish a user's connection to a particular server. This is an important consideration for administrators of e-commerce websites, where a server may have data associated with a specific user that is not dynamically shared with other servers at the site.

Persistence-based load balancing enables the network administrator to redirect requests from a client to the real server that initially handled the request. On a switch running WebOS software, persistence can be based on IP source address, cookies for HTTP requests, or SSL session ID for encrypted HTTP requests.

Cookie-Based Persistence

As viewed by the switch, a "cookie" is an HTTP header sent as part of the response to a request from a client by the switch or the server. On all subsequent requests from the client for the same IP address, the client includes the cookie, to enable the server to determine that the user is the same one that sent the original request.

There are two methods of handling cookies; each is described below:

- Passive Cookie Mode

  In this mode, there is no definition of any special persistence cookie on the server. The network administrator configures an existing cookie that the switch should look for in subsequent requests from the same client.

- Active Cookie Mode/Cookie Rewrite Mode

  In active cookie mode (or cookie rewrite mode), the switch will generate the cookie value on behalf of the server, eliminating the need for a network administrator to generate cookies for each user.

SSL Session ID-Based Persistence

Using the SSL session ID, the switch forwards the request to the real server that it bound the user to during the last session. Using the session ID will enhance network performance. Negotiating the type of encryption used during data transmission requires a lengthy protocol exchange that is bypassed with session persistence.

NOTE – For more information about the application and configuration of session persistence policies, refer to the *WebOS 8.0 Application Guide*.

Layer 4 Features

Running WebOS software, Alteon WebSystem switches support 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.

WebOS Layer 4 (Application/Protocol) features and support are summarized below, with pointers to where you can find topics discussed in greater detail.

Server Load Balancing

Through server load balancing, your Alteon Web switch is aware of the shared services provided by your server pool. The switch can then balance user session traffic among the available servers. For even greater control, traffic is distributed according to a variety of user-selectable metrics.

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

- TCP and UDP load balancing
- SSL session ID substitution
- MaxConns, back-up, and overflow server support
- Round-robin and connection-based load balancing
- Server static weighting
- Hash and min-misses load balancing

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.

WebOS server health check enhancements increase availability to a broader base of applications by allowing flexible application and content verification using customized scripts. Through the use of “send/expect” script-based health checks, you can configure the switch to send a sequence of user-specified tests to be executed on select servers and the corresponding confirmation to be verified, to ensure total application and content availability. The switch can also check availability of SSL and RADIUS functionality.

Firewall Load Balancing

Firewall load balancing with Alteon WebSystems Web switches allows firewalls to operate in parallel, giving network administrators the ability to maximize firewall productivity, scale firewall performance without forklift upgrades, and eliminate the firewall as a single point of failure.

To facilitate FWLB, Alteon WebSystems Web switches feature a powerful distributed processing architecture, and sophisticated Layer 4 through 7 switching functionality, including the ability to maintain the state of individual TCP sessions. This makes our Web switches ideally suited for the processor-intensive packet examination and manipulation required to perform FWLB for thousands or tens of thousands of packets per second.

FTP Server Load Balancing

FTP load balancing enhancements enable the switch to look deep into the FTP packet to support load balancing of FTP servers on a private network, regardless whether the servers are operating in active or passive FTP mode.

Direct Server Return

Using this WebOS feature, the server responds directly to the client in a one-arm load balancing configuration. This capability is useful for sites where large amounts of data are going from servers to clients, such as is the case with content providers or portal sites that have asymmetric traffic patterns.

Filtering

Alteon WebSystems switches support Layer 3 (IP) and Layer 4 (transport) filtering for up to 224 filters per port, giving network administrators a powerful tool to protect their server networks. Switch-wide filtering rules can be defined on each Alteon WebSystems switch, with any or all rules applied to each port.

WebOS software supports packet filtering based on IP options, ICMP message types, the IP Type-of-Service (TOS), and TCP flags. 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)

  WebOS software supports packet filtering, based on any or all ICMP message types.
- TCP flags

  WebOS software supports packet filtering, based on any or all TCP flags.
- Application source port, by name, integer or range
- Application destination port, by name, integer or range
- TOS coloring; that is, replacing TOS value with configured value per filter

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, router load balancing, and web cache redirection.

Web Cache Redirection

Web-cache redirection can help alleviate the congestion seen at your Internet router. When Application Redirection filters are properly configured for your WebOS powered switch, outbound client requests for Internet data are intercepted and redirected to a group of web-cache servers on your network. The web-cache servers duplicate and store inbound Internet data that has been requested by your clients. If the web-cache servers recognize a client’s outbound request as one that can be filled with cached information, the web-cache servers will supply the information, rather than sending the request out across the Internet.

FTP Client NAT

Alteon WebSystems switches provide Network Address Translation (NAT) services to many clients with private IP addresses. However, on switches running WebOS 6.0, clients using active FTP cannot send a request to a remote FTP server when their client IP address is private. In WebOS 8.0, an FTP enhancement now provides the capability to perform true FTP NAT for dynamic NAT.

The switch can monitor the control channel and replace the client 's private IP address with a proxy IP (PIP) address defined on the switch. When a client in active FTP mode sends a “PORT” command to a remote FTP server, the switch will look into the data part of the frame and modify the PORT command.

Virtual Matrix Architecture

Virtual Matrix Architecture (VMA) is a hybrid architecture that realizes the full potential of distributed processing by taking advantage of any unused resources within a Web switch. It combines the strengths of central and distributed processing to deliver improvements in processing power and port capacity.

With VMA, the switch makes optimal use of system resources by distributing the workload to multiple processors. Dividing the workload and using multiple processors to complete a task improves switch performance and increases the number of concurrent sessions per switch.

Characteristics of VMA are

- Based on its IP address, each client is assigned a designated port that does all the Layer 4 frame processing.
- Each client is assigned to a designated port's CPUs for L4-7 processing, regardless of where it ingresses. The algorithm ensures even distribution of traffic. Packets to and from the same client are always processed by the same CPUs.
- Session entries are kept in memory local to designated CPUs. A global session table is kept for all persistent sessions. All ports store all filtering/redirection policies.
- CPUs at all ports are actively processing load at all times.
- Memory at all 8 ports is pooled to increase storage capacity, enabling up to 512K session table entries, depending on platform and configuration; even when all traffic enters at a single port.
- Increased packet buffering capacity.

VMA provides maximum parallel processing with minimum memory search latency. It is optimized for both asymmetric and symmetric topologies, providing up to an eight-fold increase in session performance at all load levels.

Global SLB

Using Global Server Load Balancing (GSLB), you can balance server traffic load across multiple physical sites. This allows you to smoothly integrate the resources of a world-wide series of server sites and balance Web content (or other services) intelligently among them. Alteon WebSystems' GSLB takes into account individual sites' health, response time, and geographic location for a global performance perspective. WebOS supports up to 64 GSLB sites per switch.

URL-based SLB is compatible with GSLB. Cookie-based persistence is compatible with GSLB using Active Cookie Mode (Cookie Rewrite Mode).

In certain customer configurations, IANA data does not provide sufficient geographic separation of proximity information. As a result, large ISP partners cannot use their own geographic data to determine GSLB site selection based on client location. WebOS software supports client proximity tables using static "client to site" mapping. Switch managers can configure private client proximity information. The limit on the number of entries in the proximity database is 128.

NOTE – For more information about the application and configuration of Global SLB, refer to the *WebOS 8.0 Application Guide*.

Bandwidth Management

Bandwidth management enables website managers to allocate a certain portion of the available bandwidth for specific users or applications. Traffic classification can be based on user or application information. Policies can be configured to set lower and upper bounds on the bandwidth allocation.

Bandwidth management provides the following support:

- Allocation of capacity from a Web server farm by website address.
- Allocation of WAN capacity by website address.
- User-defined minimum and maximum rates.
- Data pacing technology for rate-controlling the source.
- Traffic classification based on Layer 2/3/4 attributes.
  - Nesting of rules allowed.
  - 256 to 1,024 traffic classes per switch. (More than 256 classes supported only on Alteon A184 and AD4 Web switches).
- Bandwidth management based on traffic classes.
  - Multiple levels of bandwidth per traffic class. Guaranteed bandwidth, soft ceiling, hard ceiling.
  - Classification based on user information or application.
  - Guaranteed bandwidth for mission-critical applications.
- Accounting. Usage logging per traffic class
- IP ToS bit rewrite, based on whether the traffic is over or under the soft limit.

NOTE – For more information about the application and configuration of bandwidth management contracts, refer to the *WebOS 8.0 Application Guide*.

Layer 3 Features

WebOS Layer 3 (IP) features and support are summarized below, with pointers to where you can find topics discussed in greater detail.

High Availability and VRRP

High availability refers to a network topology where no device can create a single point of failure on a network and no single device forces a point of failure on another part of the network.

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, as defined in RFC 2338, to include virtual servers, allowing for full *active-active*, *active-standby, and hot-standby* redundancy between its Layer 4 switches. Both redundant switch configurations increase application availability by removing single points of failure from networks.

Active-Standby Redundancy

In an *active-standby configuration*, shown in Figure 1-1 on page 1-15, both switches can support active traffic. However, services are not shared across the switches. That is, each switch can be active for some number of services, such as IP routing interfaces or load balancing VIP addresses, and act as a standby for other services on the other switch.

NOTE – In an active-standby configuration, the same service cannot be active simultaneously on both switches.

Active-Active Redundancy

Active-active redundancy enables more efficient network resource allocation than the hot-standby method. It also supports more complex failover topologies. In an *active-active* configuration, shown in Figure 1-2 on page 1-15, both switches can process traffic for the same service at the same time. That is, both switches can be active simultaneously for a given IP routing interface or load balancing virtual server (VIP).

When both switches are healthy, active-active configurations increase performance and capacity by allowing two or more Web switches to support the same interface and service.

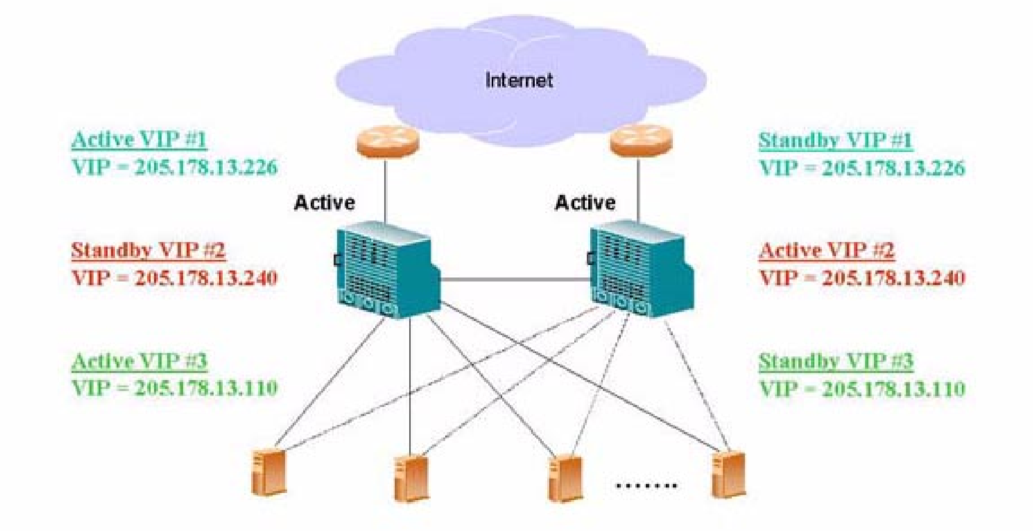

Figure 1-1 Active-Standby Redundancy

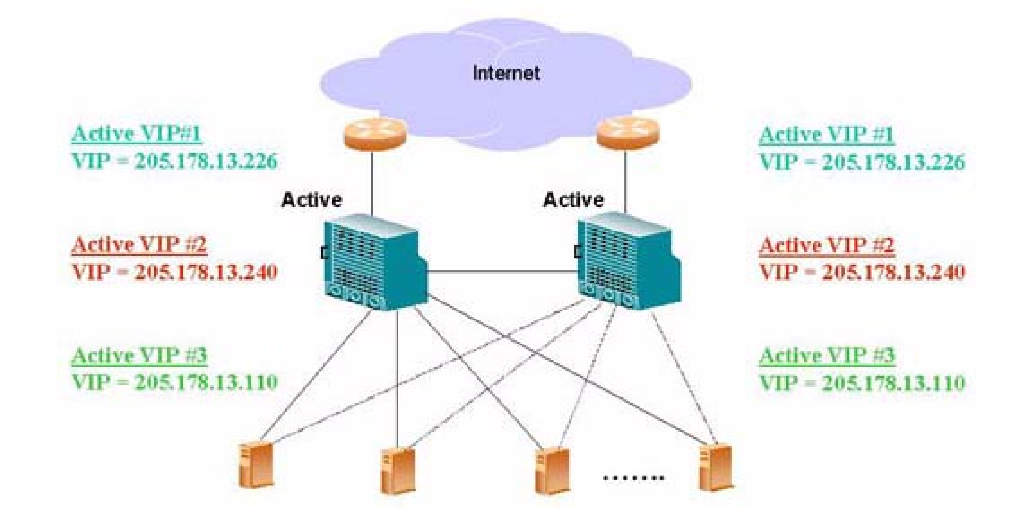

Figure 1-2 Active-Active Redundancy

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.

Border Gateway Protocol (BGP) Support

WebOS allows user Domain Name Server (DNS) requests to be resolved by the closest authoritative DNS server, based on Border Gateway Protocol (BGP) autonomous system (AS) hops. User requests can now be served by the closest site using BGP, while BGP route removal ensures that requests will not be forwarded to failed or overloaded server farms.

Alteon Web switches can advertise their IP interfaces and VIP addresses using BGP, as well as take BGP feeds from up to four BGP router peers. This gives customers more resilience and flexibility in balancing traffic from the Internet.

Layer 2 Features

WebOS Layer 2 features and support are summarized below, with pointers to where you can find topics discussed in greater detail.

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

Jumbo Frames

Alteon Websystem switches automatically and transparently forward Ethernet frames of all sizes, including optional jumbo frames of up to 9,000 bytes. Jumbo frames can reduce packet processing overhead on servers by as much as 85 percent and increase throughput on CPU-bound systems by over 100 percent.

NOTE – Jumbo frames are not supported on ports operating at 10 Mbps or ports set to half-duplex at any link speed.

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

Port Trunk Groups

Ports in a trunk group combine their bandwidth to create a single, larger virtual link. WebOS software supports EtherChannel-compatible trunk groups, enabling link-level redundancy and load sharing with other EtherChannel-compatible devices. WebOS support enables the following port trunking capabilities:

- Up to eight trunk groups can be configured per switch
- Up to six ports can be trunked together to form a single virtual link with bandwidth between 2 and 4 Gigabits per second
- IP Session ID hashing for IP (addresses?)
- MAC SA/DA hashing for non-IP (addresses?)
- 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.

WebOS supports 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 VLANs can be configured per port, enabling the network administrator to create separate VLANs for different packet types, such as IP, IPX.
- Up to 246 VLANs per switch

Spanning Tree Support

When multiple paths exist, Spanning Tree configures the network so that a switch uses only the most efficient path. Each Spanning Tree on the switch is associated with a port and can be configured with multiple VLANs.

Port Mirroring

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

NOTE – Port mirroring is supported on the AD4 and A184 Web switches.

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 1 Features

WebOS Layer 1 features and support are summarized below, with pointers to where you can find topics discussed in greater detail.

Port Link Characteristics

- 100 Mbps ports support half- and full-duplex operations and 802.3u 10/100 autonegotiation
- 1000 Mbps ports support 802.3z full-duplex operation with asymmetric flow control
- IEEE 802.3x Flow Control

Switch Management Features

Network administrators can configure and monitor all Alteon WebSystems switch functions via the WebOS Browser-Based Interface (BBI), SNMP applications, and a command-line interface (CLI) accessed from the console port, via Telnet, or via Secure Shell program (SSH). Seven levels of password protection are provided, to allow switch configuration changes and to view switch information and statistics.

NOTE – WebOS access levels and password protection are described in "Accessing the Switch" on page 2-5.

Secure Switch Administration

- Secure shell (SSH) protocol-based secured switch management on the Alteon AD4 and Alteon 184 Web switches.
- Secure copy protocol (SCP) can be used to securely upload/download switch configuration.

RADIUS Authentication

- Authentication for secured switch management supports a variable-length RADIUS secret password.
- RSA SecurID token-based authentication is supported, provided that the RADIUS server can do the RSA ACE/Server proxy.

NOTE – For information on how SSH, SCP, and RADIUS authentication is implemented, refer to the *WebOS 8.0 Application Guide*.

Network Management

Usability features of the command-line interface and SNMP are described below.

Command-Line Interface (CLI)

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.

WebOS CLI features are listed below:

- Configuration restore command

  "revert" command to remove pending changes between "apply"
- Viewing of last 10 syslog messages from console
- Option to reset (zeroing) Layer 4 and Layer 7 statistics counters via a single CLI command. The **`clear`** command is found in the `/stat/slb` and the `/stat/slb/port` menu.

  The following counters are NOT cleared by this command:

  - All operational counters that the switch used to perform Layer 4 and Layer 7 functions; for example. the current sessions on a real server.
  - All related SNMP counters.
- New configuration dump format, `/cfg/dump`

  More "readable" configuration dump, with one command per line and indentation.

Browser-Based Interface (BBI)

The WebOS BBI provides direct browser-to-switch interaction for switch configuration and monitoring.

Alteon AD4 and A184 Web switches support a private MIB and four groups of RMON on every port. Port mirroring provides for switch and server performance analysis. The switch management interface is integrated with HP OpenView 6.0 under UNIX (HPUX, Solaris) and Windows NT.

SNMP MIB Support

The SNMP agent for Alteon WebSystems Web 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 WebOS MIB, which can be compiled into MIB-based systems such as HP-OpenView.

SNMP agent features in WebOS are listed below:

- Option to enable/disable SNMP
- Option to disable SNMP SET

NOTE – For a listing of SNMP Agent MIBs supported in WebOS 8.0, refer to Appendix B.

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.

WebOS 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 your ACEnic adapter *Installation and User's Guide* for more information about this feature.

CHAPTER 2

The Command-Line Interface

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

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

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

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

Connecting to the Switch

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

- Using a console connection via the console port
- Using a Telnet connection over the network
- Using a SSH connection to securely log into another computer over a network

Establishing a Console Connection

Requirements

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

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

Table 2-1 Console Configuration Parameters

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

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

Procedure

1. **Connect the terminal to the Console port using the serial cable.**

2. **Power on the terminal.**

3. **To establish the connection, press <Enter> a few times on your terminal.**

You will next be required to enter a password for access to the switch. (For more information, see “Setting Passwords” on page 3-13).

Establishing a Telnet Connection

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

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

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

Using a BOOTP Server

By default, the WebOS software is set up to request its IP address from a `BOOTP` server. If you have a `BOOTP` server on your network, add the MAC address of the switch to the `BOOTP` configuration file located on the `BOOTP` server. The MAC address can be found on a small white label on the back panel of the switch. The MAC address can also be found in the System Information Menu (see "System Information" on page 5-15).

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 then be prompted to enter a password as explained below.

Establishing an SSH Connection

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

NOTE – SSH is only supported on Alteon AD4 and Alteon 184 Web switches.

The supported SSH encryption and authentication methods are listed below.

Server Host Authentication:	Client RSA-authenticates the switch in the beginning of every connection.
Key Exchange:	RSA
Encryption:	3DES-CBC, DES
User Authentication:	Local password authentication, Radius

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

The following SSH clients have been tested:

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

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

Running SSH

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

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

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

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

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

You will then be prompted to enter your username and password.

Accessing the Switch

To enable better switch management and user accountability, seven levels or *classes* of user access have been implemented on the 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, Telnet, or SSH, 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 3-13.

Table 2-2 User Access Levels

User Account	Description and Tasks Performed	Password
User	The User has no direct responsibility for switch management. He or she can view all switch status information and statistics, but cannot make any configuration changes to the switch.	`user`
SLB Operator	The SLB Operator manages Web servers and other Internet services and their loads. In addition to being able to view all switch information and statistics, the SLB Operator can enable/disable servers using the Server Load Balancing operation menu.	`slboper`
Layer 4 Operator	The Layer 4 Operator manages traffic on the lines leading to the shared Internet services. This user currently has the same access level as the SLB operator. and the access level is reserved for future use, to provide access to operational commands for operators managing traffic on the line leading to the shared Internet services.	`l4oper`
Operator	The Operator manages all functions of the switch. In addition to SLB Operator functions, the Operator can reset ports or the entire switch.	`oper`
SLB Administrator	The SLB Administrator configures and manages Web servers and other Internet services and their loads. In addition to SLB Operator functions, the SLB Administrator can configure parameters on the Server Load Balancing menus, with the exception of not being able to configure filters or bandwidth management.	`slbadmin`

Table 2-2 User Access Levels

User Account	Description and Tasks Performed	Password
Layer 4 Administrator	The Layer 4 Administrator configures and manages traffic on the lines leading to the shared Internet services. In addition to SLB Administrator functions, the Layer 4 Administrator can configure all parameters on the Server Load Balancing menus, including filters and bandwidth management.	`l4admin`
Administrator	The superuser Administrator has complete access to all menus, information, and configuration commands on the switch, including the ability to change both the user and administrator passwords.	`admin`

NOTE – With the exception of the "admin" user, access to each user level can be disabled by setting the password to an empty value. All user levels below "admin" will (by default) be initially disabled (empty password) until they are enabled by the "admin" user. This is done in order to avoid inadvertently leaving the switch open to unauthorized users.

CLI 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 2-1 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 7-6.

CHAPTER 3

First-Time Configuration

To help with the initial process of configuring your switch, the WebOS software includes a Setup utility. The Setup utility prompts you step-by-step to enter all the necessary information for basic configuration of the switch. This chapter describes how to use the Setup utility and how to change system passwords.

Using the Setup Utility

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

Information Needed For Setup

Setup requests the following information:

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

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

Starting Setup When You Log In

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

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

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

   ```
   Enter Password:
   ```

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

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

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

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

   NOTE – If the default `admin` login is unsuccessful, or if the administrator Main Menu appears instead, the system configuration has probably been changed from the factory default settings. If you are certain that you need to return the switch to its factory default settings, see "Selecting a Configuration Block" on page 10-4.

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

Stopping and Restarting Setup Manually

Stopping Setup

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

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

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

Restarting Setup

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

```
# /cfg/setup
```

Setup Part 1: Basic System Configuration

When Setup is started, the system prompts:

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

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

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

   If you decide not to configure VLANs during this session, you can configure them later using the configuration menus, or by restarting the Setup facility. For more information on VLANs issues, see *(*to be supplied)*.

