W H I T E P A P E R

The Next Step In Server Load Balancing

- OVERVIEW
- BENEFITS
- APPLICATIONS
- SERVER LOAD BALANCING OPERATIONS
- ADVANCED FUNCTIONS
- HIGH AVAILABILITY CONFIGURATION
- WEB SWITCH ARCHITECTURE
- SUMMARY

Alteon WebSystems, Inc.

50 Great Oaks Boulevard

San Jose, California 95119

408-360-5500

408-360-5501 fax

http://www.alteon.com

In the past few years the World Wide Web has grown from a speculative medium to a robust infrastructure that handles mission-critical business traffic. According to IDC, the number of e-commerce users over the Web increased to 142 million in 1998 and is forecast to surpass 500 million by 2003. This increase in use will drive commerce on the Internet to more that $1 trillion by 2003. To support this growth web sites must be able to set up and maintain a few million connections every second.

Further, quick response time and continuous availability are mandatory requirements as sites compete to offer users the best on-line experience. What is required in all mission-critical Web-computing infrastructures is a solution that can scale server capacity dynamically to match aggregate client demand while ensuring continuous service availability.

Load balancing switches provide the infrastructure to scale application-processing power, maximize server efficiency and ensure high application availability. The new breed of switch-based server load-balancers consolidates multiple web infrastructure functions in addition to server load balancing and multi-layer switching. Examples of these functions include redirecting traffic to caches, load balancing traffic across multiple firewalls, packet filtering and bandwidth management. Alteon Web Systems coined the term “Web Switch” to represent this new class of device that front-ends server farms and provides Internet traffic control.

Web Switches use protocol information in Layers 3, 4 and above, such as URLs, TCP or UDP port numbers, the SYN/FIN bits that mark the start and end of TCP application sessions and IP source and destination addresses, to identify and manage application-layer sessions.

First generation server load-balancers are PC-based software products with limited performance and connectivity. To meet the rapid growth in traffic volume and server population, Web Switches offer many orders of magnitude of improvement in performance, connectivity, resiliency and economy.

OVERVIEW

Web Switches dynamically distribute load across a group of servers running a common application (or set of applications) while making the group appear as one server to the client.

For instance, a number of Web servers with access to the same content can be logically combined into an HTTP cluster or “Virtual Server Group (VSG)” (a group of servers that supports a common application or set of applications). The VSG provides a “virtual” HTTP service to clients. Clients are not aware that there are a number of real servers participating in providing this service. The clients access the service using a virtual IP (VIP) address that resides in a Web Switch that front-ends the real servers. As connection requests arrive for the virtual service, the Web Switch passes these requests on to one of the real servers in the VSG based upon knowledge of the servers’ availability, load handling capability, and present load.

A key part of server load balancing is session management. Once the Web Switch assigns a session to a real server, it must recognize all successive packets associated with that session. These packets are processed and forwarded appropriately to make sure that the client continues to be associated with the same physical server for the duration of the session.

Web Switches also monitor the completion of sessions at which time the binding of the session to the physical server can be removed. This ensures that the next time the client connects he will be connected to the most available server at the time, providing the best possible service to each client. The administrator can invoke special mechanisms if the application requires "persistent connections" (successive connections to be forwarded to the same physical server), such as with FTP control and data connections from the same client, SSL (Secure Sockets Layer), and persistent HTTP used for multi-page forms and search engines.

BENEFITS OF ALTEON WEB SWITCHES

Technical benefits

Reliability—Increases application availability by removing solution dependency on the availability of a single device in the web data center infrastructure such as a server, router, etc. The switch monitors availability of physical devices and on failure of any one device automatically routes new sessions to other healthy devices. Also, spreading active sessions across multiple devices decreases the number of sessions disrupted during a device failure.

With Alteon Web Switches, backup servers can be configured to come on-line if a server or server group fails. Web Switches, can support multi-site load balancing to protect against failure of the entire site. See the Alteon Web Systems white paper "Enhancing Web User Experience with Global Server Load Balancing" for more details.

Industry Leading Performance—Delivers the best switching performance in the industry. The state-of-the-art ASICs used in the switches integrate repetitive networking tasks in hardware, resulting in wire-speed throughput for Layer 2/3 packet switching and session switching at higher layers. Using Alteon Web Switches to distribute traffic can improve performance of business-critical web sites.

Performance Scalability—Provides the infrastructure to match resources with demand. Alteon Web Switches are designed to provide linear performance increases as more ports are deployed. During periods of peak load a switch farm can dynamically and gracefully increase the capacity of the VSG by enrolling overflow servers into the load balancing process.

