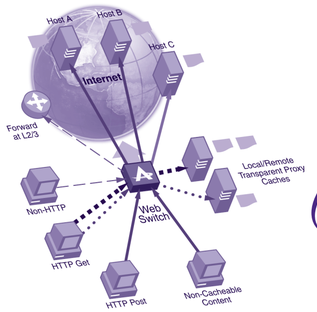
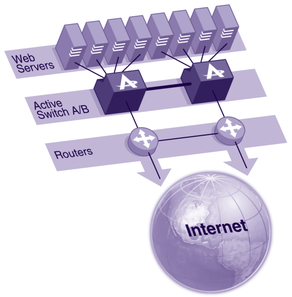
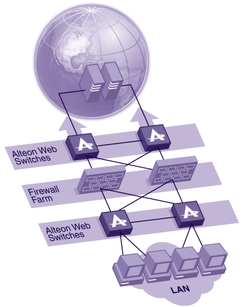


Application Redirection



High Availability



Firewall Load Balancing

ORDERING INFORMATION

STANDARD Web OS

Standard Web OS ships as an integral part of every Alteon WebSystems' switch, providing:

- Layer 2 switching
- Layer 3 IP switching and routing
- Filtering
- EtherChannel support
- Comprehensive network management including CLI, Web UI, SNMP and RMON

SERVER LOAD BALANCING (SLB) & APPLICATION REDIRECTION / (AR)

- 720004 Web OS SLB/AR (Alteon 708)
- 720001 Web OS SLB/AR (Alteon 180 only. This software is included with other Web switch models)

Web OS SLB/AR adds Layer 4 through 7 switching functionality to Alteon WebSystems' switches, including:

- Server load balancing
- Application redirection
- Content intelligence
- Server security

GLOBAL SERVER LOAD BALANCING (GSLB)

- 720006 Web OS GSLB (Alteon 708)
- 720003 Web OS GSLB (Alteon 180-Plus, Alteon 180e, Alteon 184, ACEdirector 2, ACEdirector 3, ACEdirector 4)

Web OS GSLB adds the ability to distribute loads across multiple Web farms, even in geographically diverse locations.

BANDWIDTH MANAGEMENT (BWM)

- 720010 Web OS BWM (Alteon 180e, Alteon 184, ACEdirector 3, ACEdirector 4)

Web OS BWM enables new enables service providers to offer a new class of value-added services based on application, transaction type, or service agreement. It provides the ability to allocate and meter traffic bandwidth using Layer 2 – 7 attributes, including filters, URLs, and cookies.

NORTEL NETWORKS™



Alteon Web OS

APPLICATION AND CONTENT-INTELLIGENT TRAFFIC CONTROL SERVICES

Alteon WebSystems' Web OS provides Layer 2 through 7 Web traffic control services integrated with Alteon's Web switch product lines. Web OS gives administrators unprecedented control of their IP networks by providing local and global server load balancing, application redirection, non-server (e.g., firewall, VPN, and router) load balancing, active-active high availability, bandwidth management, and server security services.

FEATURES

- High-performance content-intelligent switching using HTTP headers, cookies, and URLs allows infrastructure optimization
- TCP, UDP, and IP server load balancing, including HTTP (persistent and non-persistent), FTP and passive FTP, SSL, SMTP, LDAP, DNS, Radius, Telnet and NNTP.
- Virtual Matrix Architecture enables dynamic utilization of all processors and memory
- Gigabit-class, content-intelligent bandwidth management enables SLAs and usage-based services
- Full inspection of URLs, cookies, and host headers, up to 4,500 bytes across multiple request packets and also across multiple responses coming from the server
- Global server load balancing to distributed servers based on health, user proximity, server weights, and response times
- Persistent connections using source IP addresses, HTTP cookies, or SSL IDs
- User-scriptable server health checks enable content verification
- Complete site redundancy and failover
- Application redirection for any traffic type
- Firewall and WAN router load balancing
- Up to 32 million concurrent sessions (700 series)

Local Server Load Balancing (SLB) Increases Application Performance and Availability

Web OS local server load balancing enables virtually unbounded server capacity by transparently distributing client traffic across a server farm. Web OS ensures high application availability through the use of continuous server, application, and content health checking, server load monitoring, support for backup and overflow servers, and active-active switch failover configurations.