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

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

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

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

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

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

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

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

```
Enter year [99]:
```

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

The system displays the date and time settings:

```
System clock set to 13:56:52 Wed June 17, 1999.
```

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

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

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

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

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

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

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

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

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

The system displays the date and time settings:

```
System clock set to 13:56:52 Wed June 17, 1999.
```

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

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

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

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

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

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

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

Setup Part 2: Port Configuration

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Setup Part 3: VLANs

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

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

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

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

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

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

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

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

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

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

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

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

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

```
>
```

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

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

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

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

Setup Part 4: IP Configuration

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

IP Interfaces

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

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

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

```
IP Config:

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

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

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

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

To keep the current setting, press <Enter>.

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

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

To keep the current setting, press <Enter>.

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

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

To keep the current setting, press <Enter>.

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

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

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

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

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

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

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

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

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

Default Gateways

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

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

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

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

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

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

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

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

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

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

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

IP Routing

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

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

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

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

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

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

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

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

Setup Part 5: Final Steps

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Setting Passwords

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

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

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

Changing the Default Administrator Password

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

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

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

```
Main# cfg
```

The Configuration Menu is displayed

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

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

```
>> Configuration# sys
```

The System Menu is displayed

```
[System Menu]
      radius   - RADIUS Authentication Menu
      ntp      - NTP Server Menu
      date     - Set system date
      time     - Set system time
      usrpw    - Set user password
      admpw    - Set administrator password
      l4apw    - Set L4 administrator password
      slbpw    - Set Slb administrator password
      l4upw    - Set L4 user password
      idle     - Set timeout for idle CLI sessions
      snmp     - Set SNMP access control
      wport    - Set Web server port number
      bannr    - Set login banner
      mnet     - Set management network
      mmask    - Set management netmask
      smtp     - Set SMTP host
      bootp    - Enable/disable use of BOOTP
      http     - Enable/disable HTTP (Web) access
      cur      - Display current system-wide parameters
```

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.

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
```

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

```
System# l4apw
```

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

```
Changing L4 ADMINISTRATOR password; validation required...
Enter current administrator password:
```

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

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

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

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

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

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

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

CHAPTER 4

Menu Basics

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

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

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

The Main Menu

The Main Menu appears after a successful connection and login. Figure 4-1 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 4-1 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 memory.

- **Operations Command Menu**

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

- **Boot Options Menu**

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

- **Maintenance Menu**

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

Global Commands

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

Table 4-1 Global Commands

Command	Action
`?` *command*	Provides more information about a specific command on the current menu. When used without the *command* parameter, a summary of the global commands is displayed.
`.`	Display the current menu.
`..`	Go up one level in the menu structure.
`/`	If placed at the beginning of a command, go to the Main Menu. Otherwise, this is used to separate multiple commands placed on the same line.
`diff`	Show any pending configuration changes.
`apply`	Apply pending configuration changes.
`save`	Write configuration changes to non-volatile flash memory.
`revert`	Remove pending configuration changes between "apply" commands. Use this command to restore configuration parameters set since last "apply" command.
`exit`	Exit from the command-line interface and log out.
`ping`	Use this command to verify station-to-station connectivity across the network. The format is as follows: **`ping`** *address* [*tries* [*delay*]] Where *address* is the hostname or IP address of the device, *tries* (optional) is the number of attempts (1-32), and *delay* (optional) is the number of milliseconds between attempts. The DNS parameters must be configured if specifying hostnames (see "Domain Name System Configuration" on page 7-30).

Table 4-1 Global Commands

Command	Action
`traceroute`	Use this command to identify the route used for station-to-station connectivity across the network. The format is as follows: **`traceroute`** *address* [*max-hops* [*delay*]] Where *address* is the hostname or IP address of the target station, *max-hops* (optional) is the maximum distance to trace (1-16 devices), and *delay* (optional) is the number of milliseconds for wait for the response. As with `ping`, the DNS parameters must be configured if specifying hostnames.
`pwd`	Display the command path used to reach the current menu.
`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 4-2 Command-Line History and Editing Options

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

Command-Line Interface Shortcuts

Command Stacking

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

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

Command Abbreviation

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

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

Tab Completion

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

CHAPTER 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]
      slb     - Layer 4 Information Menu
      route   - IP Routing Information Menu
      arp     - ARP Information Menu
      fdb     - Forwarding Database Information Menu
      sys     - Show system information
      log     - Show last 10 syslog messages
      link    - Show link status
      stp     - Show STP information
      vlan    - Show VLAN information
      port    - Show port information
      ip      - Show IP information
      vrrp    - Show Virtual Router Redundancy Protocol information
      trunk   - Show Trunk Group information
      bwm     - Show Bandwidth Management information
      swkey   - Show enabled software features
      dump    - Dump all information
```

The information provided by each menu option is briefly described in Table 5-1 on page 5-2, with pointers to where detailed information can be found.

Table 5-1 Information Menu Options (/info)

Command Syntax and Usage
`slb` Displays the Layer 4 Information Menu. For details, see page 5-5.
`route` Displays the IP Routing Menu. Using the options of this menu, the system displays the following for each configured or learned route: ■ Route destination IP address, subnet mask, and gateway address ■ Type of route ■ Tag indicating origin of route ■ Metric for RIP tagged routes, specifying the number of hops to the destination (1-15 hops, or 16 for infinite hops) ■ The IP interface that the route uses For details, see page 5-8
`arp` Displays the Address Resolution Protocol (ARP) Information Menu. For details, see page 5-11.
`fdb` Displays the Forwarding Database Information Menu. For details, see page 5-13.
`sys` Displays system information, including: ■ System date and time ■ Switch model name and number ■ Switch name and location ■ Time of last boot ■ MAC address of the switch management processor ■ IP address of IP interface #1 ■ Hardware version and part number ■ Software image file and version number ■ Configuration name ■ Log-in banner, if one is configured For details, see page 5-15.
`log` Displays 10 most recent syslog messages. For details, see page 5-16.

Table 5-1 Information Menu Options (/info)

Command Syntax and Usage
`link` Displays configuration information about each port, including: ■ Port number ■ Port speed (10, 100, 10/100, or 1000) ■ Duplex mode (half, full, or auto) ■ Flow control for transmit and receive (no, yes, or auto) ■ Link status (up or down) For details, see page 5-17.
`stp` In addition to seeing if STP is enabled or disabled, you can view the following STP bridge information: ■ Priority ■ Hello interval ■ Maximum age value ■ Forwarding delay ■ Aging time You can also see the following port-specific STP information: ■ Port number and priority ■ Cost ■ State For details, see page 5-18.
`vlan` Displays VLAN configuration information, including: ■ VLAN Number ■ VLAN Name ■ Status ■ Jumbo Frame usage ■ Port membership of the VLAN For details, see page 5-20.

Table 5-1 Information Menu Options (/info)

Command Syntax and Usage
`port` Displays port status information, including: ■ Port number ■ Whether the port uses VLAN Tagging or not ■ Port VLAN ID (PVID) ■ Port name ■ VLAN membership For details, see page 5-21.
`ip` Displays IP Information. For details, see page 5-22. IP information, includes: ■ IP interface information: Interface number, IP address, subnet mask, broadcast address, VLAN number, and operational status. ■ Default gateway information: Metric for selecting which configured gateway to use, gateway number, IP address, and health status ■ IP forwarding information: Enable status, `lnet` and `lmask` ■ Port status
`vrrp` Displays the VRRP Information Menu. For details, see page 5-23.
`trunk` When trunk groups are configured, you can view the state of each port in the various trunk groups. For details, see page 5-24.
`bwm` Shows bandwidth management information. For details, see page 5-25.
`swkey` Displays a list of all the optional software packages which have been activated or installed on your switch.
`dump` Dumps all switch information available from the Information Menu (10K or more, depending on your configuration). If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump commands.

/info/slb
SLB Information

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

Layer 4 information includes the following:

Table 5-2 Layer 4 Information Menu Options (/info/slb)

Command Syntax and Usage

`sess`

Displays the Session Table Information Menu. To view menu options, refer to page 5-6.

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

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

`virt` *<virtual server number (1-256)>*

- Displays Virtual Server State: Virtual server number, IP address, virtual MAC address
- Virtual Port State: Virtual service or port, server port mapping, real server group, group backup server.

`filt` *<filter ID (1-224)>*

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

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

Displays the physical port number, proxy IP address, filter status, a list of applied filters, and client and/or server Layer 4 activity.

`gslb`

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

`dump`

Displays all Layer 4 information for the switch. For details, see page 5-7.

/info/slb/sess

Show Session Table Information

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

Table 5-3 Session Information Menu Options (/info/slb/sess)

Command Syntax and Usage

`find` *<IP address>*

Displays all session entries with source IP address.

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

Displays all session entries on port.

`dump`

Displays all session entries.

/info/slb/dump

Show All Layer 4 Information

```
Global SLB state:
  1: 220.3.78.3,  0.0.0.0,         FAILED

Real server state:
  2: 10.10.10.2, 00:60:cf:42:e4:40, vlan 1, port 8, health 4, up

Virtual server state:
  1: 10.10.10.10,     00:60:cf:40:78:ce
    HTTP Application:     virtual ports:
    http: rport http, group 1, backup none, httpslb
        real servers:
          2: 10.10.10.2, backup none, up
             exclusionary string matching: disabled

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

Port state:
  1: 0.0.0.0
     filt disabled, filters: empty
  2: 0.0.0.0
     filt disabled, filters: empty
  3: 0.0.0.0
     filt disabled, filters: empty
  4: 0.0.0.0
     filt disabled, filters: empty
```

/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 5-4 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 (eg., 192.4.17.44)>* Displays routes to a single gateway.
`type indirect|direct|local|broadcast|martian|multicast` Displays routes of a single type.

`tag fixed|static|snmp|addr|rip|icmp|broadcast`
`|martian|multicast|dynamic`
Displays routes of a single tag.

`if` *<interface number (1-256)>* Displays routes on a single interface.
`dump` Displays all routes configured in the switch. For more information, see page 5-9.

/info/route/dump

Show All IP Route Information

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

The following table describes the `Type` parameters.

Table 5-5 IP Routing Type Parameters

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

The following table describes the `Tag` parameters.

Table 5-6 IP Routing Tag Parameters

Parameter	Description
`fixed`	The address belongs to a host or subnet attached to the switch.
`static`	The address is a static route which has been configured on the 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/arp
ARP Information

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

The ARP information includes IP address and MAC address of each entry, address status flags (see Table 5-8 on page 12), VLAN and port for the address, and port referencing information.

Table 5-7 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 number (1-16)>* Displays the ARP entries on a single port.
`vlan` *<VLAN number (1-4094)>* Displays the ARP entries on a single VLAN.
`refpt` *<port number (1-16)>* Displays the ARP entries referenced by a single port.
`dump` Displays all ARP entries. including: ■ IP address and MAC address of each entry ■ Address status flag (see below) ■ The VLAN and port to which the address belongs ■ The ports which have referenced the address (empty if no port has routed traffic to the IP address shown) For more information, see page 12.
`addr` Displays the ARP address list.

/info/arp/dump

Show All ARP Entry Information

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

The `Flag` field is interpreted as follows:

Table 5-8 ARP Dump Flag Parameters

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

/info/arp/addr

ARP Address List Information

```
>> Address Resolution Protocol# addr
    IP address        IP mask        MAC address    VLAN Flags
  --------------- --------------- ----------------- ---- -----
  205.178.18.66   255.255.255.255 00:70:cf:03:20:04        P
  205.178.50.1    255.255.255.255 00:70:cf:03:20:06    1
  205.178.18.64   255.255.255.255 00:70:cf:03:20:05    1
```

/info/fdb
FDB Information.

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

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

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

Table 5-9 FDB Information Menu Options (/info/fdb)

Command Syntax and Usage
`find` *<MAC-addr>* [*<VLAN>*] Displays a single database entry by its MAC address. You are prompted to enter the MAC address of the device. Enter the MAC address using the format, `xx:xx:xx:xx:xx:xx`. For example, `08:00:20:12:34:56`. You can also enter the MAC address using the format, `xxxxxxxxxxxx`. For example, `080020123456`.
`port` *<port as number 1-16>* Displays all FDB entries for a particular port.
`vlan` *<VLAN number 1-4094>* Displays all FDB entries on a single VLAN.
`refpt` *<port as number 1-16>* Displays the FDB entries referenced by a single port.
`dump` Displays all entries in the Forwarding Database. For more information, see page 5-14.

/info/fdb/dump

Show All FDB Information

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

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

If the state for the port is listed as an interface (`IF`), the MAC address is for a standard VRRP virtual router. If the state is listed as a virtual server (`VIP`), the MAC address is for a virtual server router; 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, refer to page 11-6.

/info/sys
System Information

```
System Information at 11:40:46 Thu Apr 13, 2000

ACEdirector 4
sysName:
sysLocation:
Last boot: 14:52:59 Tue Apr 11, 2000 (reset from console)

MAC address: 00:60:cf:43:a4:70    IP (If 1) address: 172.19.1.201
Hardware Revision: B
Hardware Part No: 200009C04
Software Version 8.0.15 (FLASH image1), active configuration.
SLB Switch 01, Southwest Territory
```

System information includes:

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

/info/log
Show Last 10 Syslog Messages

```
>> Information# log
Apr 1 17:28:52 ALERT slb: cannot contact real server 215.118.113.74
Apr 1 17:29:10 NOTICE console: admin login
Apr 1 17:30:01 NOTICE  telnet/ssh-1: admin idle timeout from Telnet
Apr 1 18:55:43 NOTICE  telnet/ssh-1: admin logout from Telnet
Apr 2 12:56:35 INFO web server: new configuration applied
Apr 2 14:57:35 WARNING slb: filter 10 fired on port 4
Apr 3 7:58:03 ERR telnet: no apply needed
```

Each syslog message has a criticality level associated with it, included in text form as a prefix to the log message. One of eight different prefixes is used, depending on the condition that the administrator is being notified of, as shown below.

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

/info/link
Link Status Information

```
------------------------------------------------------------------
Port   Speed    Duplex     Flow Ctrl      Link
----   -----   --------  --TX-----RX--   ------
  1    10/100    any      yes    yes      down
  2      100    full      yes    yes      down
  3    10/100    any      yes    yes      down
  4      100    half       no     no       up
  5      100    half       no     no      down
  6      100    half       no     no      down
  7    10/100    any      yes    yes      down
  8      100    half       no     no       up
  9     1000    full      yes    yes      down
```

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

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

For example, if "**`/info/link`**" is entered from console, link status of all four Gigabit ports on slot 1 will be displayed. If no parameter is entered, link status of entire switch wire ports will be displayed. The screenshot shown above is an abbreviated link status of the entire switch wire ports. An error message will be displayed if the slot number entered is either a Management Processor or Switch Processor slot.

/info/stp
Spanning Tree Information

```
Current Root:            Path-Cost Port Hello MaxAge FwdDel Aging
 7fff 00:60:cf:40:4c:b0       15      8    2     20     15    300

Number of topology changes:        2
Time since last topology change:    0 days, 03:24:08

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

Port  Priority   Cost    State      Designated Bridge       Des Port
  1      128       0    DISABLED
  2      128       0    DISABLED
  3      128       0    DISABLED
  4      128      10   FORWARDING  8000-00:60:cf:43:a4:70   32772
  5      128       0    DISABLED
  6      128       0    DISABLED
  7      128       0    DISABLED
  8      128      10   FORWARDING  8000-00:60:cf:40:61:00   32776
  9      128       0    DISABLED
```

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

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

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

- Slot number
- Port number and priority
- Cost
- State

The following table describes the STP parameters.

Table 5-10 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/vlan
VLAN Information

```
VLAN                Name                 Status  Jumbo   Ports
---- --------------------------------    ------  -----   --------------
1      Default VLAN                        ena     n     2/1 2/2 3/4-3/9
2      VLAN 2                              ena     n     4/3
```

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

VLAN information includes:

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

/info/port

Port Information

```
------------------------------------------------------------------
Port  Tag   UnTag  PriTag  RMON             VLAN(s)
----  ----  -----  ------  ----  ------------------------------
  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:

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

/info/ip
IP Information

```
Interface information:
  1: 172.19.1.201,    255.255.255.0,   172.19.1.255,    vlan 1, up

Default gateway information: metric strict
  1: 172.19.1.1,      up

Current IP forwarding settings: OFF

Current local networks:

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

BGP Information:
  OFF, id 172.25.1.26

BGP Peer Information
* 2 205.178.18.40, id 205.178.18.40, hold 90, established
```

IP information includes:

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

/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 2, 205.178.18.210,  if  1, renter, prio 100, master, server
  2: vrid 1, 205.178.18.202,  if  1, renter, prio 100, backup
  3: vrid 3, 205.178.18.204,  if  1, renter, prio 100, master, proxy
```

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

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

/info/trunk
Trunk Group Information

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

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

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

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

```
Current Bandwidth Management setting: ON
  Policy Enforcement:enabled
Contract number 6:
  name
  policy            4
  precedence        6
  save history      enabled
  overwrite TOS     disabled
  and is currently  enabled

Policy 4:
     Hard Limit        10m
     Soft Limit         9m
     Reservation        8m
     Overlimit TOS       0
     Underlimit TOS      0
     Buffer Limit    32640

Set for
        Virtual Servers:
                1: 205.178.13.87
                  virtual ports: 80,
        Filters :
        VLANs :
        Default for Ports :
        Default for Trunk Groups :
```

/info/swkey
Software Enabled Keys

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

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

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

info/dump
Information Dump

Use the dump command to dump all switch information available from the Information Menu (10K or more, depending on your configuration). This data is useful for tuning and debugging switch performance.

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

CHAPTER 6

The Statistics Menu

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

/stats
Statistics Menu

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

Table 6-1 Statistics Menu Options (/stats)

Command Syntax and Usage

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

Displays the Port Statistics Menu for the specified port. Use this command to display traffic statistics on a port-by-port basis. Traffic statistics are included in SNMP Management Information Base (MIB) objects. To view menu options, refer to page 6-4.

`slb`

Displays the SLB (server load balancing) Menu. To view menu options, see page 6-5.

`bwm`

Displays the Bandwidth Management Menu. To view menu options, refer to page 6-21.

`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, refer to page 6-24.

`if` *<interface number (1-256)>*

Displays IP interface statistics for the management processors. To view an example of what is displayed on-screen, see page 6-25.

`ip`

Displays IP statistics. See page 6-25 for sample output.

`ICMP`

Display ICMP statistics. See page 6-26 for sample output.

`TCP`

Displays TCP statistics. See page 6-26 for sample output.

`UDP`

Displays UDP statistics. See page 6-26 for sample output.

`SNMP`

Displays SNMP statistics. See page 6-30 for sample output.

`Route`

Displays route statistics. See page 6-30 for sample output.

`ARP`

Displays ARP (Address Resolution Protocol) statistics. See page 6-30 for sample output.

Table 6-1 Statistics Menu Options (/stats)

Command Syntax and Usage
`DNS` Displays DNS (Domain Name Server) statistics. See page 6-30 for sample output.
`vrrp` When virtual routers are configured, you can display the following protocol statistics for VRRP: ■ Advertisements received (`vrrpInAdvers`) ■ Advertisements transmitted (`vrrpOutAdvers`) ■ Advertisements received, but ignored (`vrrpBadAdvers`) To view an example of what is displayed on-screen, see page 6-30.
`dump` Dumps all switch statistics. Use this command to gather data for tuning and debugging switch performance. If you want to capture dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump command. For more information, refer to page 6-30.

`/stats/port` *<port number>*
Port Statistics Menu

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

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

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

Table 6-2 Port Statistics Menu Options (/stats/port)

Command Syntax and Usage
`brg` Displays bridging (“dot1”) statistics for the port.
`ether` Displays Ethernet (“dot1”) statistics for the port.
`if` Displays interface statistics for the port.
`ip` Displays IP statistics for the port.
`link` Displays link statistics for the port.
`maint` Displays maintenance statistics for the port.

/stats/slb
Load Balancing Statistics

```
[Server Load Balancing Statistics Menu]
      port     - SLB Switch Port Stats Menu
      real     - Show real server stats
      group    - Show real server group stats
      virt     - Show virtual server stats
      filt     - Show filter stats
      gslb     - Show global SLB stats
      url      - Show URL SLB and Redirection stats
      ssl      - Show SSL SLB stats
      ftp      - Show FTP SLB parsing and NAT stats
      maint    - Show maintenance stats
      clear    - Clear non-operational Server Load Balancing stats
      dump     - Dump all SLB statistics
```

Table 6-3 SLB Statistics Menu Options (/stats/slb)

Command Syntax and Usage

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

Displays the following real server statistics:

- Number of times the real server has failed its health checks
- Number of sessions currently open on the real server
- Total sessions the real server was assigned
- Highest number of simultaneous sessions recorded for each real server
- Real server transmit/receive octets

To view an example of what is displayed on-screen, see page 6-7.

`group` *<real server group number (1-256)>*

Displays the following real server group statistics:

- Current and total sessions for each real server in the real server group.
- Current and total sessions for all real servers associated with the real server group.
- Highest number of simultaneous sessions recorded for each real server.
- Real server transmit/receive octets. For per-service octet counters, see the procedure on page 6-7.

To view an example of what is displayed on-screen, see page 6-8.

Table 6-3 SLB Statistics Menu Options (/stats/slb)

Command Syntax and Usage
`virt` *<virtual server number (1-256)>* Displays the following virtual server statistics: ■ Current and total sessions for each real server associated with the virtual server. ■ Current and total sessions for all real servers associated with the virtual server. ■ Highest number of simultaneous sessions recorded for each real server. ■ Real server transmit/receive octets. For per-service octet counters, see page 6-8. To view an example of what is displayed on-screen, see page 6-8.
`filt` *<filter ID (1-224)>* Displays the total number of times any filter has been used. See page 6-9 for sample output.
`port` *<port number (1-9)>* Displays the switch port statistics. See page 6-9 for sample output.
`gslb` Displays the Global SLB Statistics Menu. For more information, see page 6-12.
`url` Displays URL SLB and redirection statistics. See page 6-14 for sample output.
`ssl` Displays SSL server load balancing statistics. See page 6-16 for sample output.
`ftp` Displays FTP SLB parsing and NAT statistics. See page 6-16 for sample output.
`maint` Displays SLB maintenance statistics. See page 6-18 for sample output.
`clear [y/n]` Clears all non-operating SLB statistics on the switch, resetting them to zero. This command does not reset the switch and does ***not*** affect the following counters: ■ Counters required for Layer 4 and Layer 7 operation (such as current real server sessions). ■ All related SNMP counters. To view the statistics reset by this command, refer to Table 6-7 on page 6-19.
`dump` Dumps all switch SLB statistics. Use this command to gather data for tuning and debugging switch performance. To save dump data to a file, set your communication software on your workstation to capture session data prior to issuing the dump command.

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

Real Server SLB Statistics

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

NOTE – Octets are provided per server, not per service, unless configured as described below.

Per Service Octet Counters

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

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

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

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

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

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

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

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

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

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

Real Server Group Statistics

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

Real server group statistics include the following:

- Current and total sessions for each real server in the real server group.
- Current and total sessions for all real servers associated with the real server group.
- Highest number of simultaneous sessions recorded for each real server.
- Real server transmit/receive octets. For per-service octet counters, see the procedure on page 6-7.

`/stats/slb/virt` *<virtual server number>*

Virtual Server SLB Statistics