Platform Independence—Provides application server platform independence. In a group or cluster, different types of servers—Unix, Windows NT, etc.—can be combined to host the same application. providing maximum flexibility.

Optimal Resource Utilization—Improves server resource utilization by distributing the traffic to multiple servers. Clients can be directed to the least utilized server, servers best configured to handle particular types of requests, etc. This decreases server costs and improves site performance.

Operational Simplicity—Allows devices to be removed for maintenance purposes, without disruption of user services. This decreases administration cost and "Total Cost of Ownership" (TCO).

Business benefits

Increased Customer Satisfaction—Customers get faster, more consistent response time, with the Web Switch directing traffic to the least loaded and most responsive web server. Web sites can consistently deliver superior performance by managing the load on web servers and preventing web servers from getting overloaded.

Improved Resource Utilization Based on Business Policies—Preferred customers and mission critical application traffic can be given higher priority by the Web Switch. Server and network resources can be allocated for high priority users and applications with the bandwidth management feature. Mission critical applications and users accessing these applications will get consistently good performance.

The switch offers features to differentiate users and guarantee a superior level of service for preferred users. Internet Service Providers and corporate IT managers can control the infrastructure resources to deliver negotiated Service Level Agreements. Also this gives them an opportunity to increase profit margins by offering tiered service levels.

Decreased TCO—The Web Switch helps in simplifying infrastructure architecture and optimizing resource utilization. This decreases administration and server cost.

APPLICATIONS

Environments that benefit from server load balancing include web hosting services, e-commerce services, on-line service providers and corporate data centers with high availability requirements. Server load balancing supports many TCP-based or UDP-based applications where common content is available across a group of servers. Internet/Intranet applications, such as Web servers, FTP servers, DNS servers and RADIUS servers have been the first to take advantage of server load balancing to support the high growth and unpredictable volume of Web-oriented traffic.

Web-Hosting and On-Line Services

Web hosters and on-line service providers typically deploy multiple HTTP, FTP and other application servers today, with load distributed across them, via round-robin DNS. This approach is undesirable because it is not fault-tolerant and requires a high degree of administration. Server load balancing enables transparent use of multiple servers with built-in high availability support.

E-commerce Web Sites

E-commerce sites are becoming mission critical for most businesses. For many new businesses the web is the only channel for generating revenue. These sites must have superior performance and high availability.

On-line Retail Sites

These sites experience sudden surges in user traffic – due to a promotion campaign, demand for a new product, etc. Web Switches help manage the load on the site, prevent servers from getting overloaded and guarantee high application availability. Also, it helps in easily scaling the site to meet demand with minimal effort.

On-line Trading Sites

These sites have more dynamic content and experience sudden traffic surges. For example, a change in political scenario or availability of new economic data could result in large fluctuations in the financial market. This results in a sudden, multi-fold increase in on-line trading activity. Also, when servers are overloaded, the performance degrades and session completion rate rapidly degrades. Planning and over provisioning to meet these sudden traffic surges is very expensive. Web Switches help to optimize resource utilization and dynamically control the load on servers to provide good performance.

Business to Business e-commerce Sites
Enterprises are setting up extranets to automate sales, marketing and procurement processes. These extranets must be highly available and provide consistent performance. Web Switches enable enterprises to build reliable and scalable infrastructures that support mission critical solutions.

Application Hosting Services
Application Service Providers (ASP) are using the Web to deliver applications on demand. This enables enterprises to rent applications and decrease their IT expenses. To meet the requirements of large enterprises, hosted solutions must be scalable, highly available and must deliver superior performance. A Web Switch enables ASPs, to build an infrastructure that supports these business critical solutions.

SERVER LOAD-BALANICNG OPERATIONS
Any TCP, UDP or IP based application can be load-balanced as long as common content is made available across the VSG supporting the application.

Grouping Options
Servers can be grouped into VSGs based on administrative requirements. Common ways to administer server content are duplicated content and centralized content back-end data servers.

Duplicated Content
Servers with duplicated content on local disks can be grouped into a VSG. This approach is low cost and ideal for applications with mostly static content requiring little replication.

Centralized Content on Back-End Data Server
Alternatively, data can be stored directly in a back-end database server. Application servers retrieve data from the back-end server in real time through file or database-sharing techniques such as a Network File System (NFS). Each application server runs any or all of the load-sharing applications as long as it has access to the associated content on the back-end server. Shared access to a single copy of data also enables server load balancing for read/write applications.

To allow the data server to keep up with aggregate requests from multiple application servers, a general recommendation is to connect back-end data servers with a higher-speed switch port than the application servers. The same Web Switch platform allows clients to access the application servers via layer 4 switching and allows the application servers to access the database server via layer 2 or layer 3 switching.