With SLB enabled, Alteon's Web switches appear to the network as "virtual servers," each represented by a Virtual IP (VIP) address or address range. Web OS supports a virtually unlimited number of VIP addresses, with up to 256 services allowed per VIP address. For maximum flexibility, real servers can be associated with any number of virtual servers.

The switches distribute load by assigning each application request to the most available real server associated with the destination VIP address and tracking each session from start to finish.

Effective load balancing algorithms tailor request distribution based on application and server types. These algorithms include least connections and round robin, with per server weighting and maximum load thresholds, plus hashing to ensure server persistence. Excess load is directed to designated overflow servers that can be located anywhere in the network. Alteon's Web OS API also allows users to design custom server-based scripts and agents to dynamically control load distribution on Alteon Web switches.

Web OS also supports URL (Uniform Resource Locator)-based load balancing, letting administrators optimize servers for different tasks and for storing specific object types. Web OS directs user requests to the appropriate server or server group based on URL or URL sub-string matches.

Application Redirection Provides High Performance Policy Routing and Load Balancing

With application redirection, powerful filtering rules are used to intercept any type of IP traffic and redirect it to a designated server, server farm or IP interfaces. Filters can be set by source and destination MAC and IP addresses, protocols, Layer 4 source and destination port numbers, as well as URLs and cookies.

Any or all filters can be applied dynamically to each switch port to allow, deny or redirect packets. Application redirection offloads policy routing from expensive routers and increases system performance.

Application redirection also allows DNS requests to be intercepted and transparently redirected to a local DNS server or group of DNS servers regardless of the destination DNS address in the request. This minimizes bandwidth usage and increases user performance.

By using application redirection with transparent Web caches, Internet access response time and WAN bandwidth consumption is reduced without placing caches in the data path. Any or all Web page requests are transparently redirected to local Web caches.

Global Server Load Balancing (GSLB) Directs Client Requests to the Nearest, Most Responsive Server

Global server load balancing allows content and servers to be distributed to up to 64 sites around the world and directs client requests to the best site, based on server health, proximity to the client, server weights, and response times.



Alteon Web OS

CONTENT-INTELLIGENT TRAFFIC CONTROL SERVICES

MORE FEATURES

- Up to 12 million session setups and teardowns per second (700 series)
- Up to 8192 services per switch
- Distribute traffic load across 64 geographically distributed sites
- Up to 1024 servers per switch
- Up to 224 packet filtering rules per port
- Full network address translation enables multi-site load balancing and traffic redirection
- Active-active and hot-standby failover for high availability, with in-band keep-alives and configuration updates
- Advanced server security features, with multi-layer access control

www.alteonwebsystems.com

Each Web switch monitors the health, response time, and performance of other sites and exchanges this information with all other GSLB Web switches. With a global view of each site's health and performance, Alteon Web switches direct traffic to sites in proportion to their performance metrics.

As a result, the best performing sites receive a greater number of user requests than other sites, in proportion to their ability to handle additional requests. This gives users consistently good response times, regardless of where their requests are directed.

In the event all servers at a GSLB site fail or are overloaded, the Web switch at the site immediately notifies its peer switches and redirects all incoming request to the best peer site using HTTP Redirect or IP Proxy. This ensures no service disruption.

Content-Intelligent Web Cache Redirection and Server Load Balancing Optimizes Infrastructure

Cache server farm processing can be optimized through URL parsing. Using sub-string matching, requests with specific URLs or URL strings can be identified and transparently redirected to designated cache servers. If multiple caches are used, the Web switch automatically distributes client requests, using an algorithm that maximizes cache hits.

Because they look at the URL and HTTP commands in each request, Web OS switches offload Web caches by redirecting to them only HTTP GET requests for cacheable objects and forwarding other requests directly to the origin servers.