```
Virtual server 1 stats:
                      Current      Total  Highest
Real IP address      Sessions   Sessions Sessions           Octets
---- --------------- -------- ---------- --------  ---------------
  1  200.100.10.14         20         60        9           480000
  2  200.100.10.15         20         77       12           616000
---- --------------- -------- ---------- --------  ---------------
     200.100.10.20         40         137      21          1096000
```

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

Virtual server statistics include the following:

- Current and total sessions for each real server associated with the virtual server.
- Current and total sessions for all real servers associated with the virtual server.
- Highest number of simultaneous sessions recorded for each real server.
- Real server transmit/receive octets. For per-service octet counters, see page 6-7.

`/stats/slb/filt` *<filter number>*

Filter SLB Statistics

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

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

`/stats/slb/port` *<port number>*

Port SLB Statistics

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

Table 6-4 Switch Processor SLB Statistics Menu Options (/stats/slb/sp)

Command Syntax and Usage

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

Displays real server statistics for the selected port. To view an example of what is displayed on-screen, see page 6-10.

`group` *<real server group number (1-256)>*

Displays real server group statistics for the selected port. To view an example of what is displayed on-screen, see page 6-10.

`virt` *<virtual server number (1-256)>*

Displays virtual server statistics for the selected port. To view an example of what is displayed on-screen, see page 6-11.

`filt` *<filter ID (1-224)>*

Displays filter statistics for the selected port To view an example of what is displayed on-screen, seepage 6-11

`maint`

Displays maintenance statistics for the selected port. To view an example of what is displayed on-screen, see page 6-11.

Table 6-4 Switch Processor SLB Statistics Menu Options (/stats/slb/sp)

Command Syntax and Usage
`clear` Clears the following non-operating SLB statistics for this port, resetting them to zero: Real server stats: Octets, Total sessions Real server group: Octets, Total sessions Virtual server: Octets, Total sessions Total firings: Octets

`/stats/slb/port` *<number>*`/real` *<server number>*

Port Real Server SLB Statistics

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

`/stats/slb/port` *<number>*`/group` *<group number>*

Port Real Server Group SLB Statistics

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

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

Port Virtual Server SLB Statistics

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

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

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

Port Filter SLB Statistics

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

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

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

Port Maintenance SLB Statistics

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

/stats/slb/gslb
Global SLB Statistics

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

Table 6-5 Global SLB Statistics Menu Options (/stats/slb/gslb)

Command Syntax and Usage
real *<real server number (1-255)>* Where the real server number represents the real server ID on this switch, under which the remote server is configured. To view an example and description of what is displayed on-screen, see page 6-12.
group *<real server group number (1-256)>* To view an example and description of what is displayed on-screen, see page 6-13.
virtual *<virtual server number (1-256)>* To view an example and description of what is displayed on-screen, see page 6-13.
maint To view an example and description of what is displayed on-screen, see page 6-14.

/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     205.178.13.54           1240               30
        2     205.178.13.223           608               12
  ----------- --------------- ----------------  ---------------
     Totals                           1848               42
```

Real server group global statistics include the following:

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

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

Virtual Server Global SLB Statistics

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

  telnet   v1    205.178.13.55             4             31032
```

Virtual server global statistics include the following:

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

- Response time: the average time (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/url

SLB URL and Redirection Statistics

```
[URL SLB and Redirection Statistics Menu]
      redir   - Show URL Redirection stats
      lb      - Show URL SLB stats
      maint   - Show URL SLB/Redir Maintenance stats
```

/stats/slb/url/redir

URL SLB Redirection Statistics

```
Total URL based web cache redirection stats:
Total cache server hits:               73942
Total origin server hits:               2244
Total none-GETs hits:                      0
Total ’Cookie: ’ hits:                     0
Total no-cache hits:                       0
```

/stats/slb/url/lb

URL SLB Statistics

```
SLB String stats:
  ID Server Load Balance String                   Hits
   1 any                                         73881
   2 .gif                                            0
   3 /sales                                          0
   4 /xitami                                    162102
   5 /manual                                         0
   6 .jpg                                            0
```

/stats/slb/url/maint

URL Maintenance Statistics

```
URL SLB/Redir maintenance stats:
No available URL LB server:                        0
Clients reset by switch:                           0
Connection Splicing to support HTTP/1.1:           0
Half open connections:                             0
Switch retries:                                    0
Current available SP memory units:               648
Current SEQ buffer entries:     0    Highest:                    0
Current URL buffer use:         0    Highest:                    0
Current SP buffer entries:      0    Highest:                    0
Alloc Fails - Seq buffers:      0    Alloc Fails - Ubufs:        0
Max sessions per bucket:        0    Max frames per session:     0
Max bytes buffered (sess):      0
```

/stats/slb/ssl

SLB Secure /Socket Layer Statistics

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

/stats/slb/ftp

SLB File Transfer Protocol Statistics

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

/stats/slb/ftp/active

Active FTP SLB Parsing and Filter Statistics

```
>> FTP SLB parsing and Filter Statistics# dump
Total FTP :                                0
Total FTP NAT Filtered:                    0
Total new active FTP NAT Index:            0
Total new FTP SLB parsing Index:           0
FTP Active FTP NAT ACK/SEQ diff:           0
FTP SLB parsing ACK/SEQ diff:              0
```

/stats/slb/ftp/parsing

Passive FTP SLB Parsing Statistics

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

/stats/slb/ftp/maint

FTP SLB Maintenance Statistics

```
FTP Buffer copy error:                     0
FTP mode switch error:                     0
```

/stats/slb/ftp/dump

FTP SLB Statistics Dump

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

/stats/slb/maint

SLB Maintenance Statistics

```
SLB Maintenance stats:
Maximum sessions:                 516096
Current sessions:                      0
Allocation failures:                   0
TCP fragments:                         0
UDP datagrams:                         0
Non TCP/IP frames:                     0
Incorrect VIPs:                        0
Incorrect Vports:                      0
No available real server:              0
Backup server activations:             0
Overflow server activations:           0
Filtered (denied) frames:              0
VMA discards:                          0
```

SLB Maintenance statistics are described in the following table.

Table 6-6 Server Load Balancing Maintenance Statistics

Statistic	Description
Current Sessions	Number of session bindings currently in use.
Allocation Failures	Indicates instances where the switch ran out of available bindings for a port.
TCP Fragments	Indicates the number of TCP fragments encountered by the switch. Layer 4 processing might not handle TCP fragments, depending on configuration.
UDP Datagrams	Indicates that the virtual server IP address and MAC are receiving UDP frames when UDP balancing is not turned on.
Non TPC/IP Frames	Indicates the number of non-IP based frames received by the virtual server.
Incorrect VIPs	This indicates the number of times the switch has received a Layer 4 request for a virtual server which was not configured.
Incorrect Vports	This dropped frames counter indicates that the virtual server has received frames for TCP/UDP services that have not been configured. Normally this indicates a mis-configuration on the virtual server or the client, but it may be an indication of a potential security probing application like SATAN.

Table 6-6 Server Load Balancing Maintenance Statistics

Statistic	Description
No Server Available	This dropped frames counter indicates that all real servers are either out of service or at their `maxcon` limit.
Backup Server Activations	This indicates the number of times a real server failure has occurred and caused a backup server to be brought online.
Overflow Server Activations	This indicates the number of times a real server has reached the `maxcon` limit and caused an overflow server to be brought online.
Filtered (Denied) Frames	This indicates the number of frames that where dropped because they matched an active filter with the "`deny`" action set.

/stats/slb/clear

Clearing the SLB Statistics

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

Table 6-7 SLB Statistics Reset Using the **`/stats/slb/clear`** Command

Real server stats: Health check failures Total sessions Highest sessions Octets
Real server group stats: Total sessions Highest sessions Octets
Virtual server stats: Total sessions Highest sessions Octets
Filter stats: Total firings
SLB switch port stats: Per port: Real server stats: Octets, Total sessions Real server group: Octets, Total sessions Virtual server: Octets, Total sessions Total firings: Octets

Table 6-7 SLB Statistics Reset Using the **`/stats/slb/clear`** Command

- Global SLB stats:
  - Per real server:
    - DNS handoffs
    - HTTP redirects
  - Per server group:
    - DNS handoffs
    - HTTP redirects

- URL SLB and Redirection stats:
  - Redir:
    - Total cache server hits
    - Total origin server hits
    - Total none-GETs hits
    - Total ’Cookie: ’ hits
    - Total no-cache hits
  - LB:
    - ID SLB String hits

- SSL SLB stats:
  - Total Sessions
  - Highest Sessions

- FTP SLB parsing and NAT stats:
  - Total FTP
  - Total FTP NAT Filtered
  - Total new active FTP NAT Index
  - Total new FTP SLB parsing Index
  - FTP Active FTP NAT ACK/SEQ diff
  - FTP SLB parsing ACK/SEQ diff

/stats/bwm
Bandwidth Management Statistics

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

Table 6-8 Bandwidth Management Statistics Menu Options (/stats/bwm)

Command Syntax and Usage

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

Displays Switch processor Contract Statistics Menu. To view menu options, refer to page 6-22.

`cont` *<BW Contract number (1-256)>*

Displays bandwidth management contract statistics.

`rcont` *<BW Contract number (1-256)>*

Displays bandwidth management contract rate statistics.

`hist`

Displays bandwidth management history statistics.

`dump`

Displays all bandwidth management statistics.

/stats/bwm/sp

Bandwidth Management Switch Processor Statistics

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

Table 6-9 Management Processor Statistics Menu Options (/stats/bwm/sp)

Command Syntax and Usage

`cont` *<BW Contract number (1-256)>*

Displays bandwidth management contract statistics.

`rcont` *<BW Contract number (1-256)>*

Displays bandwidth management contract rate statistics.

/stats/bwm/cont *<contract number>*

Bandwidth Management Contract Statistics

```
------------------------------------------------------------------
BW Contract statistics
Contract Name              Octets    Discards  BufUsed BufMax
-------- --------------- ---------- ---------- ------- ------
     1                            0          0       0  32640
     2                            0          0       0  32640
```

Use this command to show statistics for all contracts or a specific contract.

/stats/bwm/rcont

Bandwidth Management Contract Rate Statistics

```
-----------------------------------------------------------------------
BW Contract statistics
Contract Name            Rate(Kbps)   Octets    Discards  BufUsed BufMax
-------- --------------- ---------- ---------- ---------- ------- -----
     6                            0          0          0      0 293760
   256   Default                  8    7476567          0      0 293760
     4                            0          0          0      0 293760
     6                            0          0          0      0 293760
   256   Default                  3    7477355          0      0 293760
     4                            0          0          0      0 293760
     6                            0          0          0      0 293760
   256   Default                  1    7477681          0      0 293760
     4                            0          0          0      0 293760
     6                            0          0          0      0 293760
   256   Default                 12    7480867          0      0 293760
     4                            0          0          0      0 293760
     6                            0          0          0      0 293760
   256   Default                  1    7481129          0      0 293760
```

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

/stats/bwm/hist

Bandwidth Management History Statistics

```
------------------------------------------------------------------
BW History statistics
Cont   Octets    Discards  TimeOver
---- ---------- ---------- --------
   1          0          0        0
1024          0          0        0
```

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

/stats/mp
Management Processor Statistics

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

Table 6-10 Management Processor Statistics Menu Options (/stats/mp)

Command Syntax and Usage

`mem`

Displays STEM memory statistics, showing available memory.

`amem`

Displays all STEM memory blocks in use to check for leaks.

`dma`

Displays DMA exception counts.

`pkt`

Displays packet statistics, to check for leads and load.

`tcb`

Displays all TCP control blocks (TCB) in use.

`uart`

Displays UART statistics.

/stats/if

Interface Statistics

```
IP interface 1 statistics:
ifInOctets:       2435148747   ifInUcastPkts:       1000174
ifInNUCastPkts:      2365278   ifInDiscards:              0
ifInErrors:                0   ifInUnknownProtos:        27
ifOutOctets:               0   ifOutUcastPkts:            0
ifOutNUcastPkts:           0   ifOutDiscards:             0
ifOutErrors:               0   ifStateChanges             0
```

/stats/ip

IP Statistics

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

/stats/icmp
ICMP Statistics

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

/stats/tcp
TCP Statistics

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

/stats/udp
UDP Statistics

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

/stats/snmp

SNMP Statistics

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

/stats/fdb
FDB Statistics

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

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

FDB statistics are described in the following table.

Table 6-11 Forwarding Database Statistics

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

/stats/route
Route Statistics

```
Route statistics:
ipRoutesCur:               8  ipRoutesHighWater:          8
ipRoutesMax:            1024

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

/stats/arp
ARP Statistics

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

This menu option enables you to display

/stats/dns
DNS Statistics

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

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

/stats/vrrp
VRRP Statistics

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:

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

The statistics for the VRRP LAN are displayed:

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

/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 difference are called out in the text.

To make finding information easier, the menu options under the Server Load Balancing Menu (/cfg/slb) are in Chapter 8.

/cfg
Configuration Menu

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

Table 7-1 Configuration Menu Options (/cfg)

Command Syntax and Usage

`sys`

Displays the System Configuration Menu.To view menu options, refer to page 7-6.

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

Displays the Port Configuration Menu. To view menu options, refer to page 7-13.

`ip`

Displays the IP Configuration Menu. To view menu options, refer to page 7-17.

`vlan` *<VLAN number (1-4094)>*

Displays the VLAN Configuration Menu. To view menu options, refer to page 7-32.

`stp`

Displays the Spanning Tree Configuration Menu. To view menu options, refer to page 7-34.

`snmp`

Displays the SNMP Configuration Menu. To view menu options, refer to page 7-38.

`mirr`

Displays the Mirroring Configuration Menu. Use this menu to configure mirroring operation. There can be a total of four mirroring selectors. One of these selectors will be used to configure both address and non address-based mirroring selection criteria, and the other three selectors will be used only for non address-based ones. The address-based selection criteria allows user to specify MAC DA, MAC SA, IP DA, and/or IP SA, in addition to in-port, out-port, in-VLAN ID, and/or COS. The maximum number of configurable monitoring ports is 2. To view menu options, refer to page 7-42.

`slb`

Displays the Server Load Balancing Configuration Menu. To view menu options, refer to Chapter 8, “The SLB Configuration Menu.

`trunk` *<group number (1-4)>*

Displays the Trunk Group Configuration Menu. To view menu options, refer to page 7-44.

`vrrp`

Displays the Virtual Router Redundancy Configuration Menu. To view menu options, refer to page 7-45.

Table 7-1 Configuration Menu Options (/cfg)

Command Syntax and Usage
`bwm` Displays the Bandwidth Management Configuration Menu. To view menu options, refer to page 7-56.
`setup` Step-by-step configuration set-up of the switch. For more information, refer to page 7-40.
`dump` Dumps current configuration to a script file. For more information, refer to page 7-40.
`ptcfg` *<host name or IP address>* *<filename on host>* Backs up current configuration to TFTP server. For more information, refer to page 7-41.
`gtcfg` *<host name or IP address>* *<filename on host>* Restores current configuration from TFTP server. For more information, refer to page 7-41.

Viewing, Applying, and Saving Changes

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

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

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

Viewing Pending Changes

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

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

Applying Pending Changes

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

```
# apply
```

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

NOTE – All configuration changes take effect immediately when applied, except for starting Spanning-Tree Protocol. To turn STP on or off, you must apply the changes, save them (see below), and then reset the switch (see “Resetting the Switch” on page 10-4).

Saving the Configuration

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

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

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

```
# save
```

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

```
# save n
```

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

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

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

For instructions on selecting the configuration to run at the next system reset, see "Selecting a Configuration Block" on page 10-4.

/cfg/sys
System Configuration

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

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

Table 7-2 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage
`sshd` Displays the SSH Server Menu. To view menu options, refer to page 7-8.
`radius` Displays the RADIUS Authentication Menu. To view menu options, refer to page 7-9.
`ntp` Displays the NTP Server Menu. To view menu options, refer to page 7-10.
`date` Prompts the user for the system date.
`time` Configures the system time using a 24-hour clock format.

Table 7-2 System Configuration Menu Options (/cfg/sys)

Command Syntax and Usage

`idle` *<idle timeout in minutes; affects both console and Telnet>*

Sets the idle timeout for CLI sessions, from 1 to 60 minutes. The default is 5 minutes.

`snmp`

Disables or provides read-only/write-read SNMP access.

`wport` *<TCP port number (1-65535)>*

Sets the switch port used for serving switch Web content. The default is HTTP port 80. If Global Server Load Balancing is to be used, set this to a different port (such as 8080).

`bannr` *<string, maximum 80 characters>*

Configures a login banner of up to 80 characters. When a user or administrator logs into the switch, the login banner is displayed; it is also displayed as part of the output from the `/info/sys` command.

`mnet` *<IP subnet (e.g., 192.4.17.0)>*

Sets the base source IP address allowed to access switch management through Telnet, SNMP, RIP, or the WebOS browser-based interface. A range of IP addresses is produced when used with `mmask` (below). Specify an IP address in dotted-decimal notation.

`mmask` *<IP subnet mask (e.g., 255.255.0.0)>*

This IP address mask is used with `mnet` to set a range of source IP addresses allowed access to switch management functions. Specify the mask in dotted-decimal notation.

`smtp` *<SMTP host name or IP address>*

Sets SMTP host.

`tnet disable|enable`

Enables or disables telnet access to the command line interface sessions. This command is available only from a local console connection.

`bootp disable|enable`

Enables or disables the use of `BOOTP`; if you enable `BOOTP`, the switch will query its `BOOTP` server for all of the switch IP parameters.

`http disable|enable`

Enables/disables HTTP (Web) access to the browser-based interface.

`user`

Displays the User Access Control Menu. To view menu options, refer to page 7-11.

`cur`

Displays the current system parameters.

/cfg/sys/sshd

SSH Server Configuration

```
[SSHD Menu]
      intrval - Set Interval for generating the RSA server key
      scpadm  - Set SCP-only admin password
      hkeygen - Generate the RSA host key
      skeygen - Generate the RSA server key
      ena     - Enable the SCP apply and save
      dis     - Disable the SCP apply and save
      on      - Turn SSH server ON
      off     - Turn SSH server OFF
      cur     - Display current SSH server configuration
```

For Alteon AD4 and 184 switches, this menu enables Secure Shell access from any SSH client.

Table 7-3 System Configuration Menu Options (/cfg/sys/sshd)

Command Syntax and Usage

`intrval` *<number of hours (0-24)>*

Sets the interval for automatically re-generating the RSA server key.

`scpadm`

Sets the SCP-only administrator password, up to 15 characters. The command will prompt for the required information.

`hkeygen`

Generates the RSA host key.

`skeygen`

Generates the RSA server key.

`ena`

Enables the SCP apply and save.

`dis`

Disables the SCP apply and save.

`on`

Enables the SSH server.

`off`

Disables the SSH server.

`cur`

Displays the current SSH server configuration.

/cfg/sys/radius

RADIUS Server Configuration

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

Table 7-4 System Configuration Menu Options (/cfg/sys/radius)

Command Syntax and Usage

`prisrv` *<IP address>*

Sets the primary RADIUS server address.

`secsrv` *<IP address>*

Sets the secondary RADIUS server address.

`secret` *<1-32 character secret>*

This is the shared secret between the switch and the RADIUS server(s).

`port` *<RADIUS port>*

Enter the number of the UDP port, between 1500 - 3000. The default is 1645.

`retries` *<RADIUS server retries (1-3)>*

Sets the number of failed authentication requests before switching to a different RADIUS server. The default is 3 requests.

`timeout` *<RADIUS server timeout seconds (1-10)>*

Sets the amount of time, in seconds, before a RADIUS server authentication attempt is considered to have failed. The default is 3 seconds.

`telnet disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables RADIUS backdoor for telnet.

`on`

Enables the RADIUS server.

`off`

Disables the RADIUS server.

`cur`

Displays the current RADIUS server parameters.

/cfg/sys/ntp

NTP Server Configuration

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

This menu enables you to synchronize the switch clock to a Network Time Protocol (NTP) server.

Table 7-5 System Configuration Menu Options (/cfg/sys/ntp)

Command Syntax and Usage

`server` *<IP address>*

Specifies the IP address of the NTP server you want to synchronize the switch clock to.

`intrval` *<resync interval in minutes>*

Specifies the interval, that is, how often, in minutes (1-2880), to resynchronize the switch clock with the NTP server.

`tzone` *<timezone offset, in hours>*

Specifies the timezone offset, in hours, of the switch you are synchronizing from Greenwich Mountain Time (GMT).

`dlight disable`|**`enable`** (or just **`d`**|**`e`**)

Disables/enables whether the switch is in a daylight savings time scenario. When enabled, the switch will add an extra hour to the system clock so that it is consistent with the local clock. By default, this option is disabled.

`on`

Enables the NTP synchonization service.

`off`

Disables the NTP synchonization service.

`cur`

Displays the current NTP service settings.

/cfg/sys/user

User Access Control Configuration

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

NOTE – Passwords can be a maximum of 15 characters.

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

Command Syntax and Usage

`usrpw`

Sets the user (`user`) password. The user has no direct responsibility for switch management. He or she can view switch status information and statistics, but cannot make any configuration changes.

`sopw`

Sets the SLB operator (`slboper`)password. The SLB operator manages Web servers and other Internet services and their loads. He or she can view all switch information and statistics and can enable/disable servers using the Server Load Balancing configuration menus.

Access includes "user" functions.

`l4opw`

Sets the Layer 4 operator (`l4oper`)password. The Layer 4 operator manages traffic on the lines leading to the shared Internet services. He or she can view all switch information and statistics, and can ...

Access includes "`slboper`" functions.

`opw`

Sets the operator (`oper`)password; the operator password can have a maximum of 15 characters. The operator manages all functions of the switch. He or she can view all switch information and statistics and can reset ports or the entire switch.

Access includes “`l4oper`” functions.

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

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

/cfg/port *<port number>*
Port Configuration

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

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

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

Command Syntax and Usage
fast If a port is configured to support Fast Ethernet, this option displays the Fast Ethernet Physical Link menu. To view menu options, refer to page 7-15.
gig If a port is configured to support Gigabit Ethernet, this option displays the Gigabit Ethernet Physical Link menu. To view menu options, refer to page 7-15.
pref If dual physical connectors are available on the port, this option defines the preferred physical connector. Choices are: ■ Fast Ethernet Port, RJ-45 connector ■ Gigabit Ethernet Port, SC fiber connector (default)
back If dual physical connectors are available on the port, this option defines the physical connector to use when the preferred choice fails or is unavailable. Choices are: ■ Fast Ethernet Port, RJ-45 connector (default) ■ Gigabit Ethernet Port, SC fiber connector ■ None

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

Command Syntax and Usage
`pvid` Sets the default VLAN number which will be used to forward frames which are not VLAN tagged.
`name` Sets a name for the port. The assigned port name appears next to the port number on some information and statistics screens.
`cont` *<BW Contract (1-256)>* Sets the default Bandwidth Contract for this port.
`rmon disable`|**`enable`** (or just **`d`**|**`e`**) Disables/enables RMON for this port.
`tag disable`|**`enable`** (or just **`d`**|**`e`**) Disables/enables VLAN tagging for this port.
`iponly disable`|**`enable`** (or just **`d`**|**`e`**) Disables/enables allowing only IP-related frames.
`ena` Enables the port.
`dis` Disables the port. (To temporarily disable a port without changing its configuration attributes, refer to “Temporarily Disabling a Port” on page 7-16.)
`cur` Displays current port parameters.

/cfg/port *<port number>* fast|gig

Port Link Configuration

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

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

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

Link menu options are described in Table 7-8 and appear on the `fast` and `gig` port configuration menus for the ACEswitch 180, and on the Port Menu for some models of Alteon WebSystems' switches. Using these configuration menus, you can set port parameters such as speed, flow control, and negotiation mode for the port link.

Table 7-8 Port Link Configuration Menu Options (/cfg/port *<number>* fast|gig)

Command Syntax and Usage
`speed 10`|**`100`**|**`1000`**|**`any`** (not all options are valid on all ports) Sets the link speed; the choices include: ■ "Any," for automatic detection (default) ■ 10 Mbps ■ 100 Mbps ■ 1000 Mbps
`mode full`|**`half`**|**`any`** Sets the operating mode; the choices include: ■ "Any," for autonegotiation (default) ■ Full-duplex ■ Half-duplex

Table 7-8 Port Link Configuration Menu Options (/cfg/port <*number*> fast|gig)

Command Syntax and Usage

`fctl rx|tx|both|none`

Sets the flow control; the choices include:

- Autonegotiation (default)
- Receive flow control
- Transmit flow control
- Both receive and transmit flow control
- No flow control

`auto on|off`

Enable or disable auto-negotiation for the port.