Operation of Alteon Web Switches
The Alteon 180, ACEdirector and Alteon 700 families of Web Switches integrate application layer session management and packet switching to offer concurrent multi-layer (2/3/4-7) switching, local and global server load balancing and policy-based redirection services.

The Web Switch, with server load balancing, acts as a virtual front-end processor to clusters of real servers connected via direct attachment to switch ports or indirectly through hubs and switches. A VIP address is configured for each VSG, presenting a single address for the server group to the rest of the network. Domain Name Services (DNS) must advertise this VIP address.

For example, if www.alteon.com were available on three servers and load-balanced by the Web Switch with a VIP of 215.176.12.3, DNS would advertise 215.176.12.3 as the address of www.alteon.com, not the individual IP addresses of the real servers.

Clients who wish to access the load-balanced application will be directed by DNS to send their requests to the VSG hosting the application. The Web Switch receives these packets and uses the destination TCP port number to identify the application that is to be load balanced.

Packets addressed to anything other than the configured VIPs on the Web Switch are forwarded at layer-2 or layer-3 as appropriate. This allows real servers to be accessed through the Web Switch simultaneously via their real IP addresses. This also enables non-load balanced applications such as system management and administration to access the servers directly through the switch.

TCP/IP Server Load-Balancing Operation

The Web Switch recognizes when a client is requesting a new TCP session by identifying the TCP SYN packet. The request is forwarded to the best available server based on the configured load balancing policy. Once the switch determines the best server, it binds the session to that server's real IP address. The Web Switch maintains a binding table that associates each active session with the real server to which it is assigned.

After the Web Switch binds a connection request to a real server, it performs address substitution so that the real server will transparently receive packets for that session. The switch replaces the VIP in the IP destination address with the server's real IP address and replaces the switch's MAC address in the MAC destination address field with the server's MAC address.

After performing the necessary address substitution, the Web Switch forwards the connection request to the chosen server. All subsequent packets belonging to that session undergo the same address substitution process and are forwarded to the same real server until the switch sees a session termination packet (i.e., a TCP FIN packet).

Likewise, the Web Switch intercepts packets traveling from the real server to the client and performs the reverse address substitution. It replaces the real server's actual IP address in the Network Layer source address field with the VIP and forwards each modified frame to the client (see Figure 1).

Upon receipt of a TCP FIN packet, the Web Switch performs the necessary address substitution and forwards the FIN packet to the appropriate real server, causing the server to tear down the connection. Then it removes the session-server binding from its binding table.

UDP/IP Server Load-Balancing Operation

Because UDP is a connectionless protocol, it technically does not support the concept of session. Nevertheless, Alteon Web Switches support server load balancing for UDP-based applications.

If UDP load balancing is activated on the Web Switch, the switch identifies UDP packets arriving from clients addressed to a VIP. If an incoming packet destined for a VIP has a source IP address that is not presently found in the Web Switch's binding table, the Web Switch uses the configured load-balancing algorithm to determine to which real server the “new session” should be bound.

Figure One

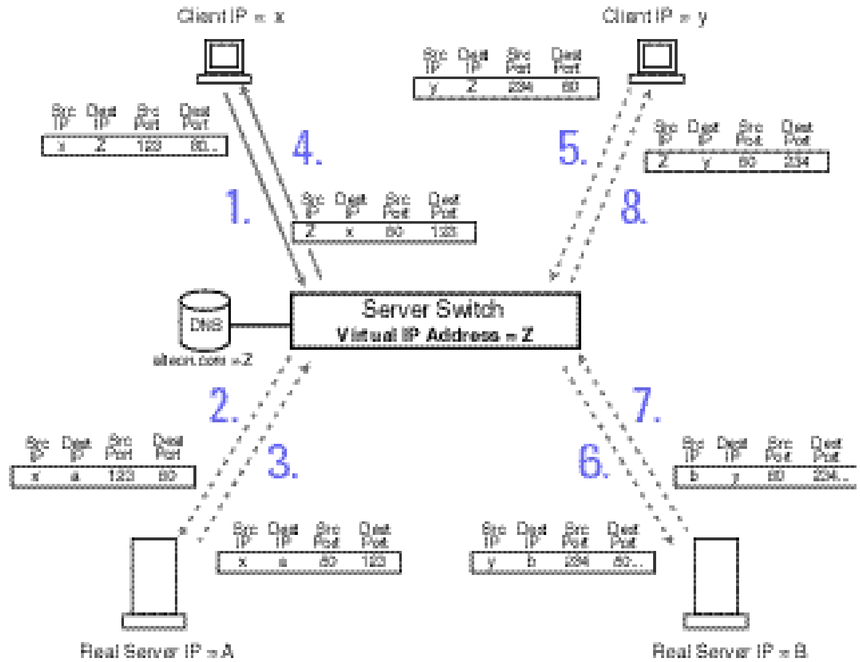

While that binding is in effect, all UDP packets from that source IP address will be sent to that same real server. When no UDP traffic from the source IP address is seen in a user-specified period of time, the “session” is removed from the switch’s binding table. In this context, “session” means the flow of packets during the time that a specific Source IP address is bound to a particular real server.