For example, non-GET requests, such as HEAD, POST, PUT, as well as CGI or ASP calls, can be automatically identified by the Web switch and bypassed from the cache - optimizing Web cache operations.

By including the host header in URL parsing, Alteon's Virtual Hosting solution makes it possible to host multiple websites on a single switch, using one virtual IP address. Client requests can also be redirected based on browser type or other information in the host header.

For example, browser-based load balancing enables service providers to distinguish wireless Internet users from customers using wireline browsers and direct clients to servers with information and services tailored to the needs of each group.

Firewall Load Balancing Speeds User Response Times, Increases Site Availability

Network managers can also use redirection filters to add the benefits of scaling and resiliency to inline packet processing devices such as firewalls. For instance, application redirection allows traffic to be distributed among multiple (up to 256), fully active firewalls to improve performance, with no changes to firewall software. Health checks are performed on the entire data path through load-balanced firewalls, increasing resilience and availability.

VPN Load Balancing Provides Scalability for Multiple Secure VPN Connections

Alteon's Web OS provides enhanced scalability with VPN load balancing, ensuring secure connectivity for VPN sessions and high availability in the event a VPN gateway fails.

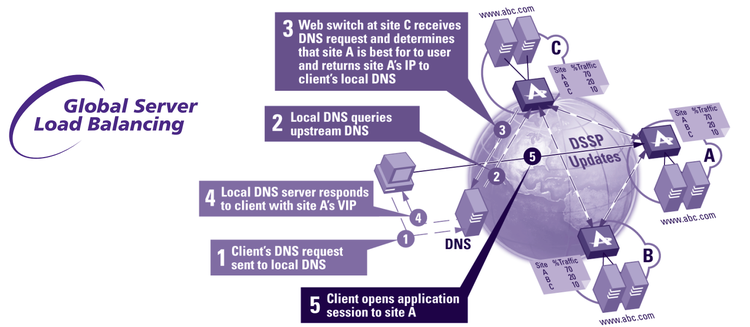
Using VPN load-balancing services, user traffic can be distributed among up to 255 active VPN gateways, enhancing scalability and ensuring security. If a VPN gateway fails, sophisticated health-checking mechanisms redirect traffic to a healthy device to maintain high availability of website services.

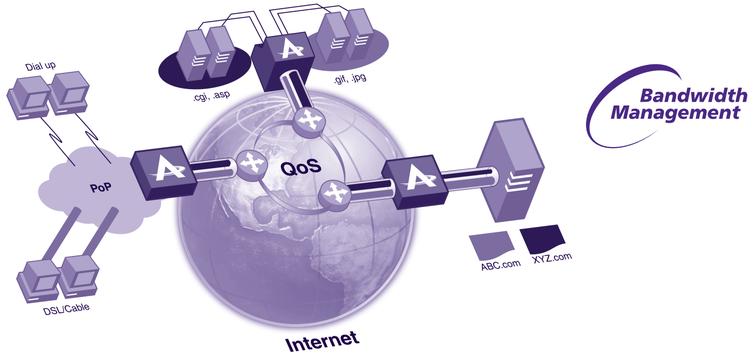
Flexible Persistence Support Enables Effective E-Commerce Infrastructure

Web OS provides a variety of persistence methods to facilitate server assignment based on specific user information. Persistence can be based on IP source address, HTTP cookies for HTTP requests, or SSL session ID for encrypted HTTPS requests. With these persistence methods, requests from new users are load balanced to the best available server, based on the configured load-balancing algorithm. Subsequent requests from the same users are automatically tracked and redirected to the server that initially handled their request.

Web OS cookie caching transparently handles mega- and micro-proxies for HTTP traffic, using the cookie values embedded in client requests and server responses to ensure session persistence. Cookies can also be used to identify one type of user from another and redirect specific client requests to a particular server group, enabling client-based services such as preferred customers and user profiling.