`cur`

Displays current port parameters.

Temporarily Disabling a Port

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

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

Because this sets a temporary state for the port, you do not need to use `apply` or `save`. See “The Operations Menu on page 9-1 for other operations-level commands.

/cfg/ip
IP Configuration

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

Table 7-9 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage

`if` *<interface number (1-256)>*

Displays the IP Interface Menu. To view menu options, refer to page 7-19.

`gw` *<default gateway number (1-4)>*

Displays the IP Default Gateway Menu. To view menu options, refer to page 7-20.

`route`

Displays the IP Static Route Menu. To view menu options, refer to page 7-21.

`frwd`

Displays the IP Forwarding Menu. To view menu options, refer to page 7-22.

`rip1`

Displays the Routing Interface Protocol Menu. To view menu options, refer to page 7-24.

`bgp`

Displays the Border Gateway Protocol Menu. To view menu options, refer to page 7-24.

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

Displays the IP Port Menu. To view menu options, refer to page 7-29.

Table 7-9 IP Configuration Menu Options (/cfg/ip)

Command Syntax and Usage
`dns` Displays the IP Domain Name System Menu. To view menu options, refer to page 7-30.
`log` *<syslog host IP address (e.g., 192.4.17.223)>* Sets the IP address of the syslog host.
`log2` *<syslog host IP address (e.g., 192.4.17.223)>* Sets the IP address of the second syslog host.
`logfac` *<syslog host local facility (0-7)>* Sets the facility of the syslog host.
`log2fac` *<syslog host local facility (0-7)>* Sets the facility of the second syslog host.
`rearp` *<2-120 minutes>* Sets the re-ARP period in minutes. The switch periodically sends ARP (Address Resolution Protocol) requests to refresh its address database. This command is used for setting the interval between ARP refreshes of the next IP address in the database.
`metrc strict|roundrobin` Sets the default gateway metric.
`cur` Displays the current IP configuration.

/cfg/ip/if *<interface number>*
IP Interface Configuration

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

The switch can be configured with up to 256 IP interfaces. Each IP interface represents the switch on an IP subnet on your network.

Table 7-10 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.
broad *<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 one VLAN, though any VLAN can have multiple IP interfaces in it.
ena Enables this interface.
dis Disables this interface.
del Removes this interface.
cur Displays the current interface settings.

`/cfg/ip/gw` *<gateway number>*

Default IP Gateway Configuration

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

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

Table 7-11 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.

`intr` *<value (0-60 seconds)>*

The switch pings the default gateway to verify that it's up. The `intr` option sets the time between health checks. The range is from 1 to 120 seconds. The default is 2 seconds.

`retry` *<attempts (1-120)>*

Sets the number of failed health check attempts required before declaring this default gateway inoperative. The range is from 1 to 120 attempts. The default is 8 attempts.

`arp`

Enables/disables ARP-only health checks.

`ena`

Enables the gateway for use.

`dis`

Disables the gateway.

`del`

Deletes this gateway from the configuration.

`cur`

Displays the current gateway settings.

Default Gateway Metrics

For information about configuring which gateway is selected when multiple default gateways are enabled, see page 7-31.

/cfg/ip/route

IP Static Route Configuration

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

Table 7-12 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, and gateway address. Enter all addresses using dotted decimal notation.
`rem` *<destination>* Removes a static route. The destination address of the route to remove must be specified using dotted decimal notation.
`cur` Displays the current IP static routes.

/cfg/ip/frwd

IP Forwarding Configuration

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

Table 7-13 IP Forwarding Options (/cfg/ip/frwd)

Command Syntax and Usage
`local` Displays the menu used to define local network for route caching.Up to five local networks (lnets) can be configured. To view menu options, refer to page 7-22.
`dirbr` Enables/disables forwarding directed broadcasts.
`on` Enables IP forwarding (routing) on the switch.
`off` Disables IP forwarding (routing) on the switch.
`cur` Displays the current IP forwarding settings.

/cfg/ip/frwd/local

Local Network Route Caching Definition

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

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

Table 7-14 IP Local Network Options (/cfg/ip/frwd/local)

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

Defining IP Address Ranges for the Local Route Cache

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

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

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

Table 7-15 Local Routing Cache Address Ranges

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

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

/cfg/ip/rip1

Routing Information Protocol Configuration

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

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

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

Table 7-16 Routing Information Protocol Menu (/cfg/ip/rip1)

Command Syntax and Usage

`updat` *<update period (1-120 seconds)>*

Sets the RIP update period in seconds.

`spply disable|enable` (or just **`d|e`**)

When enabled, the switch supplies routes to other routers.

`lsten disable|enable` (or just **`d|e`**)

When enabled, the switch learns routes from other routers.

`deflt none|lsten|spply|both`

When enabled, the switch accepts RIP default routes from other routers and gives them priority over configured default gateways. When disabled, the switch rejects RIP default routes.

`poisn disable|enable` (or just **`d|e`**)

When enabled, the switch uses split horizon with poisoned reverse. When disabled, the switch uses only split horizon.

`on`

Globally turns RIP ON.

`off`

Globally turns RIP OFF.

`cur`

Displays the current RIP configuration.

/cfg/ip/bgp

Border Gateway Protocol Configuration

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

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

The BGP Menu enables you to configure the switch to receive routes and to advertise static routes, fixed routes and virtual IP addresses with other internal and external routers. In the current WebOS implementation, we do not advertize BGP routes learned from other BGP "speakers."

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

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

Table 7-17 Border Gateway Protocol Menu (/cfg/ip/bgp)

Command Syntax and Usage
`peer` *<peer number (1 -4)>* Displays the menu used to configure each BGP *peer*; that is, each border router within an autonomous system that exchanges routing information with routers on other external networks. To view menu options, refer to page 7-26.
`filt` *<filter number (1-4)>* Displays the menu used to configure the range of IP destinations accepted by each BGP peer filter. To view menu options, refer to page 7-28.
`on` Globally turns BGP on.

Table 7-17 Border Gateway Protocol Menu (/cfg/ip/bgp)

Command Syntax and Usage

`off`

Globally turns BGP off.

`cur`

Displays the current BGP configuration.

/cfg/ip/bgp/peer

BGP Peer Configuration

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

This menu is used to configure BGP peers; that is, border routers that exchange routing information with routers on internal and external networks.

Table 7-18 BGP Peer Configuration Options (/cfg/ip/bgp/peer)

Command Syntax and Usage

`addr` *<IP address (e.g., 192.4.17.101)>*

Defines the IP address for the specified peer (border router), using dotted decimal notation. The default address is 0.0.0.0.

`ras` *<AS number (0-65535)>*

Sets the remote autonomous system number for the specified peer.

Table 7-18 BGP Peer Configuration Options (/cfg/ip/bgp/peer)

Command Syntax and Usage

`if` *<interface number (1-256)>*

Selects a switch IP interface (between 1 and 256) for the specified peer. The default value is 1.

`las` *<AS number (0-65535)>*

Sets the local autonomous system number for the specified peer.

`hold` *<hold time (0, 3-65535)>*

Sets the period of time, in seconds, that will elapse before the peer session is torn down because the switch hasn't received a "keep alive" message from the peer.

`ttl` *<number of router hops (1-255)>*

Specifies the number of router hops that the IP datagram can make. This value is used to restrict the number of "hops" the advertisement makes. It is also used to support multi-hops, which allow BGP peers to talk across a routed network.

`metric <`*metric (1-255)***`>`**

Sets the length of the AS path used when advertising eBGP routes. When advertising iBGP routes, this parameter sets the local preference.

`fixed disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables advertising fixed routes.

`static disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables advertising static routes. No default route is advertised.

`vip disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables advertising virtual server routes.

`ena`

Enables this peer configuration.

`dis`

Disables this peer configuration.

`del`

Deletes this peer configuration.

`cur`

Displays the current BGP peer configuration.

/cfg/ip/bgp/filt

BGP Filter Configuration

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

This menu enables you to configure filters that specify the routes/range of IP destinations a peer router will accept from other peers. A route must match a filter to be installed in the routing table.

Table 7-19 BGP Filter Configuration Options (/cfg/ip/bgp/filt)

Command Syntax and Usage

`addr` *<IP address (e.g., 192.4.17.101)>*

Defines the starting IP address for this filter, using dotted decimal notation. The default address is 0.0.0.0.

`mask` *<IP address>*

This IP address mask is used with `addr` to define the range of IP addresses that will be accepted by the peer when the filter is enabled.

`ena`

Enables this BGP filter.

`dis`

Disables this BGP filter.

`del`

Deletes this BGP filter.

`cur`

Displays the current BGP filter configuration.

/cfg/ip/port *<port number>*

IP Port Configuration

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

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

Table 7-20 IP Forwarding Port Options (/cfg/ip/port)

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

/cfg/ip/dns

Domain Name System Configuration

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

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

Table 7-21 Domain Name Service Menu Options (/cfg/ip/dns)

Command Syntax and Usage

`prima` *<IP address (e.g., 192.4.17.101)>*

You will be prompted to set the IP address for your primary DNS server. Use dotted decimal notation.

`secon` *<IP address (e.g., 192.4.17.101)>*

You will be prompted to set the IP address for your secondary DNS server. If the primary DNS server fails, the configured secondary will be used instead. Enter the IP address using dotted decimal notation.

`dname` *<dotted DNS notation>*|**none**

Sets the default domain name used by the switch.
For example: `mycompany.com`

`cur`

Displays the current Domain Name System settings.

`/cfg/ip/log` *<IP address>*

Syslog Host Configuration

If configured, the switch software logs the following types of messages to syslog host:

```
Apr 1 17:28:52 ALERT slb: cannot contact real server 215.118.113.74
Apr 1 17:29:10 NOTICE console: admin login
Apr 1 17:26:35 INFO web server: new configuration applied
Apr 1 17:26:35 WARNING slb: filter 10 fired on port 4
Apr 1 17:28:03 ERR telnet: no apply needed
```

Each syslog message has a criticality level associated with it, included in text form as a prefix to the log message. One of eight different prefixes is used, depending on the condition that the administrator is being notified of, as shown below.

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

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

Default Gateway Metrics

If multiple default gateways are configured and enabled, a metric can be set to determine which primary gateway is selected. There are two metrics; each is described in the table below:

Table 7-22 Default Gateway Metrics (/cfg/ip/metrc)

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

/cfg/vlan *<VLAN number>*
VLAN Configuration

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

The commands in this menu configure VLAN attributes, change the status of the VLAN, delete the VLAN, and change the port membership of the VLAN. For more information on configuring VLANs, see “Setup Part 3: VLANs” on page 3-7.

Table 7-23 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.
cont *<BW Contract (1-256)>* Sets the Bandwidth contract for this VLAN.
add *<port number>* Adds port(s) or trunk group(s) to the VLAN membership.
remove *<port number>* Removes port(s) or trunk group(s) from this VLAN.
def Defines the specified VLAN as a list of ports.
jumbo disable|**enable** (or just **d**|**e**)> Enables/disables support for Jumbo Frame support on this VLAN.
ena Enables this VLAN.

Table 7-23 VLAN Configuration Menu Options (/cfg/vlan)

Command Syntax and Usage
`dis` Disables this VLAN without removing it from the configuration.
`del` Deletes this VLAN.
`cur` Displays the current VLAN configuration.

NOTE – You cannot add a port to more than one VLAN unless the port has VLAN tagging turned on (see "/cfg/port <port number>" on page 7-13).

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` *<STP number>*
Spanning Tree Configuration

```
[Spanning Tree Group 1 Menu]
      brg     - Bridge parameter menu
      port    - Port parameter menu
      on      - Globally turn Spanning Tree ON
      off     - Globally turn Spanning Tree OFF
      cur     - Display current bridge parameters
```

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 7-24 Spanning Tree Configuration Menu (/cfg/stp)

Command Syntax and Usage

`brg`

Displays the Bridge Spanning Tree Menu. To view menu options, refer to page 7-35.

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

Displays the Spanning Tree Port Menu. To view menu options, refer to page 7-37.

`on`

Globally enables STP.

`off`

Globally disables STP.

`cur`

Displays current STP parameters.

/cfg/stp/brg

Bridge Spanning Tree Configuration

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

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

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

Table 7-25 Bridge Spanning Tree Menu Options (/cfg/stp/brg)

Command Syntax and Usage

`prior` *<new bridge priority (0-65535)>*

Configures the bridge priority. The bridge priority parameter controls which bridge on the network is the STP root bridge. To make this switch the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. The range is 0 to 65535, and the default is 32768

`hello` *<new bridge hello time (1-10 secs)>*

Configures the bridge hello time.The hello time specifies how often the root bridge transmits a configuration bridge protocol data unit (BPDU). Any bridge that is not the root bridge uses the root bridge hello value. The range is 1 to 10 seconds, and the default is 2 seconds.

`mxage` *<new bridge max age (6-40 secs)>*

Configures the bridge maximum age. The maximum age parameter specifies the maximum time the bridge waits without receiving a configuration bridge protocol data unit before it reconfigures the STP network. The range is 6 to 40 seconds, and the default is 20 seconds.

Table 7-25 Bridge Spanning Tree Menu Options (/cfg/stp/brg)

Command Syntax and Usage
`frwd` *<new bridge Forward Delay (4-30 secs)>* Configures the bridge forward delay parameter. The forward delay parameter specifies the amount of time that a bridge port has to wait before it changes from the listening state to the learning state and from the learning state to the forwarding state. The range is 4 to 30 seconds, and the default is 15 seconds.
`aging` *<new bridge Aging Time (1-65535 secs, 0 to disable)>* Configures the forwarding database aging time. The aging time specifies the amount of time the bridge waits without receiving a packet from a station before removing the station from the forwarding database. The range is 1 to 65535 seconds, and the default is 300 seconds. To disable aging, set this parameter to 0.
`current` Displays the current bridge STP parameters.

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

- $2*(\mathit{fwd}-1) \geq \mathit{mxage}$
- $2*(\mathit{hello}+1) \leq \mathit{mxage}$

/cfg/stp/port *<port number>*

Spanning Tree Port Configuration

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

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

- Port priority
- Port path cost

Table 7-26 Spanning Tree Port Menu (/cfg/stp/port *<port-number>*)

Command Syntax and Usage

`prior` *<new port Priority (0-255)>*

Configures the port priority. The port priority helps determine which bridge port becomes the designated port. In a network topology that has multiple bridge ports connected to a single segment, the port with the lowest port priority becomes the designated port for the segment. The range is 0 to 255, and the default is 128.

`cost` *<new port Path Cost (1-65535, 0 for default)>*

Configures the port path cost. The port path cost is used to help determine the designated port for a segment. Generally speaking, the faster the port, the lower the path cost. The range is 1 to 65535. The default is 10 for 100Mbps ports, and 1 for gigabit ports. A value of 0 indicates that the default cost will be computed for an autonegotiated link speed.

`on`

Enables STP on the port.

`off`

Disables STP on the port.

`cur`

Displays the current STP port parameters.

/cfg/snmp
SNMP Configuration

```
[SNMP Menu]
      name     - Set SNMP "sysName"
      locn     - Set SNMP "sysLocation"
      cont     - Set SNMP "sysContact"
      rcomm    - Set SNMP read community string
      wcomm    - Set SNMP write community string
      trap1    - Set first SNMP trap host address
      trap2    - Set second SNMP trap host address
      t1comm   - Set community string for first trap host
      t2comm   - Set community string for second trap host
      auth     - Enable/disable SNMP "sysAuthenTrap"
      linkt    - Enable/disable SNMP link up/down trap
      cur      - Display current SNMP configuration
```

The WebOS 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 7-27 SNMP Configuration Menu Options (/cfg/snmp))

Command Syntax and Usage

`name` *<new string, maximum 64 characters>*

Configures the name for the system. The name can have a maximum of 64 characters.

`locn` *<new string, maximum 64 characters>*

Configures the name of the system location. The system location can have a maximum of 64 characters.

`cont` *<new string, maximum 64 characters>*

Configures the name of the system contact. 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.

`trap1` *<new SNMP trap host IP address (e.g., 192.4.17.101)>*

Configures the IP address of the first SNMP trap host using dotted decimal notation. The SNMP trap host is the device that receives SNMP trap messages from the switch.

`trap2` *<new SNMP trap host IP address (e.g., 192.4.17.101)>*

Configures the IP address of the second SNMP trap host using dotted decimal notation.

`t1com` *<new trap host community string, maximum 32 characters>*

Configures the community string for the first trap host.

`t2com` *<new trap host community string, maximum 32 characters>*

Configures the community string for the second trap host.

`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.

`cur`

Displays the current STP port parameters.

/cfg/setup
Setup

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

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

```
Configuration# setup
```

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

/cfg/dump
Dump

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

```
Configuration# dump
```

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

The configuration script file can be used to configure other switches through a Telnet connection. When using Telnet to configure a new switch, paste the configuration commands from the script file at the command line prompt of the switch. The active configuration can also be saved or loaded via TFTP, as described on page 7-41.

`/cfg/ptcfg` *<TFTP server> <filename>*
Saving the Active Switch Configuration

When the `ptcfg` command is used, the switch’s active configuration commands (as displayed using `/cfg/dump`) will be uploaded to the specified script configuration file on the TFTP server. To start the switch configuration upload, at the `Configuration#` prompt, enter:

```
Configuration# ptcfg <server> <filename>
```

Where *server* is the TFTP server IP address or hostname, and *filename* is the name of the target script configuration file.

NOTE – If the TFTP server is running SunOS or the Solaris operating system, the specified `ptcfg` file must exist prior to executing the `ptcfg` command, and must be writable (set with proper permission, and not locked by any application). The contents of the specified file will be replaced with the current configuration data.

`/cfg/gtcfg` *<TFTP server> <filename>*
Loading the Active Switch Configuration

When the `gtcfg` command is used, the active configuration will be replaced with the commands found in the specified configuration file. The file can contain a full switch configuration or a partial switch configuration. The configuration loaded using `gtcfg` is not activated until the `apply` command is used. If the `apply` command is found in the configuration script file loaded using this command, the apply action will be performed automatically.

To start the switch configuration download, at the `Configuration#` prompt, enter:

```
Configuration# gtcfg <server> <filename>
```

Where *server* is the TFTP server IP address or hostname, and *filename* is the name of the target script configuration file.

/cfg/mirr/port
Port Mirroring Configuration

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

NOTE – Port mirroring menu options are accessible only to the Alteon AD4 and Alteon 184 Web switches.

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

There can be a total of four mirroring selectors. One of these selectors will be used to configure both address and non address-based mirroring selection criteria, and the other three selectors will be used only for non address-based ones. The address-based selection criteria allows user to specify MAC destinationA, MAC SA, IP DA, and/or IP SA, in addition to in-port, out-port, in-VLAN ID, and/or COS. The maximum number of configurable monitoring ports is 2.

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

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

Table 7-28 Port Mirroring Options (/cfg/mirr/port)

Command Syntax and Usage
`to` *<input port to be mirrored>* This defines the monitoring port. When port mirroring is enabled, packets received and/ or transmitted by the mirrored port will be duplicated to the switch port specified in this command.
`from` *<input port to be mirrored>* This defines the mirrored port. When port mirroring is enabled, packets received and/or sent by the port specified in this command will be sent to the monitor port.
`dir` This determines which type of packets will be sent to the monitor port: `in` = packets received at the mirrored port `out` = packets sent from the mirrored port `both` = packets sent and received by the mirrored port
`tmout` *<seconds after which mirroring gets disabled (1-86400, 0 for no timeout)>* Port mirroring will be automatically disabled (regardless of port state) after the time-out period specified in this command. Valid times are from 0 (does not time-out) to 86400 seconds.
`dis` Turns port mirroring off.
`ena` Turns port mirroring on.
`cur` Displays the current parameter settings.

/cfg/trunk
Trunk Configuration

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

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

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

Table 7-29 Trunk Configuration Menu Options (/cfg/trunk)

Command Syntax and Usage
`cont` *<BW Contract (1-256)>* Sets the default Bandwidth Contract for this trunk group.
`add` *<port number (1-9)>* Adds a physical port to the current trunk group.
`rem` *<port number (1-9)>* Removes a physical port from the current trunk group.
`ena` Enables the current trunk group.
`dis` Turns the current trunk group off.
`del` Removes the current trunk group configuration.
`cur` Displays current trunk group parameters.

/cfg/vrrp
VRRP Configuration

```
[Virtual Router Redundancy Protocol Menu]
      vr      - VRRP Virtual Router menu
      group   - VRRP Virtual Router Group menu
      if      - VRRP Interface menu
      track   - VRRP Priority Tracking menu
      hotstan - Enable/disable hot-standby processing
      on      - Globally turn VRRP ON
      off     - Globally turn VRRP OFF
      cur     - Display current VRRP configuration
```

Virtual Router Redundancy Protocol (VRRP) support on Alteon 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 7-30 Virtual Router Redundancy Protocol Options (/cfg/vrrp)

Command Syntax and Usage

`vr` *<virtual router number (1-256>*

Displays the VRRP virtual router menu. This menu is used for configuring up to 256 virtual routers on this switch. To view menu options, refer to page 7-46.

`group`

Displays the VRRP virtual router group menu, used to combine all virtual routers together as one logical entity. Group options must be configured when using two or more Alteon switches in a hot-standby failover configuration; where only one switch is active at any given time. To view menu options, refer to page 7-51.

`if` *<interface number (1-256)>*

Displays the VRRP virtual router interface menu. To view menu options, refer to page 7-53.

Table 7-30 Virtual Router Redundancy Protocol Options (/cfg/vrrp)

Command Syntax and Usage

`track` *<interface number (1-256)>*

Displays the VRRP tracking menu. This menu is used for weighting the criteria used when modifying priority levels in the master router election process. To view menu options, refer to page 7-54.

`hostan disable|enable` (or just **`d|e`**)

Enables/disables hot standby processing.

`on`

Globally enables VRRP on this switch.

`off`

Globally disables VRRP on this switch.

`cur`

Displays the current VRRP parameters.

/cfg/vrrp/vr *<router number>*
Virtual Router Configuration

```
[VRRP Virtual Router 1 Menu]
      track    - Priority Tracking Menu
      vrid     - Set virtual router ID
      addr     - Set IP address
      if       - Set interface number
      prio     - Set renter priority
      adver    - Set advertisement interval
      preem    - Enable/disable preemption
      share    - Enable/disable sharing
      ena      - Enable virtual router
      dis      - Disable virtual router
      del      - Delete virtual router
      cur      - Display current VRRP virtual router configuration
```

This menu is used for configuring up to 256 virtual routers for this switch. A virtual router is defined by its virtual router ID and an IP address. On each VRRP-capable routing device participating in redundancy for this virtual router, a virtual router will be configured to share the same virtual router ID and IP address.