Real Server Health Monitoring

The most important consideration when making load-balancing decisions is whether or not the servers and their load-balanced applications are working properly. The Web Switch uses special mechanisms to monitor servers and ensure that connection requests are directed to healthy servers. The different mechanisms available for health checking are physical connection monitoring, ICMP (Ping) monitoring, TCP connection monitoring, active content verification and dynamic application verification.

Physical Connection Monitoring

The Web Switch monitors the physical link status of switch ports connected to real servers. If the physical link to a real server goes down, the Web Switch immediately places that server in the “Server Failed” state and does not forward new connection requests to that server. In addition, the switch does not forward any more traffic associated with existing sessions to that real server and immediately removes all existing sessions bound to that real server from its binding table. This ensures that existing connections are closed quickly with minimal impact to the user.

ICMP Monitoring

At the most basic level, the Web Switch can monitor server health by sending ICMP ping requests to servers. Based on the server’s, the switch can verify the health of the server’s networking stack. This is recommended when all IP traffic or a UDP-based application other than DNS is being load balanced.

The rate at which ICMP Pings are sent is a user-configurable parameter. If a user-configurable number of consecutive Ping requests fail, the target server is placed in the "Server Failed" state. While in this state, no new connection requests are sent to the real server.

TCP Connection Monitoring

A more sophisticated health checking mechanism is for the Web Switch to send TCP connection requests (i.e., TCP SYN requests) to the real servers and determine that the server responds. This is recommended when a TCP-based application for which the switch does not support content-based health checks is being load balanced.

A more sophisticated health checking mechanism is for the Web Switch to send TCP connection requests (i.e., TCP SYN requests) for each load-balanced application to each real server and determine that the server responds. These connection requests identify both failed servers and failed services on a healthy server. The rate at which these connection requests are sent is a user-configurable parameter.

If a user-configurable number of consecutive connection requests to an application fail, the application on target server is placed in the "Service Failed" state. While in this state, no new connection requests are sent to the application on that real server. However, connection requests continue to be sent to other, healthy applications running on the same server.

When an application on a real server is in the "Service Failed" state, the Web Switch continues to perform health checks. When the Web Switch has been successful in connecting to a load-balanced service, the real server is slowly brought back into service, using mechanisms that don't overwhelm the newly available server.

Active Content Verification (ACV)

To verify the availability of a Web service and associated content, users can configure a reference URL to instruct the Web Switch to access the data represented by the URL during its periodic server health checks.

To check HTTP service availability, the Web Switch makes a request for the specific content via an HTTP GET and then verifies the received content and return code. This test not only checks the TCP connection, but also checks the web server and content server. In the example in figure 2, the HTTP request is sent through Web Server A and File Server A. If the verification request is successful, it proves that there is a path from the switch to the data. If the verification request failed, the switch redirects the request through another path.

The ACV feature supports not only HTTP but also other services including NNTP, FTP, SMTP, POP3, IMAP, DNS and RADIUS.

Dynamic Application Verification (DAV)

Dynamic application verification expands upon Active Content Verification. With Dynamic Application Verification, the Web Switch can verify the availability of dynamic applications like .asp applications, cgi-scripts, forms, etc. on web servers. This feature allows users to write their own tests that emulate user actions and supply the switch with a URL of the test script.

The DAV feature can also be used to monitor the health of other Internet services like mail service.