The session persistence required for secure eCommerce transactions is intelligently handled by





Alteon's iSD-SSL Web optimization platform. The Alteon iSD-SSL accelerates "shopping cart" sessions by offloading Secure Sockets Layer (SSL) processing from local servers without imposing delays on other traffic in the data path.

Bandwidth Management and Quality of Service Enables New Services and Improves User Experience

Using switch-based bandwidth management, Web hosters and service providers can offer their customers a new class of value-added services based on specific applications, transaction types, or service agreements. Web OS bandwidth management provides granular control of bandwidth and traffic into and out of servers while enabling a high level of security. Because Web OS-based switches are designed to connect directly to servers and have multi-processor computing horsepower, they are uniquely positioned to meter server output and enable usage-based services.

Traffic bandwidth can be allocated and metered based on any filtering rule, as well as URLs and HTTP cookies and headers. Classification policies allow bandwidth limitations to be applied to particular applications. Different classes of service can be created based on unique user attributes such as browser type or language.

Website managers can meter, control, and account for bandwidth use – by client, server farm, virtual service, application, user class, content type, and many other Layer 2 through 7 attributes.

Additionally, bandwidth contracts can be implemented to guarantee service rates and improve the transmission quality of delay- or loss-sensitive applications. For each traffic class, the website manager assigns a bandwidth policy composed of three bandwidth rates: a Committed Information Rate which is the guaranteed minimum bandwidth, a Soft Limit which is the metered rate used when bandwidth is available, and a Hard Limit which is the maximum burst rate.

Web OS-equipped switches provide QoS (Quality of Service) across the Internet by tagging traffic with IP ToS settings, allowing upstream routers and switches to transport traffic at the appropriate service levels.

Web OS logs a variety of flow statistics that can be used to track usage by bandwidth contracts, traffic classes, and bandwidth policies, providing data for accounting applications and capacity planning.

Multi-Level Resiliency Services Ensure Application Availability

Web OS enables non-stop content access in the event of network, switch, server, application, or even entire site failures. Alteon Web switches inherently support fully-meshed topologies - eliminating all system-wide single points of failure. Redundant switches can be deployed in an active-active configuration, with both switches simultaneously servicing users. Active/active redundancy improves resource availability, scalability, and network performance, minimizing user impact when failures occur. In the active standby mode, Stateful Failover allows the active Web switch to share session information, including SSL session state, HTTP cookie state, and FTP session state, with the standby switch to ensure session persistence.

Web OS-equipped switches constantly monitor server, application, and content availability, bypassing unhealthy servers and services when distributing new sessions.

Intelligent application health checking ensures integrity of the entire data path, including content retrieval for HTTP, SSL, POP3, SMTP, IMAP, NNTP, FTP, Radius and DNS services. In addition, "send/expect" script-based health checking allows users to dynamically verify content availability.

Effective Server Attack Protection

Web OS server security protects websites against malicious attacks and unwanted intrusion while providing continuous service for legitimate connection requests.

Web OS-equipped switches can thwart TCP SYN attacks without blocking legitimate session requests by terminating all TCP connection requests at the switch.

This is superior to the traditional approach, which forces a threshold for total connections at which the switch stops accepting TCP SYN requests, including those from legitimate users. Other Web OS server security measures include Network Address Translation (NAT) to allow implementation of private Internet addresses, access list filtering and secure switch administration.

Secure Switch Management

Web OS support for Secure Shell (SSH) and Secure Copy (SCP) protocols ensures that messages between a remote administrator and the switch use secure tunnels, protecting network data from IP spoofing, IP source routing, and interception of clear text passwords and other data by intermediate hosts.

APPLICATION EXAMPLES

- Traffic distribution across multiple Web servers and Web server farms
- Scaling firewall and Web server performance for secure e-Commerce sites
- Transparent caching infrastructure for ISPs
- Web, email, News, FTP, DNS, and Radius load balancing for ISP online services
- Host header parsing for simplified Virtual Hosting
- High-performance bandwidth management for usage-based and preferential services
- Multi-site load distribution for content publishers
- High-availability services for Web hosters

www.alteonwebsystems.com