Table 7-31 VRRP Virtual Router Options (/cfg/vrrp/vr)

Command Syntax and Usage
`track` Displays the VRRP priority tracking menu for this virtual router. Tracking is an Alteon WebSystems proprietary extension to VRRP, used for modifying the standard priority system used for electing the master router. Tracking is not needed if sharing (`share`) is enabled. The default value is "nothing." To view menu options, refer to page 7-49.
`vrid` *<virtual router ID (1-255)>* Defines the virtual router ID. This is used in conjunction with `addr` (below) to define a virtual router on this switch. To create a pool of VRRP-enabled routing devices which can provide redundancy to each other, each participating VRRP device must be configured with the same virtual router: one that shares the same `vrid` and `addr` combination. The `vrid` for standard virtual routers (where the virtual router IP address is not the same as any virtual server) can be any integer between 1 and 255. The default value is 1. All `vrid` values must be unique within the VLAN to which the virtual router's IP interface (see `if` below) belongs.
`addr` *<IP address (e.g., 192.4.17.101)>* Defines the IP address for this virtual router using dotted decimal notation. This is used in conjunction with the `vrid` (above) to configure the same virtual router on each participating VRRP device. The default address is 0.0.0.0.
`if` *<interface number (1-256)>* Selects a switch IP interface (between 1 and 256). If the IP interface has the same IP address as the `addr` option above, this switch is considered the "owner" of the defined virtual router. An owner has a special priority of 255 (highest) and will always assume the role of master router, even if it must preempt another virtual router which has assumed master routing authority. This preemption occurs even if the `preem` option below is disabled. The default value is 1.

Table 7-31 VRRP Virtual Router Options (/cfg/vrrp/vr)

Command Syntax and Usage
`prio` *<priority (1-254)>* Defines the election priority bias for this virtual server. This can be any integer between 1 and 254. The default value is 100. During the master router election process, the routing device with the highest virtual router priority number wins. If there is a tie, the device with the highest IP interface address wins. If this virtual router's IP address (`addr`) is the same as the one used by the IP interface, the priority for this virtual router will automatically be set to 255 (highest). When priority tracking is used (`/cfg/vrrp/track` or `/cfg/vrrp/vr` #`/track`), this base priority value can be modified according to a number of performance and operational criteria.
`adver` *<seconds (1-255)>* Defines the time interval between VRRP master advertisements. This can be any integer between 1 and 255 seconds. The default value is 1.
`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). By default, this option is enabled.
`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. By default, this option is enabled.
`ena` Enables this virtual router.
`dis` Disables this virtual router.
`del` Deletes this virtual router from the switch configuration.
`cur` Displays the current configuration information for this virtual router.

/cfg/vrrp/vr *<router number>* **/track**

Virtual Router Priority Tracking Configuration

```
[VRRP Virtual Router 1 Priority Tracking Menu]
      vrs     - Enable/disable tracking 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
      cur     - Display current VRRP virtual router configuration
```

This menu is used for modifying the priority system used when electing the master router from a pool of virtual routers. Various tracking criteria can be used to bias the election results. Each time one of the tracking criteria is met, the priority level for the virtual router is increased by an amount defined through the VRRP Tracking Menu (see page 7-54).

Criteria are tracked dynamically, continuously updating virtual router priority levels when enabled. If the virtual router preemption option (see `preem` in Table 7-31 on page 7-47) is enabled, this virtual router can assume master routing authority when its priority level rises above that of the current master.

Some tracking criteria (`vrs`, `ifs`, and `ports` below) apply to standard virtual routers, otherwise called “virtual interface routers.” Other tracking criteria (`l4pts`, `reals`, and `hsrp`) apply to “virtual server routers,” which perform Layer 4 Server Load Balancing functions. A virtual *server* router is defined as any virtual router whose IP address (`addr`) is the same as any configured virtual server IP address.

Table 7-32 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 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 7-32 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. A port is considered "active" if it has a link and is forwarding traffic. This helps elect the virtual routers with the most available ports as the master.

`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. This helps elect the switch closest to the master HSRP router as the master, optimizing routing efficiency.

`cur`

Displays the current configuration for priority tracking for this virtual router.

`/cfg/vrrp/group` *<router group number>*

Virtual Router Group Configuration

```
[VRRP Virtual Router Group Menu]
      track   - Priority Tracking Menu
      vrid    - Set virtual router ID
      if      - Set interface number
      prio    - Set renter priority
      adver   - Set advertisement interval
      preem   - Enable/disable preemption
      share   - Enable/disable sharing
      ena     - Enable virtual router
      dis     - Disable virtual router
      del     - Delete virtual router
      cur     - Display current VRRP virtual router configuration
```

This menu is used for associating all virtual routers into a single logical virtual router, which forces all virtual routers on the switch to either be master or backup as a group. A virtual router is defined by its virtual router ID and an IP address. On each VRRP-capable routing device participating in redundancy for this virtual router, a virtual router will be configured to share the same virtual router ID and IP address.

NOTE – This option is required to be configured only when using at least two Alteon switches in a hot-standby failover configuration; where only one switch is active at any given time.

Table 7-33 VRRP Virtual Router Group Options (/cfg/vrrp/group)

Command Syntax and Usage

`track`

Displays the VRRP priority tracking menu for the virtual router group. Tracking is an Alteon WebSystems proprietary extension to VRRP, used for modifying the standard priority system used for electing the master router. Tracking is not needed if sharing (`share`) is enabled. The default value is "nothing."
To view menu options, refer to page 7-54.

`vrid` *<virtual router ID (1-255)>*

Defines the virtual router ID.

The `vrid` for standard virtual routers (where the virtual router IP address is not the same as any virtual server) can be any integer between 1 and 255. All `vrid` values but must be unique within the VLAN to which the virtual router's IP interface (see `if` below) belongs.

Table 7-33 VRRP Virtual Router Group Options (/cfg/vrrp/group)

Command Syntax and Usage

`if` *<interface number (1-256)>*

Selects a switch IP interface (between 1 and 256).

`prio` *<priority (1-254)>*

Defines the election priority bias for this virtual router group. This can be any integer between 1 and 254. The default value is 100.

During the master router election process, the routing device with the highest virtual router priority number wins. If there is a tie, the device with the highest IP interface address wins. If this virtual router's IP address (`addr`) is the same as the one used by the IP interface, the priority for this virtual router will automatically be set to 255 (highest).

When priority tracking is used (`/cfg/vrrp/track` or `/cfg/vrrp/vr #/track`), this base priority value can be modified according to a number of performance and operational criteria.

`adver` *<1-255 seconds>*

Defines the time interval between VRRP master advertisements. This can be any integer between 1 and 255 seconds. The default is 1.

`preem disable|enable` (or just **`d`**|**`e`**)

Enables or disables master preemption. When enabled, if the virtual router group is in backup mode but has a higher priority than the current master, this virtual router will preempt the lower priority master and assume control. Note that even when `preem` is disabled, this virtual router will always preempt any other master if this switch is the owner (the IP interface address and virtual router `addr` are the same).

`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. By default, this option is enabled.

`ena`

Enables the virtual router group.

`dis`

Disables the virtual router group.

`del`

Deletes the virtual router group from the switch configuration.

`cur`

Displays the current configuration information for the virtual router group.

/cfg/vrrp/if *<interface number>*
VRRP Interface Configuration

NOTE – The *interface-number* (1 to 256) represents the IP interface on which authentication parameters must be configured.

```
[VRRP Interface 1 Menu]
      auth    - Set authentication types
      passw   - Set plain-text password
      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 7-34 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` *<key>*

Defines a plain-text password up to eight characters long. This password will be added to each VRRP packet transmitted by this interface when password authentication is chosen (see `auth` above).

`del`

Clears the authentication configuration parameters for this IP interface. The IP interface itself is not deleted.

`cur`

Displays the current configuration for this IP interface's authentication parameters.

/cfg/vrrp/track

VRRP Tracking Configuration

```
[VRRP Tracking Menu]
      vrs     - Set priority increment for virtual router tracking
      ifs     - Set priority increment for IP interface tracking
      ports   - Set priority increment for VLAN switch port tracking
      l4pts   - Set priority increment for L4 switch port tracking
      reals   - Set priority increment for L4 real server tracking
      hsrp    - Set priority increment for HSRP tracking
      cur     - Display current VRRP Priority Tracking configuration
```

This menu is used for setting weights for the various criteria used to modify priority levels during the master router election process. Each time one of the tracking criteria is met (see VRRP Virtual Router Priority Tracking Menu on page 7-49), the priority level for the virtual router is increased by an amount defined through this menu..

Table 7-35 VRRP Tracking Options (/cfg/vrrp/track)

Command Syntax and Usage

`vrs` *<0-254>*

Defines the priority increment value (1 through 254) for virtual routers in master mode detected on this switch. The default value is 2.

`ifs` *<0-254>*

Defines the priority increment value (1 through 254) for active IP interfaces detected on this switch. The default value is 2..

`ports` *<0-254>*

Defines the priority increment value (1 through 254) for active ports on the virtual router’s VLAN. The default value is 2.

`l4pts` *<0-254>*

Defines the priority increment value (1 through 254) for physical switch ports with active Layer 4 processing. The default value is 2.

`reals` *<0-254>*

Defines the priority increment value (1 through 254) for healthy real servers behind the virtual server router. The default value is 2.

Table 7-35 VRRP Tracking Options (/cfg/vrrp/track)

Command Syntax and Usage

`hsrp`

Defines the priority increment value (1 through 254) for switch ports with Layer 4 client-only processing that receive HSRP broadcasts. The default value is 10.

`cur`

Displays the current configuration of priority tracking increment values.

NOTE – These priority tracking options only define increment values. These options do not affect the VRRP master router election process until options under the VRRP Virtual Router Priority Tracking Menu (see page 7-49) are enabled.

/cfg/bwm
Bandwidth Management Configuration

```
[Bandwidth Management Menu]
      cont     - Contract Menu
      policy   - Policy Menu
      user     - Set SMTP server user name
      force    - Enable/disable enforce policies
      on       - Globally turn Bandwidth Management processing ON
      off      - Globally turn Bandwidth Management processing OFF
      cur      - Display current Bandwidth Management configuration
```

NOTE – Up to 1024 bandwidth management contracts can be configured on the Alteon AD4 and Alteon 184 Web switches.

Table 7-36 Bandwidth Management Options (/cfg/bwm)

Command Syntax and Usage
`cont` *<contract number (1-256)>* Displays the Bandwidth Management Contract menu. To view menu options, refer to page 7-58.
`policy` *<policy number (1-64)>* Displays the Bandwidth Management Policy menu. To view menu options, refer to page 7-59.
`user` *<user name>* Sets the SMTP user name; that is, the user to whom the history statistics will be mailed.
`force disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables enforces policies. When disabled, no bandwidth limits will be applied on queues. By default, this option is enabled.
`on` Globally enables Bandwidth Management on this switch.
`off` Globally disables Bandwidth Management on this switch.
`cur` Displays the current Bandwidth Management configuration.

/cfg/bwm/cur

Bandwidth Management Current Configuration

```
Current Bandwidth Management setting: ON
  Policy Enforcement: enabled
  SMTP server user name:

Contract Name             Policy Prec Hist TOS State
     4                       4      4   E   D    E
     6                       4      6   E   D    E
   256   Default             --     0   E   D    E
*Default contract gets all the BW that is available on
 a port after the active contracts reserved BW is taken.

Policy     Hard  Soft  Resv oTOS uTOS Buffer
    1     2000k 1000k  500k    0    0   32640
    2       40m   35m   30m    0    0   32640
    3     2000k 1000k  500k    0    0   32640
    4       10m    9m    8m    0    0   32640
    5     2000k 1000k  500k    0    0   32640
    6     2000k 1000k  500k    0    0   32640
    7     2000k 1000k  500k    0    0   32640
    8     2000k 1000k  500k    0    0   32640
    9     2000k 1000k  500k    0    0   32640
   10     2000k 1000k  500k    0    0   32640
   11     2000k 1000k  500k    0    0   32640
   12     2000k 1000k  500k    0    0   32640
   13     2000k 1000k  500k    0    0   32640
   14     2000k 1000k  500k    0    0   32640
   15     2000k 1000k  500k    0    0   32640
   16     2000k 1000k  500k    0    0   32640
   17     2000k 1000k  500k    0    0   32640
   18     2000k 1000k  500k    0    0   32640
   19     2000k 1000k  500k    0    0   32640
   20     2000k 1000k  500k    0    0   32640
   21     2000k 1000k  500k    0    0   32640
   22     2000k 1000k  500k    0    0   32640
   23     2000k 1000k  500k    0    0   32640
   24     2000k 1000k  500k    0    0   32640
   25     2000k 1000k  500k    0    0   32640
   26     2000k 1000k  500k    0    0   32640
   27     2000k 1000k  500k    0    0   32640
   28     2000k 1000k  500k    0    0   32640
   29     2000k 1000k  500k    0    0   32640
   30     2000k 1000k  500k    0    0   3264
```

`/cfg/bwm/cont` *<contract number>*

Bandwidth Management Contract Configuration

```
[BW Contract 1 Menu]
      name    - Set Contract name
      policy  - Set Contract Policy
      prec    - Set Contract Precedence
      history - Enable/disable Saving Contract stats history
      wtos    - Enable/disable overwriting IP TOS for this Contract
      ena     - Enable BW Contract
      dis     - Disable BW Contract
      del     - Delete BW Contract
      cur     - Display current BW Contract configuration
```

Table 7-37 Bandwidth Management Policy Menu Options (/cfg/bwm/cont)

Command Syntax and Usage

`name` *<15 character name>*

Sets the name for this bandwidth contract.

`policy` *<bandwidth policy number (1-64)>*

Sets the policy number for this bandwidth contract.

`prec` *<bandwidth precedence value (1-255)>*

Sets the precedence value for this bandwidth contract.

`history disable|enable` (or just **`d|e`**)

Disables/enables saving statistics for this contract on the server.

`wtos disable|enable` (or just **`d|e`**)

Disables/enables overwriting the IP Type of Service (TOS) for this contract.

`ena`

Enables this bandwidth contract.

`dis`

Disables this bandwidth contract.

`del`

Removes this contract from the switch.

`cur`

Displays the current bandwidth contract configuration.

`/cfg/bwm/pol` *<policy number>*

Bandwidth Management Policy Configuration

```
[Policy 1 Menu]
      hard    - Set hard Limit
      soft    - Set soft Limit
      resv    - Set Reservation Limit
      utos    - Set underlimit (soft limit) TOS
      otos    - Set overlimit (soft limit) TOS
      buffer  - Set Buffer Limit
      cur     - Display current Policy configuration
```

Table 7-38 Bandwidth Management Policy Menu Options (/cfg/bwm/pol/)

Command Syntax and Usage

`hard` *< 250K-5000K|1MB-1000MB>*

Sets the hard bandwidth limit for this policy. This is the highest amount of bandwidth available to this policy. The default value is 2000 kbps.

`soft` *< 250K-5000K|1MB-1000MB>*

Sets the soft bandwidth limit for this policy. The default value is 1000 kbps.

`resv` *< 250K-5000K|1MB-1000MB>*

Sets the reserve limit for this policy. This is the amount of bandwidth always available to this policy. The default value is 500Kbytes.

`utos` *<BW Policy ToS (0-255)>*

Sets the new `utos` value to overwrite the original ToS value if the traffic for this contract is under the soft limit. With this option set to the default value of “0,” the switch will not overwrite the ToS value.

`otos` *<BW Policy ToS (0-255)>*

Sets the new `otos` value to overwrite the original ToS value if the traffic for this contract is over the soft limit. With this option set to the default value of “0,” the switch will not overwrite the ToS value.

`buffer` *<Maximum buffer space (bytes) (8192-512000)>*

Sets the buffer limit for this policy. The default value is 32640 bytes.

`cur`

Displays the current bandwidth policy configuration.

CHAPTER 8
The SLB Configuration Menu

This chapter discusses how to use the command-line interface (CLI) for configuring Server Load Balancing (SLB) on the switch.

/cfg/slb
SLB Configuration

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

Table 8-1 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage
`real` *<real server number (1-256)>* Displays the menu for configuring real servers. To view menu options, refer to page 8-3.
`group` *<real server group number (1-256)>* Displays the menu for placing real servers into real server groups. To view menu options, refer to page 8-7.
`virt` *<virtual server number (1-256)>* Displays the menu for defining virtual servers. To view menu options, refer to page 8-12.

Table 8-1 Server Load Balancing Configuration Menu Options (/cfg/slb)

Command Syntax and Usage

`filt` *<filter ID (1-224)>*

Displays the menu for Filtering and Application Redirection. To view menu options, refer to page 8-20.

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

Displays the menu for setting physical switch port states for Layer 4 activity. To view menu options, refer to page 8-29.

`gslb`

Displays the menu for configuring Global Server Load Balancing. To view menu options, refer to page 8-31.

`url`

URL Redirection and Load Balance Menu. To view menu options, refer to page 8-38.

`sync`

Displays the Synch Peer Switch Menu. To view menu options, refer to page 8-41.

`adv`

Displays the Layer 4 Advanced Menu. To view menu options, refer to page 8-43.

`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 9-9). Enabling Layer 4 services is not necessary for using filters only to allow, deny, or NAT traffic (see note below).

`off`

Globally disables Layer 4 services. All configuration information will remain in place (if `applied` or `saved`), but the software processes will no longer be active in the switch

`cur`

Displays the current Server Load Balancing configuration.

Filtering and Layer 4

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>*
Real Server SLB Configuration

NOTE – The real-server-number (1 to 255) 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
      tmout   - Set minutes inactive connection remains open
      backup  - Set backup real server
      inter   - Set interval between health checks
      retry   - Set number of failed attempts to declare server DOWN
      restr   - Set number of successful attempts to declare servr UP
      addlb   - Add URL path for URL load balance
      remlb   - Remove URL path for URL load balance
      remote  - Enable/disable remote site operation
      proxy   - Enable/disable client proxy operation
      submac  - Enable/disable source MAC address substitution
      nocook  - Enable/disable no available URL cookie operation
      exclude - Enable/disable exclusionary string matching
      ena     - Enable real server
      dis     - Disable real server
      del     - Delete real server
      cur     - Display current real server configuration
```

This menu is used for configuring information about 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 8-2 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 `PING`ed to determine if the server is up, and the administrator will be warned if the server does not respond.

Table 8-2 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage

`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.

`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 8-10).

`maxcon` *<maximum connections (0-200000)>*

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 sent to the backup/overflow server or backup/overflow server group. If no backup servers/server group are configured, client requests will be dropped by the virtual server.

`tmout` *<even number of minutes (2-30)>*

Sets the number of minutes an inactive session remains open (in even numbered increments).

Every client-to-server session being load balanced is recorded in the switch's Session Table. When a client makes a request, the session is recorded in the table, the data is transferred until the client ends the session, and the session table entry is then removed.

In certain circumstances, such as when a client application is abnormally terminated by the client's system, TCP/UDP connections will remain registered in the switch's binding table. In order to prevent table overflow, these orphaned entries must be aged out of the binding table.

Using the tmout option, you can set the number of minutes to wait before removing orphan table entries. Settings must be specified in even numbered increments between 2 and 30 minutes. The default setting is 10.

This option is also used with the Persistent option (see /cfg/slb/virt/pbind). When persistent is activated, this option sets how long an idle client is allowed to remain associated with a particular server.

Table 8-2 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage

`backup` *<real server number (1-256)>*|**none**

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

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

The backup server is also used in overflow situations. If the real server reaches its `maxcon` (maximum connections) limit, the backup comes online to provide additional processing power until the original server becomes desaturated.

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

`inter` *<number of seconds between health checks (0-60)>*

Sets the interval between real server health verification attempts.

Determining the health of each real server is a necessary function for Layer 4 switching. For TCP services, the switch verifies that real servers and their corresponding services are operational by opening a TCP connection to each service, using the defined service ports configured as part of each virtual service. For UDP services, the switch pings servers to determine their status.

The `intr` option lets you choose the time between health checks. The range is from 1 to 60 seconds. The default interval is 2 seconds. An interval of “0” disables health checking for the server.

`retry` *<number of consecutive health checks (1-63)>*

Sets the number of failed health check attempts required before declaring this real server inoperative. The range is from 1 to 63 attempts. The default is 4 attempts

`restr` *<number of consecutive health checks (1-63)>*

Sets the number of successful health check attempts required before declaring a UDP service operational. The range is from 1 to 63 attempts. The default is 8 attempts

`addlb`

Add URL path for URL load balance.

`remlb`

Remove URL path for URL load balance.

`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.

Table 8-2 Real Server Configuration Menu Options (/cfg/slb/real)

Command Syntax and Usage

`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).`submac  - Enable/disable source MAC address substitution`

`submac disable`|**`enable`** (or just **`d`**|**`e`**)

Enables or disables source MAC address substitution.

`nocook disable`|**`enable`** (or just **`d`**|**`e`**)

Enables or disables no available URL cookie operation.

`exclude disable`|**`enable`** (or just **`d`**|**`e`**)

Enables or disables exclusionary string matching.

`enable`

You *must* perform this command to enable this real server for Layer 4 service. When enabled, the real server can process virtual server requests associated with its real server group. This option, when the `apply` and `save` commands are used, enables this real server for operation until explicitly disabled.

See `/oper/slb/ena` on page 9-6 for an operations-level command

`dis`

Disables this real server from Layer 4 service. Any disabled server will no longer process virtual server requests as part of the real server group to which it is assigned. This option, when the `apply` is are used, disables this real server until it is explicitly re-enabled. This option *does not* perform a graceful server shutdown.

See `/oper/slb/dis` on page 9-6 for an operations-level command.

`del`

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

`cur`

Displays the current configuration information for this real server.

`/cfg/slb/group` *<group number>* Real Server Group SLB Configuration

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

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

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

Table 8-3 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
`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 8-10.
`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 `healt` option (see below).

Table 8-3 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage

`health icmp|tcp|http|dns|pop3|smtp|nntp|ftp|imap|radius|sslh |script<n>`

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

- `icmp` For Layer 3 health checking, `ping` the server.
- `tcp` For TCP service, open and close a TCP/IP connection to the server.
- `http` For HTTP service, uses HTTP 1.1 `GETS` when a `HOST:` header is required to check that the URL content specified in `content` is accessible on the server. Otherwise, an `HTTP/1.0 GET` occurs.
- `dns` For Domain Name Service, check that the domain name specified in `content` can be resolved by the server.
- `pop3` For user mail service, check that the *user*:*password* account specified in `content` exists on the server.
- `smtp` For mail-server services, check that the user specified in `content` is accessible on the server.
- `nntp` For newsgroup services, check that the newsgroup name specified in `content` is accessible on the server.
- `ftp` For FTP services, check that the filename specified in `content` is accessible on the server through anonymous login.
- `imap` For user mail service, check that the *user:password* value specified in `content` exists on the serve
- `radius` For RADIUS remote access server authentication, check that the *user:password* value specified in `content` exists on the switch and the server. To perform application health checking to a RADIUS server, the network administrator must also configure the `/cfg/slb/secrt` parameter. The `secrt` value is a field of up to 32 alphanumeric characters that is used by the switch to encrypt a password during the RSA Message Digest Algorithm (MD5) and by the RADIUS server to decrypt the password during verification.
- `sslh` Enables the switch to query the health of the SSL servers by sending an SSL client "Hello" packet and then verify the contents of the server's "Hello" response. During the handshake, the user and server exchange security certificates, negotiate an encryption and compression method, and establish a session ID for each session.
- `script` Enables the use of script-based health checks in send/expect format to check for application and content availability.