Figure Two

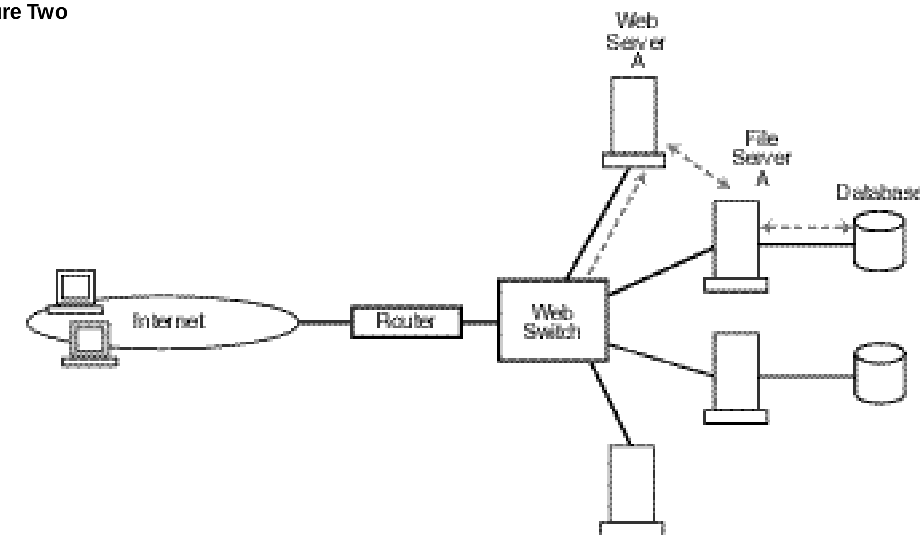

Backup Servers
A server or link failure, as well as administrative action, can cause a Web Switch to remove a server from the load-balancing VSG. Users may configure the Web Switch to introduce a backup server into a VSG when any or all of the servers within the VSG fail.

A backup server introduced into the VSG after the removal of a single server processes connections until the Web Switch determines that the removed server is once again operational. At that point, the backup server is removed from the mix.

In the case where the switch introduces the backup server into the mix after the removal of all servers in a VSG, the backup server is taken out of the VSG when the Web Switch determines that all servers in the group are once again operational. Note that a backup server being taken out of the mix means that no new connections are made to it. Existing connections remain operational until they are terminated by normal mechanisms.

Load Balancing Policies
The Web Switch makes decisions regarding which server within a VSG to assign a new connection based on user configured load-balancing policies. The different policy options available in the Alteon Web Switch are:

- Least utilized or "most available" server selection
- Persistence-based server selection
- URL-based load balancing

"Most Available" Server Policies
Various policies can be used to drive the load balancing decision. On the Web Switch, the load balancing policy for each VSG is a user configurable parameter. "Most Available" policies include:

- Simple Least Connections
- Weighted Least Connections
- Simple Round-Robin
- Weighted Round-Robin

In addition, a maximum-connections threshold can be configured for any or all real servers in conjunction with any of the above policies.

Simple Least Connections

In simple least connections, the number of active connections being processed by each real server is tracked. When a request for a new connection is received, it is forwarded to the real server in the VSG with the fewest active connections. This is often viewed as the fairest policy because servers that close connections faster get more connection requests forwarded to them over time. Implicitly, this algorithm takes into account feedback from the real servers in that faster or less loaded servers will close connections more quickly, indicating that they are available for more workload.

Weighted Least Connections

In weighted least connections, a weighting function is added to the simple least connections policy. The number of connections to each server is normalized based on each server's static weighting and each connection request is directed to the server with the fewest active normalized (as opposed to actual) connections. The idea is that servers with greater inherent capacity should support a larger number of active connections.

Simple Round Robin

With this algorithm, new connection requests are forwarded to the real servers in a round robin fashion such that, over time, each server in the application server group gets the same number of connections. This doesn't mean that each server will have the same number of active connections, because some will close connections faster than others. Round robin is commonly used when the servers in a VSG are roughly equal in capacity.

Weighted Round Robin

Weighted round robin load balancing is similar to round robin load balancing, but each server in the application VSG is assigned a static weight based on some view of the capacity of each server. Servers are presented connection requests in proportion to their weighting.

Persistence Policies

Many e-commerce sites maintain stateful information for each customer on Web servers. It is important for the Web Switch to forward all incoming requests from a particular user to the same server, until the completion of the transaction. Also, users returning to an e-commerce session after a temporary disconnect, should be redirected to the same server to avoid having to restart their shopping activity.

The persistence policies that are supported on the Alteon Web Switches are:

- Hash
- Minimum Misses
- SSL Session ID Tracking
- Cookie-Based Tracking

Hash
With this policy, the server is chosen based on source IP address. The IP address is used to generate an index into a table that contains all servers in VSG. Since server selection is based on source IP, all requests from a given user are sent to the same server. This is particularly useful in e-commerce applications and firewall load balancing where session state must be maintained

With this policy, the server selection table in recomputed every time a server leaves or enters the VSG. No existing connections are rehashed unless server they go to has gone down.

The hash policy is better than minimum misses (see below) for most applications, as it generally offers a more even distribution of connections across servers. If server load balancing statistics indicate that using hash causes one server to processing more requests over time than others, consider using minimum misses, SSL session ID tracking or cookie-based tracking. One of the biggest challenges with the hash algorithm (as well as min misses) is that all clients coming from behind a proxy may use the same source IP address and therefore cannot be differentiated. In this case, traffic from all clients is sent to the same server, defeating load balancing.