Table 8-3 Real Server Group Configuration Menu Options (/cfg/slb/group)

Command Syntax and Usage
`backup r`<real server number *(1-256)*>|**g**<group number>|**none** Sets the real server or real server group used as the backup/overflow server/server group for this real server group. To prevent loss of service if the entire real server group fails, use this option to assign a backup real server/real server group number. Then, if the real server group becomes inoperative, the switch will activate the backup real server /server group until one of the original real servers becomes operative again. The backup server/server group is also used in overflow situations. If all the servers in the real server group reach their `maxcon` (maximum connections) limit, the backup server/server group comes online to provide additional processing power until one of the original servers becomes desaturated. The same backup/overflow server/server group may be assigned to more than one real server group at the same time.
`name` *<string, maximum 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` Sets real server failure threshold.
`add` *<real server number (1-256)>* Adds a real server to this real server group. You will be prompted to enter the number (1 to 256) of the real server to add to this group.
`rem` *<real server number (1-256)>* Remove a real server from this real server group. You will be prompted for the ID number for the real server to remove from this group.
`del` Deletes this real server group from the Layer 4 software configuration. This removes the group from operation under all virtual servers it is assigned to. Use this command with caution: if you remove the only group assigned to a virtual server, the virtual server will become inoperative.
`cur` Displays the current configuration parameters for this real server group.

Server Load Balancing Metrics

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

Table 8-4 Real Server Group Metrics

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

Table 8-4 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 8-3), a higher proportion of connections are given to servers with higher weights. This can improve load balancing among servers of different performance levels. Weights are not applied when using `hash` or `minmisses`.

`/cfg/slb/virt` *<server number>* Virtual Server SLB Configuration

NOTE – The *virtual-server-number* (1 to 256) represents the number of the virtual server that you wish to configure.

```
[Virtual Server 1 Menu]
      service - Virtual Service Menu
      vip     - Set IP addr of virtual server
      dname   - Set domain name of virtual server
      cont    - Set BW Contract
      layr3   - Enable/disable layer 3 only balancing
      ftpp    - Enable/disable FTP SLB parsing for virtual server
      ena     - Enable virtual server
      dis     - Disable virtual server
      del     - Delete virtual server
      cur     - Display current virtual configuration
```

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

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

Table 8-5 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
`service` *<virtual port or name, from 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 menu options, refer to page 8-14.
`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 `ARP`s and `PING`s from network ports as if it was a normal server. Client requests directed to the virtual server's IP address will be balanced among the real servers available to it through real server group assignments.

Table 8-5 Virtual Server Configuration Menu Options (/cfg/slb/virt)

Command Syntax and Usage
`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, etcetera). To define the hostname, see `hname` below. To clear the `dname`, specify the name as **`none`**.
`cont <`*BW contract (1-256)***`>`** Enter a new Bandwidth Contract for this virtual server.
`layr3 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.
`ftpp disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables FTP SLB parsing for this virtual server. When this option is enabled, the switch modifies the appropriate FTP method/command to support FTP servers on a private network for both active and passive FTP modes. To do this, the switch looks deeper into the packet and modifies the **`PORT`** command for active FTP or the “entering the passive mode” command for passive FTP.
`ena` Enables this virtual server. This option activates the virtual server within the switch so that it can service client requests sent to its defined IP address.
`dis` This option disables the virtual server so that it no longer services client requests.
`del` This command removes this virtual server from operation within the switch and deletes it from the Layer 4 switching software configuration. Use this command with caution, as it will delete the options that have been set for this virtual server.
`cur` Displays the current configuration of the specified virtual server.

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

Virtual Server Service Configuration

```
[Virtual Server 1 http Service Menu]
      group   - Set real server group number
      rport   - Set real port
      hname   - Set hostname
      httpslb - Set HTTP SLB processing
      cont    - Set BW contract for this virtual service
      pbind   - Set persistent binding type
      udp     - Enable/disable UDP balancing
      frag    - Enable/disable remapping UDP server fragments
      nonat   - Enable/disable only substituting MAC addresses
      del     - Delete virtual service
      cur     - Display current virtual service configuration
```

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

Table 8-6 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage
`group` *<real server group number>* Sets a real server group for this service. You will be prompted to enter the number (1 to 256) of the real server group to add to this service.
`rport` *<real server port (0-65534)>* Defines the real server TCP or UDP port assigned to this service. By default, this is the same as the virtual port (service virtual port). If `rport` is configured to be different than the virtual port defined in **`/cfg/slb/virt/service`** *<virtual port>*, the switch will map the virtual port to this real port.
`hname` *<hostname>*|**none** Sets the hostname for a service added. This is used in conjunction with dname (above) to create a full host/domain name for individual services. The format for this command is as follows: # **`hname`** *<hostname>* For example, to add a hostname for Web services, you could specify “www” as the hostname. If a `dname` of “foocorp.com” was defined (above), “www.foocorp.com” would be the full host/domain name for the service. To clear the hostname for a service, use the following command: # **`hname none`**

Table 8-6 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage

`httpslb disable|enable` (or just **`d|e`**)

Enables/disables HTTP-based server load balancing. Once enabled, you can proceed to enable one of the following application options and their applicable parameters:

`urlslb|host|cookie|browser|urlhash|others`

- `urlslb`: Enable/disable URL SLB
- `host`: Enable/disable for virtual hosting
- `cookie`: Enable/disable cookie-based SLB

  For cookie-based preferential load balancing. You will be prompted for the following: Cookie name, starting point of the cookie value, number of bytes to be extracted, enable/disable checking for cookie in URI.
- `browser`: Enable/disable SLB, based on browser type
- `urlhash`: Enable/disable URL hashing based on URI
- `others`: Requires inputs for a particular header field

`cont` *<BW Contract (0-256)>*

Sets a Bandwidth Contract for this virtual service.

`pbind clientip|cookie|sslid|disable`

Enable/disable persistent bindings for a real server. This may be necessary for some server applications where state information about the client system is retained on the server over a series of sequential connections, such as with SSL (Secure Socket Layer, HTTPS), website search results, or multi-page Web forms.

- The `clientip` option uses the client IP address as an identifier, and associates all connections from the same client with the same real server until the client becomes inactive and the connection is aged out of the binding table. The connection timeout value (set in the Real Server Menu) is used to control how long these inactive but persistent connections remain associated with their real servers. When the client resumes activity *after* their connection has been aged out, they will be connected to the most appropriate real server based on the load balancing metric.
  An alternative approach may be to use the real server group metrics `minmisses` or `hash` (see "Server Load Balancing Metrics" on page 8-10).
- The `cookie` option uses a cookie defined in the HTTP header or placed in the URI for hashing. A *permanent cookie* gets stored on the client's browser, as part of the response from a site's server. It will be sent by the browser when the client makes subsequent requests to the same site, even after the browser has been shut down. A *temporary cookie* is only valid for the browser session. The temporary cookie expires when the browser is closed.
- Secure Sockets Layer (SSL) is a set of protocols built on top of TCP/IP that allow an application server and user to communicate over an encrypted HTTP session, providing authentication, non-repudiation, and security. The session ID is a value comprising 32 random bytes chosen by the SSL server that gets stored in a session hash table. By enabling the sslid option, all subsequent SSL sessions which present the same session ID will be directed to the same real server.
- The `disable` option enables you to disable presistent binding, if it has previously been enabled for a particular application.

Table 8-6 Virtual Server Service Configuration Options (/cfg/slb/virt/service)

Command Syntax and Usage
`udp disable`|`enable`|`stateless` (or just **`d`**|**`e`**|**`s`**) Enable/disable 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/disables remapping server fragments for virtual port.
`nonat disable`|`enable` (or just **`d`**|**`e`**) Enables/disables only substituting MAC address of the real server. This option does not substitute IP addresses.
`del` This command removes this virtual service from operation within the switch and deletes it from the Layer 4 switching software configuration. Use this command with caution, as it will delete the options that have been set for this virtual service.
`cur` Displays the current configuration of services on the specified virtual server.

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 8-29).

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

Port Mapping

NOTE – Layer 4 port mapping is not supported if Direct Access Mode is enabled on a server.

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.

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.

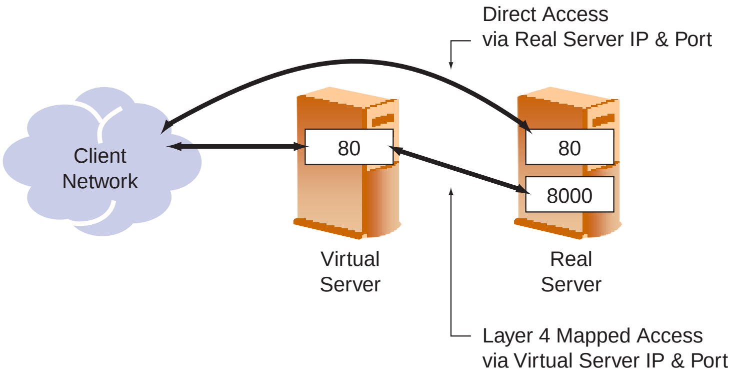

Figure 8-1 Mapped and Non-Mapped server access

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.

The format for the mapping command is as follows:

```
Virtual server 1# service virtual-server-port|real-server-port
```

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

`/cfg/slb/filt` *<filter number>* SLB Filter Configuration

```
[Filter 1   Menu]
      adv     - Filter Advanced 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
      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
      inver   - Enable/disable filter inversion
      ena     - Enable filter
      dis     - Disable filter
      del     - Delete filter
      cur     - Display current filter configuration
```

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

The required 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 8-7 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
`adv` Displays the Filter Advanced Menu. To view menu options, refer to page 8-25.
`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
89	ospf
112	vrrp

Table 8-7 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage

`sport any`|*<name>*|*<port>*|*<port>*-*<port>*

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
20	ftp-data
21	ftp
22	ssh
23	telnet
25	smtp
37	time
42	name
43	whois
53	domain
69	tftp
70	gopher
79	finger
80	http
109	pop2
110	pop3
111	sunrpc
119	nntp
123	ntp
143	imap
144	news
161	snmp
162	snmptrap
179	bgp
194	irc
220	imap3
389	ldap
443	https
520	rip
554	rtsp
1985	hsrp

`dport any`|*<name>*|*<port>*|*<port>*-*<port>*

If defined, traffic with the specified real server TCP or UDP destination port will be affected by this filter. The port number, range, name, or "**any**" can be specified, just as with `sport` above.

Table 8-7 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
`action` Specify the action this filter takes: `allow` Allow the frame to pass. `deny` Discard frames that fit this filter's profile. This can be used for building basic security profiles. `redir` Redirect frames that fit this filter's profile, such as for web-cache redirection. In addition, Layer 4 processing must be activated (see the `/cfg/slb/on` command on page 8-1). `nat` Perform generic Network Address Translation (NAT). This can be used to map the source or destination IP address and port information of a private network scheme to/from the advertised network IP address and ports. This is used in conjunction with the `nat` option below and can also be combined with proxies.
`group` *<real server group number (1-256)>* This option applies only when `redir` is specified at the filter action. Define a real server group (1 to 256) to which redirected traffic will be sent.
`rport` *<real server port (0-65535)>* This option applies only when `redir` is specified at the filter action. This defines the real server TCP or UDP port to which redirected traffic will be sent. For valid Layer 4 health checks, this must be configured whenever TCP protocol traffic is redirected. Also, if transparent proxies are used for Network Address Translation (NAT) on the switch (see the `pip` option in Table 8-12 on page 8-29), `rport` must be configured for all Application Redirection filters.
`nat source`|`dest` When `nat` is set as the filter action (see above), this command specifies whether the source or the destination information is re-mapped. If **`source`** is specified, the frame's source IP address (`sip`) and port number (`sport`) are replaced with the `dip` and `dport` values. If **`dest`** is specified, the frame's destination IP address (`dip`) and port number (`dport`) are replaced with the `sip` and `sport` values.
`inver disable`|`enable` (or just **`d`**|**`e`**) Inverts the filter logic. If the conditions of the filter are met, *don't* act. If the conditions for the filter are *not met,* perform the assigned action.
`ena` Enables this filter.

Table 8-7 Filter Configuration Menu Options (/cfg/slb/filt)

Command Syntax and Usage
`dis` Disables this filter.
`del` Deletes this filter.
`cur` Displays the current filter configuration.

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 8-8 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/adv

Advanced Filter Configuration

```
[Filter 1 Advanced Menu]
      tcp     - TCP Flags Advanced Menu
      tos     - Set IP Type of Service
      tmask   - Set IP TOS mask
      newtos  - Set new IP TOS
      option  - Enable/disable IP option matching
      icmp    - Set ICMP message type
      cont    - Set BW contract
      proxy   - Enable/disable client proxy
      cache   - Enable/disable caching sessions that match filter
      log     - Enable/disable logging
      ack     - Enable/disable TCP ACK or RST matching
      fwlb    - Enable/disable firewall redirect hash method
      urlp    - Enable/disable URL parsing
      ftpa    - Enable/disable active FTP NAT
      cur     - Display current advanced filter configuration
```

Table 8-9 Advanced Filter Menu (/cfg/slb/filt/adv)

Command Syntax and Usage

`tcp`

Displays the TCP Flags Advanced Menu. To view menu options, refer to page 8-27.

`tos` *<0-255>*

Sets the IP Type of Service. This option is used to match ToS in frames.

`tmask` *<0-255>*

Sets the IP Type of Service mask. This option is used to match ToS in frames.

`newtos` *<0-255>*

Sets the new IP Type of Service for allow filters. A value of “0” means that the ToS does not change.

`option disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables IP option matching.

`icmp any`|*<number>*|**`name`**

Sets the ICMP message type. For a list of ICMP message types, refer to Table 8-11 on page 8-28.

Table 8-9 Advanced Filter Menu (/cfg/slb/filt/adv)

Command Syntax and Usage

`proxy disable|enable` (or just **`d|e`**)

Enables/disables client proxy. This option applies only when `redir` or `nat` is specified as the filter action. Enable or disable proxy IP address translation for traffic matching the filter criteria. By default, this is enabled. If disabled, any proxy defined for the switch port using the `pip` command (see page 8-29) is not performed for traffic meeting the filter criteria. This is useful when certain traffic must retain original IP address information, or when other forms of translation (such as Application Redirection or NAT) are preferred.

`cache disable|enable` (or just **`d|e`**)

Enables/disables caching sessions that match filter.

`cont` *<BW Contract (1-256)>*

Sets the Bandwidth Contract.

`log disable|enable` (or just **`d|e`**)

Enables/disables logging.

`ack disable|enable` (or just **`d|e`**)

Enable/disable 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 source 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.

If cache is disabled, it will filter out ona per-packet basis. If the cache is enabled, then filtering is performed on a per-session basis.

`fwlb disable|enable` (or just **`d|e`**)

To ensure that the "stateful inspection" behavior of firewalls is maintained, a hashing algorithm is used to ensure that inbound packets and outbound packets for a pair of IPSA/IPDA traverse through the same firewall. If the dport is 80 or 21, enabling this option changes the hash of the filter from a WCR hash to a FWLB hash.

`urlp disable|enable` (or just **`d|e`**)

Enables/disables URL parsing.

`ftpa disable|enable` (or just **`d|e`**)

Enables/disables active FTP Client NAT (Network Address Translation). When a client in active FTP mode sends a **`PORT`** command to a remote FTP server, the switch will look into the data part of the frame and and replace the client 's private IP address with a proxy IP (`PIP`) address. The real server port (`RPORT`) will be replaced with a proxy port (`PPORT`), that is `PIP`:`PPORT`.

`cur`

Displays the current advanced filter configuration.

/cfg/slb/filt/adv/tcp

Advanced Filter TCP Configuration

```
[TCP flags advanced Menu]
      urg     - Enable/disable TCP URG flag matching
      ack     - Enable/disable TCP ACK flag matching
      psh     - Enable/disable TCP PSH flag matching
      rst     - Enable/disable TCP RST flag matching
      syn     - Enable/disable TCP SYN flag matching
      fin     - Enable/disable TCP FIN flag matching
      cur     - Display current ACL TCP filter configuration
```

Table 8-10 Advanced Filter TCP Menu (/cfg/slb/filt/adv/tcp)

Command Syntax and Usage
`urg disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables TCP URG flag matching.
`ack disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables TCP ACK flag matching.
`psh disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables TCP PSH flag matching.
`rst disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables TCP RST flag matching.
`syn disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables TCP SYN flag matching.
`fin disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables TCP FIN flag matching.
`cur` Displays the current Access Control List TCP filter configuration.

ICMP Message Types

The following ICMP message types are used with the `/cfg/slb/filt/adv/icmp` command:

Table 8-11 ICMP Message Types

Type #	Message Type	Description
0	echorep	ICMP echo reply
3	destun	ICMP destination unreachable
4	quench	ICMP source quench
5	redir	ICMP redirect
8	echoreq	ICMP echo request
9	rtradv	ICMP router advertisement
10	rtrsol	ICMP router solicitation
11	timex	ICMP time exceeded
12	param	ICMP parameter problem
13	timereq	ICMP timestamp request
14	timerep	ICMP timestamp reply
15	inforeq	ICMP information request
16	inforep	ICMP information reply
17	maskreq	ICMP address mask request
18	maskrep	ICMP address mask reply

`/cfg/slb/port` *<port number>*
Port SLB Configuration

```
[SLB port 1 Menu]
      client   - Enable/disable client processing
      server   - Enable/disable server processing
      hotstan  - Enable/disable hot-standby processing
      intersw  - Enable/disable inter-switch processing
      proxy    - Enable/disable use of PIP for ingress traffic
      pip      - Set Proxy IP address for port
      filt     - Enable/disable filtering
      add      - Add filter to port
      rem      - Remove filter from port
      cur      - Display current port configuration
```

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 8-12 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.
`hotstan disable|enable` (or just **d**|**e**) Enables/disables hot-standby processing. Use this option and the intersw option in conjunction with VRRP hot-standby failover.

Table 8-12 Port Configuration Menu Options (/cfg/slb/port)

Command Syntax and Usage
`intersw disable|enable` (or just **`d|e`**) Enables/disables inter-switch processing. This option is enabled for ports connected to a peer switch.
`proxy disable|enable` (or just **`d|e`**) Enables or disables a proxy on this port.. When the PIP is defined, client address information in Layer 4 requests is replaced with this proxy IP address. In Server Load Balancing applications, this forces response traffic to return through the switch, rather than around it, as is possible in complex routing environments. Proxies are also useful for Application Redirection and Network Address Translation (NAT). When `pip` is used with Application Redirection filters, each filter's `rport` parameter must also be defined (see `rport` on page 8-21).
`pip` *<proxy IP address>* Sets the proxy IP address for this port, using dotted decimal notation. When the PIP is defined, client address information in Layer 4 requests is replaced with this proxy IP address.
`filter disable|enable` (or just **`d|e`**) Enables or disables filtering on this port.
`add` *<filter ID (1-224)>* Adds a filter for use on this port.
`rem` *<filter ID (1-224)>* Removes a filter from use on this port.
`cur` Displays current system parameters.

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 9-4 on page 9-6).

/cfg/slb/gslb
Global SLB Configuration

```
[Global SLB Menu]
   site     - Remote Site Menu
   lookup   - Network Preference Lookup Menu
   ttl      - Set Time To Live of DNS resource records
   mincon   - Set minimum number of site connections
   inter    - Set interval between remote site updates
   weight   - Set local weight
   dns      - Enable/disable DNS handoffs
   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
   on       - Globally turn Global SLB ON
   off      - Globally turn Global SLB OFF
   cur      - Display current Global SLB configuration
```

NOTE – The `local`, `one`, `alway`, and `geo` options have no effect on lookup.

Table 8-13 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage

`site` *<remote site (1-64)>*

Displays the Remote Site Menu for one of up to 64 remote sites. To view menu options, refer to page 8-34.

`lookup`

Displays the Global SLB Lookup Menu. The options in this menu will overwrite the geographic awareness (IANA table) during DNS queries. To view menu options, refer to page 8-35.

`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 8-13 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.

`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.

`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 number of available connections equal to the total from each server's configured `maxcons` value, minus the server's current number of connections. When the real servers for the local virtual 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/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/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.

Table 8-13 Global SLB Menu Options (/cfg/slb/gslb)

Command Syntax and Usage

`geo disable|enable` (or just **`d|e`**)

Enables/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/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>
```

`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 9-9).

`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.

`cur`

Displays current Global SLB configuration.

/cfg/slb/gslb/site *<site number>*

GSLB Remote Site Configuration

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

Up to 64 remote sites can be configured.

Table 8-14 GSLB Remote Site Menu Options (/cfg/slb/gslb/site)

Command Syntax and Usage

`prima` *<server IP address>*

Defines the IP interface IP address of the primary switch at the remote site used for Global Server Load Balancing. Use dotted decimal notation.

`secon` *<server IP address>*

If the remote site is configured with a redundant switch, enter the IP address of the IP interface for the remote secondary switch here. If the remote site primary switch fails, the local switch will address the remote site secondary switch instead.

`name` *<string, maximum 15 characters>*

Sets the name of the remote site.

`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.

`ena`

Enables this remote site for use with Global Server Load Balancing.

`dis`

Disables this remote site. The switch will no longer use this remote site for Global Server Load Balancing.

Table 8-14 GSLB Remote Site Menu Options (/cfg/slb/gslb/site)

Command Syntax and Usage

`del`

Removes this remote site from operation and deletes its configuration.

`cur`

Displays the current remote site configuration.

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 7-2 on page 7-6).

/cfg/slb/gslb/lookup

GSLB Lookup Configuration

```
[Global SLB Lookup Menu]
      network - Internet Network Preference Menu
      dname   - Set domain name for internal lookup table
      lookups - Enable/disable network preference lookups
      cur     - Display current lookup configuration
```

Table 8-15 GSLB Lookup Menu Options (/cfg/slb/gslb/lookup)

Command Syntax and Usage

`network` *<preference number (1-128)>*

Displays the Internet Network Preference Menu. If enabled, the switch responds to DNS requests based on the configured `dname` and Internet Preference Menu option settings. To view menu options, refer to page 8-36.

`dname` *<domain name>*|**none**

Sets the domain name for the internal lookup table.

`lookups disable`|**`enable`** (or just **d**|**e**)

Enables or disables network preference lookups.

`cur`

Displays the current lookup configuration.

/cfg/slb/gslb/lookup/network *<preference number>*

GSLB Internet Network Preference Lookups Configuration

```
[Network 1 Menu]
      sip     - Set Source IP address
      mask    - Set net mask
      vip1    - Set VIP address
      vip2    - Set VIP address
      del     - Delete internet network entry
      cur     - Display current internet network entry configuration
```

You can overwrite the IANA table by defining client networks, using the options in this menu. You should use regular GSLB to respond to a DNS request under the following conditions:

- Queried domain is not matched.
- Client IP address doesn’t match address in the Network Preference Menu and no default entry is configured.
- There is an entry match in the Network Preference Menu; however, VIP1 and VIP2 are not healthy, that is they are down or over the `mincon`.

The *default entry* is one where the source IP address and mask are not configured (both are 0.0.0.0) and only the VIP1 and VIP 2 are configured. All client networks not in the Network Preference Menu will use this entry to respond to a DNS request.

Table 8-16 GSLB Internet Network Preference Menu Options (/cfg/slb/gslb/lookup/network)

Command Syntax and Usage

`sip` *<IP address>*

Sets the source IP address. Specify an IP address in dotted decimal notation, or “**`any`**”. A range of IP addresses is produced when used with the `mask` option.

`mask` *<IP address>*

This IP address mask is used with the `sip` to find a correct virtual IP address to respond to a DNS request.

`vip1` *<IP address>*

Sets the first virtual IP address. The VIP can either be a local or remote virtual server. The switch returns the VIP with the least response time that is over the `mincon` (minimum number of available connections).

Table 8-16 GSLB Internet Network Preference Menu Options (/cfg/slb/gslb/lookup/network)

Command Syntax and Usage
`vip2` *<IP address>* Sets the second VIP address.
`del` Deletes the specified network entry.
`cur` Displays the current Internet network entry configuration.

/cfg/slb/url
URL Resource Definition

```
[URL Resource Definition Menu]
    redir    - Web Cache Redirection Menu
    lb       - Server Load Balancing Menu
```

Table 8-17 URL Resource Definition Menu Options (/cfg/slb/url)

Command Syntax and Usage
`redir` Displays the Web Cache Redirection Menu. To view menu options, refer to page 8-38.
`lb` Displays the Server Load Balancing Menu. To view menu options, refer to page 8-40.

/cfg/slb/url/redir
Web Cache Redirection Configuration

```
[Web Cache Redirection Menu]
    add      - Add URL expression
    rem      - Remove URL expression
    urlal    - Enable/disable auto-ALLOW for non-GETs to origin servers
    cookie   - Enable/disable auto-ALLOW for Cookie to origin servers
    nocache  - Enable/disable no-cache control header to origin servers
    hash     - Enable/disable URL hashing based on URI
    header   - Enable/disable server loadbalance based on HTTP header
    cur      - Display current URL expression table
```

Table 8-18 Web Cache Redirection Menu Options (/cfg/slb/url/redir)

Command Syntax and Usage
`add` *<string>* Adds the URL expression.
`rem` *<string>* Removes the URL expression.

Table 8-18 Web Cache Redirection Menu Options (/cfg/slb/url/redir)

Command Syntax and Usage

`urlal disable|enable` (or just **d**|**e**)

Enables/disables auto-ALLOW for non-GETs to origin servers.

- If this command is enabled, the switch will redirect all non-GET requests to the origin server.
- If this command is disabled, the switch will compare the URI against the expression table to determine whether all non-GET requests should be redirected to a cache server or origin server.

`cookie disable|enable` (or just **d**|**e**)

Enables/disables auto-ALLOW for cookie to origin servers.

- If this command is enabled, the switch will redirect all requests that contain "Cookie:" in the HTTP header to the origin server.
- If this command is disabled, the switch will compare the URI against the expression table to determine whether it should redirect all requests that contain "Cookie:" in the HTTP header to a cache server or origin server.

`nocache disable|enable` (or just **d**|**e**)

Enables/disables no-cache control header to origin servers.

- If this command is enabled, the switch will redirect all requests that contain "Cache-Control: no-cache" in HTTP/1.1 header or "Pragma: no-cache" in HTTP/1.0 header to the origin server.
- If this command is disabled, the switch will compare the URI against the expression table to determine whether it should redirect requests that contain "Cache-Control: no-cache" in HTTP/1.1 header or "Pragma: no-cache" in HTTP/1.0 header to a cache server or origin server.

`hash disable|enable` (or just **d**|**e**)

Enables/disables URL hashing based on the URI.

- If hashing is enabled, you can set the length of URI that will be used to hash into the cache server.
- If hashing is disabled, the switch will only use the host header field to calculate the hash key.

`header disable|enable` (or just **d**|**e**)

Enables/disables server load balancing based on HTTP header.

`cur`

Displays the current URL expression table.

/cfg/slb/url/lb

Server Load Balance Resource Configuration

```
[Server Loadbalance Resource Menu]
      message - Set error message
      add     - Add URL path for load balance
      rem     - Remove URL path for load balance
      cur     - Display current URL paths
```

Table 8-19 URL Cache Redirection Menu Options (/cfg/slb/url/lb)

Command Syntax and Usage
`message` *<64 byte error message>* Sets an error message.
`add` *<URL path string>* Adds the URL path for load balancing.
`rem` *<URL path ID>* Removes the URL path for load balancing.
`cur` Displays the current URL paths.

/cfg/slb/sync
Synchronize Peer Switch Configuration

```
[Config Synchronization Menu]
      peer     - Synch peer switch menu
      filt     - Enable/disable syncing filter configuration
      ports    - Enable/disable syncing port cofiguration
      prios    - Enable/disable syncing VRRP priorities
      pips     - Enable/disable syncing proxy IP addresses
      cur      - Display current Layer 4 sync configuration
```

To synchronize the configuration between two switches, a peer must be configured on each switch. Switches being synchronized must use the same administrator password. Peers are sent SLB, FILT, and VRRP configuration updates using **`/oper/slb/synch`**.

Table 8-20 Synchonization Menu Options (/cfg/slb/sync)

Command Syntax and Usage

`peer` *<peer switch number>*

Displays the Sync Peer Switch Menu. To view menu options, refer to page 8-42.

`filt disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables syncing filter configuration.

`ports disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables syncing Layer 4 port cofiguration.

`prios disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables syncing VRRP priorities.

`pips disable`|**`enable`** (or just **`d`**|**`e`**)

Enables/disables syncing proxy IP addresses.

`cur`

Displays the current Layer 4 synchronization configuration.

/cfg/slb/sync/peer

Peer Switch Configuration

```
[Peer Switch 1 Menu]
      addr    - Set peer switch IP address
      ena     - Enable peer switch
      dis     - Disable peer switch
      del     - Delete peer switch
      cur     - Display current peer switch configuration
```

To synchronize the configuration between two switches, a peer must be configured on each switch. Switches being synchronized must use the same administrator password.

Table 8-21 Synch Peer Switch Menu Options (/cfg/slb/sync/peer)

Command Syntax and Usage

`addr` *<IP address>*

Sets the peer switch IP address.

`ena`

Enables the peer for this switch.

`dis`

Disables the peer for this switch.

`del`

Deletes the peer for this switch

`cur`

Displays the current peer switch configuration.

/cfg/slb/adv
Advanced Layer 4 Configuration

```
[Layer 4 Advanced Menu]
      script   - Scriptable Health Check Menu
      imask    - Set virtual and real IP address mask
      mnet     - Set managment network
      mmask    - Set management subnet mask
      pmask    - Set persistent mask
      secret   - Set RADIUS secret
      direct   - Enable/disable Direct Access Mode
      grace    - Enable/disable graceful real server failure
      matrix   - Enable/disable Virtual Matrix Architecture
      cur      - Display current Layer 4 advanced configuration
```

Table 8-22 Layer 4 Advanced Menu Options (/cfg/slb/adv)

Command Syntax and Usage
`script` Displays the Scriptable Health Check Menu. To view menu options, refer to page 8-45.
`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 8-46.
`mnet` *<IP address>* If defined, management traffic with this source IP address will be allowed direct (non-Layer 4) access to the real servers. Specify an IP address in dotted decimal notation. A range of IP addresses is produced when used with the `mmask` option.
`mmask` *<IP subnet mask (e.g., 255.255.255.0)>* This IP address mask is used with the `mnet` to select management traffic which is allowed direct real server access.
`pmask` *<IP subnet mask (e.g., 255.255.255.0)>* Sets persistent mask.
`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/secret` value and the `cntnt` 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) and by the RADIUS server to decrypt the password during verification.

Table 8-22 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, refer to "Direct Access Mode" on page 8-46.
`grace  disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables graceful real server failure. Allows existing connections to newly failed server to gracefully continue.
`matrix disable`|**`enable`** (or just **`d`**|**`e`**) Enables/disables the use of Virtual Matrix Architecture on the switch.
`cur` Displays the current Layer 4 advanced configuration.

/cfg/slb/adv/script

Scriptable Health Checks Configuration

```
[Health Script 1 Menu]
      open     - Add open command to end of script
      send     - Add send command to end of script
      expect   - Add expect command to end of script
      close    - Add close command to end of script
      rem      - Remove last command from script
      del      - Delete script
      cur      - Display current script configuration
```

Table 8-23 Scriptable Health Check Menu Options (/cfg/slb/adv/script)

Command Syntax and Usage

`open` *<TCP port-number>*

Sets the TCP port to be opened.

`send`

ASCII string to send through open TCP port.

`expect`

ASCII string expected for successful health check on open TCP port.

`close`

Closes TCP connection.

`rem`

Remove the last entered line from the script.

`del`

Delete the current script.

`cur`

List the current script configuration.

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.

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/direc`) 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 is not supported in some configurations.

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.

Virtual Matrix Architecture

Virtual Matrix Architecture (VMA) is a hybrid architecture that takes advantage of any unused resources within a Web switch by distributing the workload to multiple processors. Dividing the workload and using multiple processors to complete a task increases the number of concurrent sessions per switch.

NOTE – When VMA is enabled and a proxy IP address is configured, you must configure proxy IP addresses (PIPs) on all switch ports.

When `matrix` (VMA) is enabled, each client is assigned to a designated port's CPUs for Layer 4-7 processing, regardless of where it ingresses. The algorithm ensures even distribution of traffic. Packets to and from the same client are always processed by the same CPUs. Memory at all eight ports is pooled to increase storage capacity, enabling up to 512K session table entries, depending on platform and configuration; even when all traffic enters at a single port.

CHAPTER 9

The Operations Menu

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

/oper
Operations Menu

```
[Operations Menu]
      port     - Operational Port Menu
      mirr     - Operational Mirroring Menu
      slb      - Operational Server Load Balancing Menu
      vrrp     - Operational Virtual Router Redundancy Menu
      ip       - Operational IP Menu
      swkey    - Enter key to enable software feature
      rmkey    - Enter software feature to be removed
```

The commands of the Operations Menu enable you to alter switch operational characteristics without affecting switch configuration.

Port Mirroring menu options are accessible only to the Alteon AD4 and Alteon 184 Web switches.

Table 9-1 Operations Menu Options (/oper)

Command Syntax and Usage

`port` *<port as number (1-16)>*

Displays the Operational Port menu. To view menu options, refer to page 9-3.

`mirr`

Displays the Operational Mirroring menu. To view menu options, refer to page 9-4.

Table 9-1 Operations Menu Options (/oper)

Command Syntax and Usage

`slb`

Displays the Operational Layer 4 menu. To view menu options, refer to page 9-6.

`vrrp`

Displays the Operational Virtual Router Redundancy menu. To view menu options, refer to page 9-7.

`ip`

Displays the IP Operations menu, which has one sub-menu/option, the Operational Border Gateway Protocol Menu To view menu options, refer to page 9-7.

`swkey` *<16-hex-digit key to enable software feature>*

Enter key to enable software feature. For more information, refer to page 9-9.

`rmkey` *<software feature to be removed>*

Enter software feature to be removed. For more information, refer to page 9-10.

`/oper/port` *<port number>*
Operations-Level Port Options

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

Operations-level port options are used for temporarily disabling or enabling a port, and for changing RMON status on a port.

Table 9-2 Operations-Level Port Menu Options (/oper/port)

Command Syntax and Usage

`rmon`

Temporarily enables/disables RMON on the port. The port will be returned to its configured operation mode when the switch is reset.

`ena`

Temporarily enables the port. The port will be returned to its configured operation mode when the switch is reset.

`dis`

Temporarily disables the port. The port will be returned to its configured operation mode when the switch is reset.

`cur`

Displays the current settings for the port.

/oper/mirr
Operations-Level Port Mirroring Options

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

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

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

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

.

Table 9-3 Port Mirroring Menu Options (/oper/mirr)

Command Syntax and Usage

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

This defines the monitoring port. When port mirroring is enabled, packets received and/ or transmitted by the mirrored port will be duplicated to the switch port specified in this command.

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

This defines the mirrored port. When port mirroring is enabled, packets received and/or sent by the port specified in this command will be sent to the monitor port.

`dir` *<in|out|both>*

This determines which type of packets will be sent to the monitor port:

- `in` = packets received at the mirrored port
- `out` = packets sent from the mirrored port
- `both` = packets sent and received by the mirrored port

`tmout` *<seconds (0-86400)>*

Port mirroring will be automatically disabled (regardless of port state) after the time-out period specified in this command. Valid times are from 0 (does not time-out) to 86400 seconds.

`dis`

Turns port mirroring off.

`ena`

Turns port mirroring on.

`cur`

Displays the current parameter settings.

/oper/slb
Operations-Level SLB Options

```
[Server Load Balancing Operations Menu]
    ena      - Enable real server
    dis      - Disable real server
    synch    - Synchronize SLB, FILT, and VRRP configuration on peers
    clear    - Clear session table on port
    cur      - Current SLB operational state
```

When the optional Layer 4 software is enabled, the operations-level Server Load Balancing options are used for temporarily disabling or enabling real servers and synchronizing the configuration between the active/active switches..

Table 9-4 Server Load Balancing Operations Menu Options (/oper/slb)

Command Syntax and Usage

`ena` *<real server number (1-256)>*

Temporarily enables a real server. The real server will be returned to its configured operation mode when the switch is reset.

`dis` *<real server number (1-256)>*

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.

`synch`

Synchronizes the SLB, filter, and VRRP configuration on a peer switch (a switch that owns the IP address). To take effect, peers must be configured on the switches and the administrator password on the switches must be identical.

`clear`

Clears all session tables and allows port filter changes to take effect immediately. Note: This disrupts current Server Load Balancing and Application Redirection sessions.

`cur`

Displays the current SLB operational state.

/oper/vrrp
Operations-Level VRRP Options

```
[VRRP Operations Menu]
        back    - Set virtual router to backup
```

This menu is used to force a master virtual router to become backup router.

Table 9-5 Virtual Router Redundancy Operations Menu Options (/oper/vrrp)

Command Syntax and Usage
`back` *<virtual router number>* Forces the specified master virtual router on this switch into backup mode. This is generally used for passing master control back to a preferred switch once the preferred switch has been returned to service after a failure. When this command is executed, the current master gives up control and initiates a new election by temporarily advertising its own priority level as 0 (lowest). After the new election, the virtual router forced into backup mode by this command will resume master control in the following cases: ■ This switch owns the virtual router (the IP addresses of the virtual router and its IP interface are the same) ■ This switch's virtual router has a higher priority and preemption is enabled. ■ There are no other virtual routers available to take master control.

/oper/ip
Operations-Level IP Options

```
[IP Operations Menu]
      bgp      - Operational Border Gateway Protocol Menu
```

Table 9-6 IP Operations Menu Options (/oper/ip)

Command Syntax and Usage

bgp

Displays the BGP Operations Menu, shown below.

/oper/ip/bgp
Operations-Level BGP Options

```
[Border Gateway Protocol Operations Menu]
      start    - Start peer session
      stop     - Stop peer session
      cur      - Current BGP operational state
```

Table 9-7 IP Operations Menu Options (/oper/ip)

Command Syntax and Usage

start *<peer number (1-4)>*

Starts the peer session.

stop *<peer number (1-4)>*

Stops the peer session.

cur

Displays the current BGP operational state.

`/oper/swkey` Activating Optional Software

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

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

To obtain a software key, you must register each License Certificate with Alteon WebSystems, and provide the MAC address of the 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: [SLB]|GSLB|WCR: SLB
```

CHAPTER 10
The Boot Options Menu

To use the Boot Options Menu, you must be logged in to the switch as the administrator. The Boot Options Menu provides options for:

- Selecting a switch software image to be used when the switch is next reset
- Selecting a configuration block to be used when the switch is next reset
- Downloading a new software image to the switch via TFTP

To access the Boot Options Menu, at the Main Menu prompt, enter:

```
Main# boot
```

The Boot Options Menu is displayed:

```
[Boot Options Menu]
        image - Select software image to use on next boot
        conf  - Select config block to use on next boot
        tftp  - Download new software image via TFTP
        reset - Reset switch [WARNING: Restarts Spanning Tree]
        cur   - Display current boot options
```

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

Updating the Switch Software Image

The switch software image is the executable code running on the switch. A version of the image ships with the switch, and comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch.

Upgrading the software image on your switch requires the following:

- Loading the new image onto a TFTP server on your network
- Downloading the new image from the TFTP server to your switch
- Selecting the new software image to be loaded into switch memory the next time the switch is reset

Downloading New Software to Your Switch

The switch can store up to two different software images, called `image1` and `image2`, as well as boot software, called `boot`. When you download new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

To download a new software to your switch, you will need the following:

- The image or boot software loaded on a TFTP server on your network
- The hostname or IP address of the TFTP server
- The name of the new software image or boot file

NOTE – The DNS parameters must be configured if specifying hostnames. See "Domain Name System Configuration" on page 7-30).

When the above requirements are met, use the following procedure to download the new software to your switch.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# tftp
```

2. **Enter the name of the switch software to be replaced:**

```
Enter name of switch software image to be replaced
 ["image1"/"image2"/"boot"]:
```

3. **Enter the hostname or IP address of the TFTP server.**

```
Enter hostname or IP address of TFTP server:
```

4. **Enter the name of the new software file on the server.**

```
Enter name of file on TFTP server:
```

The exact form of the name will vary by TFTP server. However, the file location is normally relative to the TFTP directory (usually `/tftpboot`).

5. **The system prompts you to confirm your request.**

You should next select a software image to run, as described below.

Selecting a Software Image to Run

You can select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot.

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# image
```

2. **Enter the name of the image you want the switch to use upon the next boot.**

The system informs you of which image is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use switch software "image1" on next reset.
Specify new image to use on next reset ["image1"/"image2"]:
```

Selecting a Configuration Block

When you make configuration changes to the switch, you must save the changes so that they are retained beyond the next time the switch is reset. When you perform the `save` command, your new configuration changes are placed in the *active* configuration block. The previous configuration is copied into the *backup* configuration block.

There is also a *factory* configuration block. This holds the default configuration set by the factory when your switch was constructed. Under certain circumstances, it may be desirable to reset the switch configuration to the default. This can be useful when a custom-configured switch is moved to a network environment where it will be reconfigured for a different purpose.

Use the following procedure to set which configuration block you want the switch to load the next time it is reset:

1. **At the `Boot Options#` prompt, enter:**

```
Boot Options# conf
```

2. **Enter the name of the configuration block you want the switch to use:**

The system informs you of which configuration block is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use active configuration block on next reset.
Specify new block to use ["active"/"backup"/"factory"]:
```

Resetting the Switch

You can reset the switch to make your software image file and configuration block changes occur.

NOTE – Resetting the switch causes the Spanning-Tree Protocol to restart. This process can be lengthy, depending on the topology of your network.

To reset the switch, at the `Boot Options#` prompt, enter:

```
>> Boot Options# reset
```

You are prompted to confirm your request.

CHAPTER 11

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]
      uudmp   - Uuencode FLASH dump
      ptdmp   - tftp put FLASH dump to tftp server
      cldmp   - Clear FLASH dump
      panic   - Dump state information to FLASH and reboot
      sys     - System Maintenance Menu
      fdb     - Forwarding Database Manipulation Menu
      debug   - Debugging Menu
      arp     - ARP Cache Manipulation Menu
      route   - IP Route Manipulation Menu
      tsdmp   - 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 11-1 Maintenance Menu Options (/maint)

Command Syntax and Usage
`uudmp` Displays dump information in uuencoded format. For more information, refer to page 11-3.
`ptdmp` Saves the system dump information via TFTP. For more information, refer to page 11-4.
`cldmp` Clears dump information from flash memory. For more information, refer to page 11-4.
`panic` Dumps MP information to FLASH and reboots. For more information, refer to page 11-5.
`sys` Displays the System Maintenance Menu. To view menu options, refer to page 11-6.
`fdb` Displays the Forwarding Database Manipulation Menu. To view menu options, refer to page 11-6.
`debug` Displays the Debugging Menu. To view menu options, refer to page 11-7.
`arp` Displays the ARP Cache Manipulation Menu. To view menu options, refer to page 11-8.
`route` Displays the IP Route Manipulation Menu. To view menu options, refer to page 11-10.
`tsdmp` Tech Support dump.

/maint/uudmp
Uuencode Flash Dump

Using this command, dump information is presented in uuencoded format. This format makes it easy to capture the dump information as a file or a string of characters. You can then contact Alteon WebSystems Customer Support for help analyzing the information.

If you want to capture dump information to a file, set your communication software on your workstation to capture session data prior to issuing the `uudmp` command. This will ensure that you do not lose any information. Once entered, the `uudmp` command will cause approximately 1460 lines of data to be displayed on your screen and copied into the file.

Using the `uudmp` command, dump information can be read multiple times. The command does not cause the information to be updated or cleared from flash memory.

NOTE – Dump information is not cleared automatically. In order for any subsequent dump information to be written to flash memory, you must manually clear the dump region. For more information on clearing the dump region, see page 11-4.

To access dump information, at the `Maintenance#` prompt, enter:

```
Maintenance# uudmp
```

The dump information is displayed on your screen and, if you have configured your communication software to do so, captured to a file. If the dump region is empty, the following appears:

```
No FLASH dump available.
```

/maint/ptdmp *<server> <filename>*
TFTP System Dump Put

Use this command to `put` (save) the system dump via TFTP.

NOTE – If the TFTP server is running SunOS or the Solaris operating system, the specified `ptdmp` file must exist *prior* to executing the `ptdmp` command, and must be writable (set with proper permission, and not locked by any application). The contents of the specified file will be replaced with the current dump data.

To save dump information via TFTP, at the `Maintenance#` prompt, enter:

```
Maintenance# ptdmp server filename
```

Where *server* is the TFTP server IP address or hostname, and *filename* is the target dump file.

/maint/cldmp
Clearing Dump Information

To clear dump information from flash memory, at the `Maintenance#` prompt, enter:

```
Maintenance# cldmp
```

The switch clears the dump region of flash memory and displays the following message:

```
FLASH dump region cleared.
```

If the flash dump region is already clear, the switch displays the following message:

```
FLASH dump region is already clear.
```

/maint/panic
Panic Command

The `panic` command causes the switch to immediately dump state information to flash memory and automatically reboot.

To select `panic`, at the `Maintenance#` prompt, enter:

```
Maintenance# panic
```

Enter **y** to confirm the command:

```
Confirm dump and reboot [y/n]: y
```

The following messages are displayed:

```
Starting system dump...done.

Reboot at 11:54:08 Thursday June 26, 1997...

Boot version 1.0.1

Alteon ACEswitch 180

Rebooted because of console PANIC command.

Booting complete 11:55:01 Thursday June 26, 1997:
```

Unscheduled System Dumps

If there is an unscheduled system dump to flash memory, the following message is displayed when you log on to the switch:

```
Note: A system dump exists in FLASH. The dump was saved
      at 13:43:22 Fri Jun 27, 1997. Use /maint/uudmp to
      extract the dump for analysis and /maint/cldmp to
      clear the FLASH region. The region must be cleared
      before another dump can be taken.
```

/maint/sys System Maintenance Options

This menu is reserved for use by Alteon WebSystems Customer Support. The options are used to perform system debugging.

/maint/fdb Forwarding Database Options

```
[FDB Manipulation Menu]
      find     - Show a single FDB entry by MAC address
      port     - Show FDB entries for a single port
      vlan     - Show FDB entries for a single VLAN
      refpt    - Show FDB entries referenced by a single port
      dump     - Show all FDB entries
      del      - Delete an FDB entry
      clear    - Clear entire FDB
```

The Forwarding Database Manipulation Menu can be used to view information, and to delete a MAC address from the forwarding database or clear the entire forwarding database. This is helpful in identifying problems associated with MAC address learning and packet forwarding decisions.