Minimum Misses
This policy is similar to hash in that the source IP address is used to generate an index into a table that contains all servers in VSG. However, unlike hash, when a server leaves the VSG, the server selection table is recomputed only to reassign users associated with the failed server. This policy results in less perturbation of server assignments with a server fails.

SSL Session Tracking
Many e-commerce sites use secure connections for transporting private information about clients. When a client connects to a server using an encrypted SSL session, a unique SSL session ID is assigned. Using SSL session IDs to maintain server persistence is the most accurate way to bind all of a client's connections during an SSL session to the same server.

SSL session tracking is most useful in scenarios where other persistence policies cannot be used because one source IP address (belonging to a proxy) is used to represent a large number of clients. For example, in large networks, proxy firewalls typically change users' IP addresses on outgoing packets for security reasons, overloading the source IP address. If persistence is determined solely by source IP address, many users will be redirected to the same real server and load balancing in the VSG will be ineffective. In these environments, using SSL Session ID persistence for load balancing is the best algorithm.

When SSL session ID tracking is enabled, Alteon Web Switches examine the TCP SYN handshake and subsequent packets to examine the SSL session ID and determine if it belongs to an existing SSL session or a new one. If the session is new, the Web Switch assigns it to a real server based on the configured load balancing algorithm (least connections or round robin). If the packets are associated with an existing session, the connection is assigned to the same server that was involved in previous portions of the SSL session.

Cookie-based Session Tracking

This algorithm uses HTTP cookie information to direct the client connection to the appropriate server. This load balancing mechanism offers more granularity than IP address, because the switch can now identify a specific user to send to a server.

With cookie-based session tracking, the Web Switch sends the first incoming request to the most available server. The server modifies the cookie and inserts its IP address. Based on this information, any subsequent request from this user is forwarded to the same server.

Maximum Connections Option

The maximum connections feature, configurable in conjunction with any of the load balancing options described above, allows users to set the maximum number of active connections to be assigned to a particular server.

When a server reaches its maximum connection limit, it will receive no more connections until it drops back below its maximum connections limit. If an overflow server has been assigned it will be brought into service. If all of the servers in the VSG reach their maximum connections limit and no overflow servers are available, no additional connection setup requests will be handled until at least one server drops below its limit.

The maximum connections option enables users to enforce service quality by preventing key servers from being overloaded.

URL Based Load Balancing

URL-based load balancing allows the Web Switch to observe each client URL request and forward the request to the appropriate server based on predefined rules. Typically every page download consists of multiple TCP requests. URL based load-balancing looks into incoming HTTP requests and, based on the URL information, forwards the request to the appropriate VSG. Server selection within a VSG is based on user defined polices and dynamic load on the server.

The example in Figure 3 shows a large e-publishing site that has graphic images (.gif and .jpg) and script files (.cgi, .bin and .exe) on separate servers. Also, the static files are stored on a separate server farm under the /product, /company and /information directories. The Web Switch observes the incoming HTTP request and forwards it to the appropriate server based on user defined policies – I/O intensive requests, like reading image files, video files, etc. are sent to VSG A and compute intensive tasks, like executing applications, are sent to VSG B. Based on this configured policy , all requests for images (.GIF) files or (.JPG) files are sent to VSG A. The requests are also load balanced among servers in VSG A, based on a user-configured "Most Available " server policy. All requests with .bin, .exe in the URL are sent to VSG B. All other requests for static pages are sent to VSG C.

Figure Three

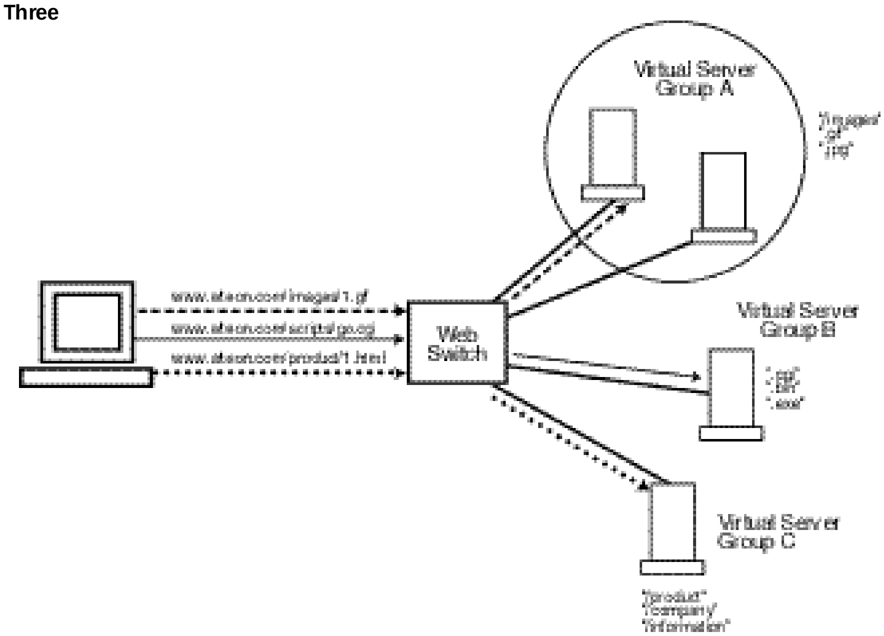

Some of the benefits of using this feature for large web site managers are:

Ease of Management—Large web sites have so much content associated with their domain name that splitting it up across multiple file systems on multiple servers makes it easier to manage, improves site performance and availability. With URL based load balancing incoming requests can be easily redirected to the server group with the correct content.

Resource Optimization—Partitioning content across multiple servers enables web sites to use servers optimized for a specific task. Servers that are optimized for compute intensive functions are better at serving web pages with dynamic content, while servers optimized for disk I/O functions are better at delivering large static pages. With URL based load balancing, the switch forwards dynamic requests to powerful server groups optimized to handle compute intensive requests and static requests to I/O optimized servers. This improves performance and optimizes resource utilization.

User Differentiation and Service Level Provision—Partitioning users among different Virtual Servers Groups gives Web site managers the ability to use different types of servers to service different classes of users. URL parsing helps in forwarding incoming clients to different servers based on user type, type of request etc. This feature can be used to give preferred customers a larger share of system resources and better response time than best effort users.

Prioritization of Transactions—URL based load balancing can also be used to give certain types of transactions higher priority, so a payment transaction, or an order transaction can be given a larger share of Virtual Server resources.

Parallel Execution—Large Web sites have their content partitioned across multiple servers. With this feature, processing tasks for each session and page request can be split and distributed to multiple servers for parallel operations. For example we can retrieve the image and text portion simultaneously from different servers. Parallel processing improves session performance and scalability

Management of Traffic Surges

The Web Switch can be configured to introduce overflow servers when any or all of the servers in a VSG hit their maximum connections threshold. It can also bring in backup servers when it detects that any or all of the servers in a VSG have suffered from a service or physical link failure.

A server that normally performs another function can be configured with the load-balanced application(s) and given access to the associated data. Under normal circumstances, the Web Switch will not forward any connections for the load-balanced application to this server.

However, if one or all of the real servers supporting the load-balanced application hit their maximum connections limits, the Web Switch introduces the overflow server into the load balancing virtual service group. When the load on the real server(s) falls to an appropriate level, the Web Switch stops forwarding connection requests to the overflow server, removing it from the load balancing service group.

HIGH AVAILABILITY CONFIGURATION

The Web Switch monitors the real servers and the load balanced applications to ensure that client requests are forwarded only to healthy servers. To further improve fault tolerance for the entire system, Web Switches can be put into active-active load balancing configurations to build topologies with no system-wide single point of failure.

The main advantages of an active-active configuration are:

Improved Resource Utilization—Both Web Switches are actively performing load balancing activities for the same virtual service. This doubles the virtual service's effective usable performance and capacity.

Improved Session Availability—Since both Web Switches are active, if one fails, only half of the user sessions to the Web server farm are affected. With traditional hot-standby configurations, only one of the switches is being used and when it fails, all the user sessions to the web server farm are lost.

Ease of Configuration—Since both switches are active all the time, it is much easier to configure the switches in this mode.

Improved Performance During Traffic Surges—Since both switches are being used, any sudden surges in traffic load are handled more easily.

More details about active-active configurations are available in the "Increasing Web Site Performance and Availability with Active/Active Web Switching" white paper.

WEB SWITCH ARCHITECTURE

Load balancing thousands or tens of thousands of connections per second over dozens of real servers requires vast amounts of processing capacity and memory. Traditionally, three types of processing take place in a server load-balancer: per-frame, per-session and background processing.

Background processing includes tasks such as server and service health checks, topology communications and statistics reporting. These tasks tend to be single-threaded and the load is light. A co-processor inside a switch is able to do the job.

Per-session processing includes a variety of activities such as identifying incoming session requests, determining the best server for each session and creating a binding, recognizing session terminations and unbinding each session, handling persistence and timing out idle sessions. Session processing is CPU-intensive, particularly when traffic to be load-balanced arrives simultaneously from many 100 Mbps or 1Gbps ports. Executing per-session processing from a centralized processor limits number of total session throughput of the load-balancer. In addition, routing each session packet to and from the centralized processor puts an additional load on the backplane of any server load balancing device.