Table 11-2 FDB Manipulation Menu Options (maint/fdb)

Command Syntax and Usage
`find` *<MAC address>* [*<VLAN>*] Displays a single database entry by its MAC address. You are prompted to enter the MAC address of the device. Enter the MAC address using the format, `xx:xx:xx:xx:xx:xx`. For example, `08:00:20:12:34:56`. You can also enter the MAC address using the format, `xxxxxxxxxxxx`. For example, `080020123456`.
`port` *<port number (1-16)>* Displays all FDB entries for a particular port.
`vlan` *<VLAN number (1-4094)>* Displays all FDB entries on a single VLAN.

Table 11-2 FDB Manipulation Menu Options (maint/fdb)

Command Syntax and Usage

`refpt` *<port number>*

Displays all FDB entries reference by a single port.

`dump`

Displays all entries in the Forwarding Database. For more information, refer to page 5-13.

`del` *<MAC address>*

Removes a single FDB entry.

`clear`

Clears the entire Forwarding Database from switch memory.

/maint/debug
Debugging Options

```
[Miscellaneous Debug Menu]
      tbuf    - Display MP trace buffer
      snap    - Display MP snap (or post-mortem) trace buffer
      sptb    - 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 the Switch Processor (SP)
- Events traced to a buffer area when a reset occurs

If the switch resets for any reason, the MP trace buffer and SP trace buffers are saved into the snap trace buffer area.

The output from these commands can be interpreted by the Alteon WebSystems Customer Support organization.

Table 11-3 Miscellaneous Debug Menu Options (maint/debug

Command Syntax and Usage

`tbuf` *<IP address>*

Displays the Management Processor trace buffer. Header information similar to the following is shown:

```
MP trace buffer at 18:27:37 Mon Dec 29, 1997; mask: 0x2ffff748
```

The buffer information is displayed after the header.

`snap` *<IP address>*

Displays the Managment Processor snap (or post-mortem) trace buffer. This buffer contains information traced at the time that a reset occurred.

`sptb` *<port number>*

Displays the Switch Processor trace buffer. Header information similar to the following is shown:

```
Port 1 trace buffer at 18:27:41 Mon Dec 29, 1997; mask:0x018007e
```

The buffer information is displayed after the header.

/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 11-4 Address Resolution Protocol Menu Options (maint/arp)

Command Syntax and Usage

`find` *<IP address>*

Shows a single ARP entry by IP address.

`port` *<port number>*

Shows ARP entries on a single port.

`vlan` *<VLAN ID>*

Shows ARP entries on a single

`refpt` *<port number>*

Shows all ARP entries referenced by a single port.

`dump`

Shows all ARP entries.

`add` *<IP address>*

Adds a single ARP entry from switch memory.

`del` *<IP address>*

Removes a single ARP entry from switch memory.

`clear`

Clears the entire ARP list from switch memory.

`addr`

Shows the list of IP addresses 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 5-11.

/maint/route
IP Route Manipulation

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

Table 11-5 IP Route Manipulation Menu Options (maint/route)

Command Syntax and Usage

`find` *<IP address>*

Shows a single route by destination IP address.

`gw` *<gateway>*

Shows routes to a single gateway.

`type` *<type>*

Shows routes of a single type.

`tag` *<tag number>*

Shows routes of a single tag.

`if` *<interface number>*

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 5-8.

Appendix A
WebOS Syslog Messages

The following syntax is used when outputting syslog messages:

<Time stamp><Log Label>`WebOS`*<Thread ID>*`:`*<Message>*

where

- *<Timestamp>*

  The time of the message event is displayed in month day hour:minute:second format. For example: `Aug 19 14:20:30`

- *<Log Label>*

  The following types of log messages are recorded: `LOG_EMERG`, `LOG_ALERT`, `LOG_CRIT`, `LOG_ERR`, `LOG_WARNING`, `LOG_NOTICE`, `LOG_INFO`, and `LOG_DEBUG`

- *<Thread ID>*

  This is the software thread that reports the log message. The following thread IDs are recorded: `stp`, `ip`, `slb`, `console`, `telnet`, `vrrp`, `system`, `web server`, `ssh`, and `bgp`

- *<Message>*: The log message

Following is a list of potential syslog messages. To keep this list as short as possible, only *<Thread ID>* and *<Message>* are shown. The messages are sorted by *<Log Label>*.

Where the *<Thread ID>* is listed as “mgmt,” one of the following may be shown: `console`, `telnet`, `web server`, or `ssh`.

LOG_ALERT

<mgmt>: ERROR: Synchronization from non-configured peer <ip_address> was blocked

<mgmt>: new synch configuration did not apply (rc=<error_code>)

<mgmt>: new synch configuration did not save (rc=<error_code>)

<mgmt>: new synch configuration did not validate (rc=<error_code>)

<mgmt>: Sync Password Failed-No Password Line

<mgmt>: Synch Password Failed-Bad Password

<mgmt>: WARNING: Synchronization from non-configured peer <ip_address>

bgp: notification (<reason>) received from <BGP peer ip_address>

bgp: session with <BGP peer ip_address> failed (<reason>)

ip: cannot contact default gateway <ip_address>

slb: cannot contact real server <ip_address>

slb: cannot contact real service <ip_address:real_port>

slb: real server <ip_address> disabled through configuration

slb: real server <ip_address> has reached maximum connections

slb: real server failure threshold (<threshold>) has been reach for group <group_id>

slb: received update from <ip_address> for unknown remote server

slb: received update from <ip_address> for unknown virtual service

stp: own BPDU received from port <port_id>

vrrp: received errored advertisement from <ip_address>

vrrp: received incorrect addresses from <ip_address>

vrrp: received incorrect advertisement interval <seconds> from <ip_address>

vrrp: received incorrect password from <ip_address>

LOG_CRIT

system: can't allocate memory in load_MP_INT

system: internal power supply failed

system: redundant power supply failed

system: temperature at sensor <sensor_id> exceeded threshhold

LOG_ERR

<mgmt>: <"apply"|"save"> is issued by another user. Try later

<mgmt>: <"Apply"|"Save"> not done

<mgmt>: A hot-standby port cannot also be an inter-switch port

<mgmt>: At least one virtual router must be enabled when group is enabled

<mgmt>: BGP peer <bgp_peer_id> have same address as IP interface <ip_interface_id>

<mgmt>: BGP peer <bgp_peer_id> IP interface <ip_interface_id> is not enabled

<mgmt>: BGP peer <bgp_peer_id> must have an IP address

<mgmt>: BGP peers <bgp_peer_id> and <bgp_peer_id> have same address

<mgmt>: Broadcast address for IP interface <interface_id> is invalid

<mgmt>: Client bindings are not supported with proxy IP addresses

<mgmt>: DAM must be turned on or a PIP must be enabled for port <port_id> in order for virtual server %lu to support FTP parsing

<mgmt>: DAM must be turned on or a PIP must be enabled for port <port_id> in order for virtual server <server_id> to support URL parsing

<mgmt>: DAM must be turned on or a PIP must be enabled for ports <port_id> in order to do URL based redirection

<mgmt>: Direct access mode is not supported with default gateway load balancing

<mgmt>: domain name must be configured

<mgmt>: duplicate default entry

<mgmt>: Dynamic NAT filter <filter_id> must be cached

<mgmt>: Enabled external lookup IP address has no IP address

<mgmt>: Enabled real server <server_id> has no IP address

<mgmt>: Enabled virtual server <server_id> has no IP address

<mgmt>: Error writing BGP changes to FLASH

<mgmt>: Error writing BWM changes to FLASH

<mgmt>: Error writing FILT changes to FLASH

<mgmt>: Error writing GSLB changes to FLASH

<mgmt>: Error writing HCS changes to FLASH

<mgmt>: Error writing IP changes to FLASH

<mgmt>: Error writing NAME changes to FLASH

<mgmt>: Error writing NTP changes to FLASH

<mgmt>: Error writing RSA changes to FLASH

<mgmt>: Error writing SLB changes to FLASH

LOG_ERR (continued)

<mgmt>: Error writing SSH changes to FLASH

<mgmt>: Error writing to FLASH

<mgmt>: Error writing URL changes to FLASH

<mgmt>: Error writing URL changes to FLASH

<mgmt>: Error writing VRRP changes to FLASH

<mgmt>: Extracting length has to set to 8 or 16 for cookie rewrite mode

<mgmt>: Filter with ICMP types configured (<icmp_type>) must have IP protocol configure to ICMP

<mgmt>: Filter with L4 ports configured <port_id> must have IP protocol configured

<mgmt>: For Global SLB, Web server must be moved from TCP port 80

<mgmt>: Hot-standby must be enabled when a virtual router has a PIP address

<mgmt>: intrval input value must be in the range [0-24]

<mgmt>: IP Interfaces <interface_id> and <interface_id> are on the same subnet

<mgmt>: Loadbalance string must be added to real server <server_id> in order to enable exclusionary string matching

<mgmt>: multiple static routes have same destination

<mgmt>: NAT filter <filter_id> cannot have port ranges

<mgmt>: NAT filter <filter_id> dest range includes RIP <server_id>

<mgmt>: NAT filter <filter_id> dest range includes VIP <server_id>

<mgmt>: NAT filter <filter_id> must be cached

<mgmt>: NAT filter <filter_id> must have same smask and dmask

<mgmt>: Network <static_network_id> has no VIP address

<mgmt>: New Path Cost for Port <port_id> is invalid

<mgmt>: No apply is needed, although a save is needed

<mgmt>: No apply is needed, although there are saved changes

<mgmt>: No apply needed

<mgmt>: Not all ports in trunk group <trunk_id> are in VLAN <vlan_id>

<mgmt>: Only <MAX_SLB_SERVICES> remote services are supported

<mgmt>: Only <MAX_SLB_SITES> remote servers are allowed per group

<mgmt>: Please configure primary RADIUS server address

<mgmt>: Port filtering must be disabled on port <port_id> in order to support cookie based persistence for virtual server <server_id>

<mgmt>: Port Mirroring changes are not applied

<mgmt>: Primary and secondary remote site <site_id> switches must differ

LOG_ERR (continued)

<mgmt>: PVID <vlan_id> for port <port_id> is not created

<mgmt>: RADIUS secret must be 1-32 characters long

<mgmt>: Real server <server_id> (Backup for <server_id>) is not enabled

<mgmt>: Real server <server_id> and group %u cannot both have backups configured

<mgmt>: Real server <server_id> cannot be added to same group

<mgmt>: Real server <server_id> cannot be backup server for both real server <server_id> and group <group_id>

<mgmt>: Real server <server_id> has same IP address as IP interface <interface_id>

<mgmt>: Real server <server_id> has same IP address as real server <server_id>

<mgmt>: Real server <server_id> has same IP address as switch

<mgmt>: Real server <server_id> has same IP address as virtual server<server_id>

<mgmt>: Real server group <group_id> cannot backup itself

<mgmt>: Redirection filter <filter_id> must be cached

<mgmt>: Remote site <site_id> and real server <server_id> must use different addresses

<mgmt>: Remote site <site_id> and virtual server <server_id> must use different addresses

<mgmt>: Remote site <site_id> does not have a primary IP address

<mgmt>: Remote sites <site_id> and <site_id> must use different addresses

<mgmt>: RS <server_id> can't exist for VS <server_id> vport <virtual_port>

<mgmt>: Save not done

<mgmt>: Save the configuration before resetting switch

<mgmt>: SLB Radius secret must be 16 characters long

<mgmt>: STP changes can't be applied since STP is OFF

<mgmt>: Switch cannot support more than <MAX_SMT> real services

<mgmt>: Switch cannot support more than <MAX_VIRT_SERVICES> virtual services

<mgmt>: Switch port <port_id> has same IP address as IP interface <interface_id>

<mgmt>: Switch port <port_id> has same proxy IP address as port <port_id>

<mgmt>: Switch reset is required to turn STP on/off

<mgmt>: There must be at least one inter-switch port if any hot-standby port exist

<mgmt>: Trunk group (<trunk_id>) ports must all have a PIP

<mgmt>: Trunk group (<trunk_id>) ports must have same L4 config

<mgmt>: Trunk group <trunk_id> contains no ports but is enabled

<mgmt>: Trunk group <trunk_id> contains ports with different PVIDs

<mgmt>: Trunk group <trunk_id> has more than <max_trunk_ports> ports

LOG_ERR (continued)

<mgmt>: Trunk groups <trunk_id> and <trunk_id> can not share the same port

<mgmt>: Two services have same hostname, <host_name>.<domain_name>

<mgmt>: Two services have same hostname, <host_name>.<domain_name>

<mgmt>: Virtual router <vr_id> cannot have same IP address as <ip_address>

<mgmt>: Virtual router <vr_id> cannot have same VRID and VLAN as <vlan_id>

<mgmt>: Virtual router <vr_id> corresponding virtual server <server_id> is not enabled

<mgmt>: Virtual router <vr_id> IP interface should be <interface_id>

<mgmt>: Virtual router <vr_id> must have an IP address

<mgmt>: Virtual router <vr_id> must have sharing disabled when hotstandby is enabled

<mgmt>: Virtual router group must be enabled when hotstandby is enabled

<mgmt>: Virtual router group must have preemption enabled when hotstandby is enabled

<mgmt>: Virtual router group must have sharing disabled when hotstandby is enabled

<mgmt>: Virtual server %lu: support nonat IP but not layer 3 bindings

<mgmt>: Virtual server <server_id> has same IP address and vport as virtual server <server_id>

<mgmt>: Virtual server <server_id> has same IP address as IP interface <interface_id>

<mgmt>: Virtual server <server_id> has same IP address as switch

<mgmt>: Virtual server <server_id>: port mapping but Direct Access Mode

<mgmt>: Virtual server <server_id>: port mapping but layer3 bindings

<mgmt>: Virtual server <server_id>: UDP service <virtual_port> with out-of-range port number

<mgmt>: Virtual servers <server_id> and <server_id> that include the same real server <server_id> cannot map the same real port or balance UDP

<mgmt>: Virtual servers <server_id> and <server_id> with same IP address must support same layr3 configuration

<mgmt>: Virtual servers: all that support IP must use same group

<mgmt>: VLAN <vlan_id> has a member port that can not support jumbo frame

<mgmt>: With VMA, ports 1-8 must all have a PIP if any one does

ip: cannot contact NTP server <ip_address>

ip: unable to listen to NTP port

LOG_NOTICE

<mgmt>: boot config block changed

<mgmt>: boot image changed

<mgmt>: second syslog host changed to <ip_address>

<mgmt>: switch reset from CLI

<mgmt>: syslog host changed to <ip_address>

system: internal power supply ok

system: rebooted <last_reset_information>

system: rebooted <last_reset_information> administrator logged in <mgmt>:
    Next boot will use new software image<1|2>

system: redundant power supply present and ok

LOG_WARNING

slb: filter <filter_id> fired on port <port_id>, <source_ip> -> <dest_ip>

Appendix B
WebOS SNMP Agent

The WebOS SNMP agent supports SNMP Version 1. Security is provided through SNMP community strings. The default community strings are “public” for `SNMP GET` operation and “private” for `SNMP SET` operation. The community string can be modified only through the command-line interface (CLI). Alteon WebSystems is registered as Vendor 1872. Detailed SNMP MIBs and trap definitions of the WebOS SNMP agent can be found in the following Alteon WebSystems enterprise MIB documents:

- Altroot.mib - Alteon product registrations which are returned as sysObjectID.
- Altswitch.mib - Alteon enterprise MIB definitions.
- Alttrap.mib - Alteon enterprise trap definitions.

Users may specify up to two trap hosts for receiving SNMP Traps. The agent will send the SNMP Trap to the specified hosts when appropriate. Traps will not be sent if there is no host specified.

WebOS SNMP agent supports the following standard MIBs:

- RFC 1213 - MIB II (System, Interface, Address Translation, IP, ICMP, TCP, UDP, SNMP Groups)
- RFC 1573 - MIB II Extension (IFX table)
- RFC 1643 - EtherLike MIB
- RFC 1493 - Bridge MIB
- RFC 1757 - RMON MIB (Statistics, History, Alarm, Event Groups)

WebOS SNMP agent supports the following generic traps as defined in RFC 1215:

- ColdStart
- WarmStart
- LinkDown
- LinkUp
- AuthenticationFailure

The SNMP agent also supports two Spanning Tree traps as defined in RFC 1493:

- NewRoot
- TopologyChange

The following are the enterprise SNMP traps supported in WebOS:

Table B-1 WebOS-Supported Enterprise SNMP Traps

Trap Name	Description
altSwPrimaryPowerSuppylFailure	Signifies that the primary power supply failed.
altSwRedunPowerSuppylFailure	Signifies that the redundant power supply failed.
altSwDefGwUp	Signifies that the default gateway defined is alive.
altSwDefGwDown	Signifies that the default gateway defined is down
altSwDefGwInService	Signifies that the default gateway is up and in service
altSwDefGwNotInService	Signifies that the default gateway is alive but not in service
altSwSlbRealServerUp	Signifies that the real server is up and operational
altSwSlbRealServerDown	Signifies that the real server is down and out of service
altSwSlbRealServerMaxConnReached	Signifies that the real server has reached maximum connections
altSwSlbBkupRealServerAct	Signifies that the backup real server is activated due to availablity of the primary real server
altSwSlbBkupRealServerDeact	Signifies that the backup real server is deactivated due to the primary real server is available
altSwSlbBkupRealServerActOverflow	Signifies that the backup real server is deactivated due to the primary real server is overflowed
altSwSlbBkupRealServerDeactOverflow	Signifies that the backup real server is deactivated due to the primary real server is out from overflow situation
altSwSlbFailoverStandby	Signifies that the switch is now a standby switch
altSwSlbFailoverActive	Signifies that the switch is now an active switch
altSwSlbFailoverSwitchUp	Signifies that the failover switch is alive
altSwSlbFailoverSwitchDown	Signifies that the failover switch is down
altSwfltFilterFired	Signifies that the packet received on a switch port matches the filter rule
altSwSlbRealServerServiceUp	Signifies that the service port of the real server is up and operational
altSwSlbRealServerServiceDown	Signifies that the service port of the real server is down and out of service

Glossary

RIP (Real Server)
An IP addresses that the switch load balances to when requests are made to a virtual IP address (VIP).

Real Server Group
Group of real servers that are associated with a VIP or filter.

VIP (Virtual IP Address)
An IP address that the switch owns and uses to load balance particular service requests (like HTTP) to other servers.

SIP (Source IP Address)
The source IP address of a frame.

DIP (Destination IP Address)
The destination IP address of a frame.

SPort (Source Port)
The source port (application socket; for example, HTTP-80/HTTPS-443/DNS-53).

Dport (Destination Port)
The destination port (application socket; for example, http-80/https-443/DNS-53)

Proto (Protocol)
The protocol of a frame. Can be any value represented by a 8-bit value in the IP header adherent to the IP specification (for example, TCP, UDP, OSPF, ICMP, and so on.)

NAT (Network Address Translation)
Any time an IP address is changed from one SIP or DIP to another address, network address translation can be said to have taken place. In general, half NAT is when the DIP or SIP is changed from one address to another and full NAT is when both addresses are changed from one address to another. VIP based load balancing uses half NAT by design since it NAT's the DIP (destination IP address) from the VIP (virtual IP address) to that of one of the RIP's (real servers).

Virtual Server Load Balancing

Classic load balancing. Requests destined for a virtual server IP address (VIP), which is owned by the switch, are load balanced to a real server contained in the group associated with the VIP. Network address translation is done back and forth, by the switch, as requests come and go.

Frames come to the switch destined for the VIP. The switch then replaces the VIP and with one of the real server IP addresses (RIP's), updates the relevant checksums, and forwards the frame to the server for which it's now destined. This process of replacing the destination IP (VIP) with one of the real server addresses is called half NAT. If the frames were not half NAT'ed to the address of one of the RIPs, a server would receive the frame that was destined for it's MAC address, forcing the packet up to Layer 3. The server would then drop the frame, since the packet would have the DIP of the VIP and not that of the server (RIP).

Redirection or Filter-Based Load Balancing

A type of load balancing; one that operates differently from VIP-based load balancing. With this type of load balancing, requests are transparently intercepted and "redirected" to a server group. "Transparently" meaning that requests are not specifically destined for a VIP that the switch owns. Instead, a filter is configured in the switch. This filter intercepts traffic based on certain IP header criteria and load balances it.

Filters can be configured to filter on the SIP/Range (via netmask), DIP/Range (via netmask), Protocol, SPort/Range or DPort/Range. The action on a filter can be Allow, Deny, Redirect to a Server Group, or NAT (either the SIP or DIP). When doing redirection based load balancing the DIP is NOT NAT'ed to that of one of the real servers. Therefore, redirection based load balancing is designed to be used to load balance devices that normally operate transparently in your network such as a firewall, spam filter, or transparent Web cache.

VRRP (Virtual Router Redundancy Protocol)

A protocol that acts very similarly to Cisco's proprietary HSRP address sharing protocol. The reason for both of these protocols is so devices have a next hop or default gateway that is always available. The way it works is two or more devices sharing an IP interface are either advertising or listening for advertisements. These advertisements are sent via a broadcast message to address 224.0.0.18.

With VRRP one switch is considered the master and the other the backup. The master is always advertising via the broadcasts. The backup switch is always listening for the broadcasts. Should the master stop advertising the backup will take over ownership of the VRRP IP and MAC addresses as defined by the specification. The switch announces this change in ownership to the devices around it by way of a Gratuitous ARP and advertisements. If the backup switch didn't do the Gratuitous ARP the Layer 2 devices attached to the switch would not know that the MAC address had moved in the network. For a more detailed description, refer to RFC 2338.

Virtual Router

A shared address between two devices utilizing VRRP, as defined in RFC 2338. One virtual router is associated with an IP interface. This is one of the IP interfaces that the switch is assigned. All IP interfaces on the Alteon switches must be in a VLAN. If there is more than one VLAN defined on the switch then the VRRP broadcasts will only be sent out on the VLAN for which the associated IP interface is a member of.

Priority

The value given to a Virtual Router to determine it's ranking with it's peer(s). Minimum value is 1 and maximum value is 254. Default is 100. A higher number will win out for master designation.

VRID (Virtual Router Identifier)

A value between 1 and 255 that is used by each virtual router to create it's MAC address and identify it's peer for which it is sharing this VRRP address. The VRRP MAC address as defined in the RFC is 00-00-5E-00-01-{VRID}. If you have a VRRP address that two switches are sharing, then the VRID number needs to be identical on both switches so each virtual router on each switch knows who to share with.

VIR (Virtual Interface Router)

A VRRP address that is an IP interface address shared between two or more virtual routers.

VSR (Virtual Server Router)

A VRRP address that is a shared VIP address. This is Alteon's proprietary extension to the VRRP spec. The switches must be able to share a VIPs as well as IP interfaces. If they didn't the two switches would fight for ownership of the VIP and the ARP tables in the devices around them would get very confused since they would have two ARP entries with the same IP address but different MAC addresses.

Preemption

Pre-emption will cause a Virtual Router that has a lower priority to go into backup should a peer Virtual Router start advertising with a higher priority.

Tracking

A method to increase the priority of a virtual router and thus master designation (with preemption enabled). Tracking can be very valuable in an active/active configuration.

You can track the following:

- Vrs: Virtual Routers in Master Mode (increments priority by 2 for each)
- Ifs: Active IP interfaces on the switch (increments priority by 2 for each)
- Ports: Active ports on the same VLAN (increments priority by 2 for each)
- l4pts: Active Layer 4 Ports, client or server designation (increments priority by 2 for each
- reals: healthy real servers (increments by 2 for each healthy real server)
- hsrp: HSRP announcements heard on a client designated port (increments by 10 for each)

Index