Finally, per-frame processing includes session address substitution (IP, MAC and Layer 4 port number), recomputing checksums, and in some cases, examining actual frame data. For example, load balancing FTP sessions requires that the load-balancer examine every packet sent by an FTP server to the client over the control channel so that it can derive the dynamic port number that will be used for data transmission. Per-frame processing can be extremely CPU-intensive and is best distributed to dedicated processors on each data path. Executing per-frame processing from a centralized processor on traditional server load-balancers limits the total frame throughput of the device.

There are many other applications that have unique session establishment characteristics. This further heightens the need to equip a load-balancer with enough capacity and flexibility to support future application load-balancing needs.

Web Switches from Alteon Web Systems use a distributed processing architecture that is ideally designed for processor-intensive session management. Each port on the Web Switch integrates a switching ASIC that consists of a hardware-assisted forwarding engine and dual, 180-MHz RISC processors. Two additional RISC processors provide support for switch-wide management functions (see Figure 4).

Figure Four

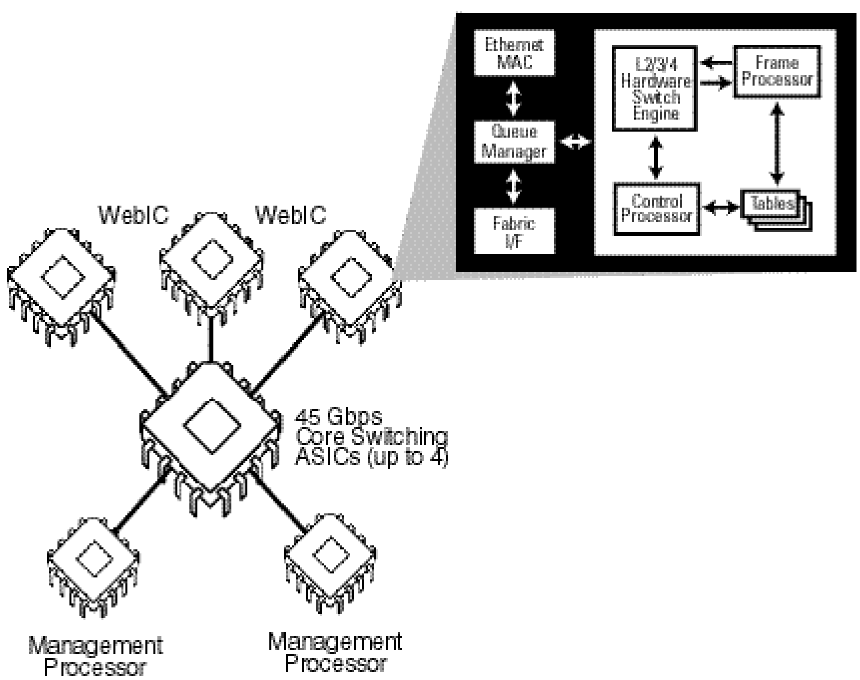

On each switch port, the forwarding engine in the switching ASIC handles packet forwarding with hardware support for header insertion and removal, TCP/UDP/IP checksum computation, queuing and statistics management.

The RISC processors handle session state management, server assignment, session binding and unbinding, session address substitution and when necessary, packet data examination. The central management processors handle background tasks such as server and service health monitoring, load tracking and bringing servers in and out of service.

With this architecture, processing tasks for each session are distributed to different processors for parallel operations. For example, the RISC processors on the client ports handle session binding and unbinding while the processors on the server ports take care of session address substitution. The management processors synchronize all tables, including the binding tables and server status tables in real time. Parallel processing increases session performance and scales the Web Switch's load balancing capacity with port density.

A proof-point for this architecture can be found in a performance test conducted by Network Computing in October 1998. The performance comparison chart in Figure 5 shows the ACEswitch 180 performance is much better than the competitors at low traffic bandwidth and the only one that can support sites that require high traffic bandwidth.

SUMMARY

In the past few years the number of users accessing web sites has grown rapidly. Also enterprises are using their web sites to increase revenue, improve customer service and optimize business processes. Server load balancing enables IT manager's to build an infrastructure that can support mission critical solutions. For web-based solutions, the switch provides – high availability, performance scalability, platform independence, resource optimization and operational simplicity. The business benefits of a Web Switch based solution are increased customer satisfaction, improved resource utilization and decreased total cost of ownership.

The features of Web Switches that are important for any web solution are layer 3 & 4 load balancing, URL based load balancing, real server health monitoring, peak load or traffic surge management, and policy based service level management.

Alteon's Web Switch provides the capacity, flexibility, reliability and ease of use to accommodate the rapid growth and mission critical requirements of large Web sites.