



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Firewall Load Balancers

Megan Restuccia

November 21, 2000

The need for high availability has led the industry to focus on eliminating any single points of failure in their network. A very common single point of failure is the firewall server. Many times firewalls will have a limited traffic load capacity, thus having a single firewall may lead to interruption of connectivity when upgrades or service need to be completed. One solution to this problem would be implementing firewall load balancers with multiple firewall servers.

Load balancers are hardware or software devices that control the flow of packets to the machines for which they are balancing the load. Load balancers can automatically detect when a server becomes unstable or unavailable and control the flow of traffic accordingly. This type of flow management also enables growth in the environment, by allowing more servers to be added transparently, as needed to meet growth requirements.

Not all load balancers are the same, and the types of firewalls they each support vary as well. The best choice will depend greatly on the design of the network it will be introduced into, and is beyond the scope of this paper.

Based on the functionality of the firewall, they can traditionally be separated into three categories: packet filters, stateful inspection and application-level/proxy servers. Packet filter firewalls examine IP packets and make a decision based on specific criteria. The criteria are generally the source and destination IP addresses and source and destination TCP and UDP port numbers. The packet filter firewall looks solely at the header information of a packet and not at the data contained within the packet. Stateful inspection firewalls examine the packet headers also, but they also look at some of the contents by transparently intercepting the packet and reviewing the contents. The criteria this firewall uses is based upon the packet header information and the data within the packet. The final category of firewalls, the application-level gateway/proxy, accepts and initiates connections on behalf of the LAN clients. Because of the nature of this type of server, requiring the initiation of two sessions for each request (one to the host and one to the client), the firewall can examine all of the data within a packet.

No matter which of category of firewalls is chosen, an added layer of complexity is brought to light when there are multiple firewalls. The multiple firewalls can either be synchronous or asynchronous. Synchronous firewalls allow either directional traffic to flow through multiple firewalls by exchanging connection information. Because of this sharing, the connection does not need to be revalidated each time a different firewall is used. Asynchronous firewalls do not exchange information, and the connection must be revalidated each time a new firewall is used. When multiple active firewalls are installed, it is important to understand its function in the network and the possible consequences.

Firewall load balancers can help to alleviate some of these issues. Load balancers can be hardware or software based. Each of these types have their own set of pros and cons as well as a common set of advantages and disadvantages with regards to the network.

Software solutions typically involve additional software to be installed on each of the servers in the cluster. There is also the possibility that a standalone machine, with the load balancing software loaded, must act as a master in the server cluster. Installing one or more pieces of software on each server can certainly have advantages, such as in-depth analysis of the operating system or the ability to synchronize multiple servers in the cluster. They can give the administrators an impressive ability to look into vital statistics on each server including CPU usage and memory utilization.

However, there are several disadvantages to the various software solutions. An unknown amount of resources could be used by the installation of another piece of software. Resource allocation might also be affected by the running of multiple software products on each of the servers. Server power could be depleted for the load balancing tasks if the implementation is task intensive. A new piece of software could also present an extra point of failure on each server. Since every new server that is added to the cluster must have the software installed, the software solution could limit scalability. The software solution could also potentially have an operating system dependency. Compounding the possible licensing costs or software adjustments and the possible complications of upgrades, a software solution could possibly eliminate a good growth plan.

The hardware solutions offered through the various vendors vary greatly, although they also consistently maintain specific characteristics. A hardware solution is a physical box that handles the flow and directing of traffic between the clients and the servers. Since a hardware solution is typically a standalone machine, there is very little change in the server environment. The server applications can continue to process data as before, while the load balancing hardware handles the traffic management. Because the solution is a standalone machine, it would typically allow for

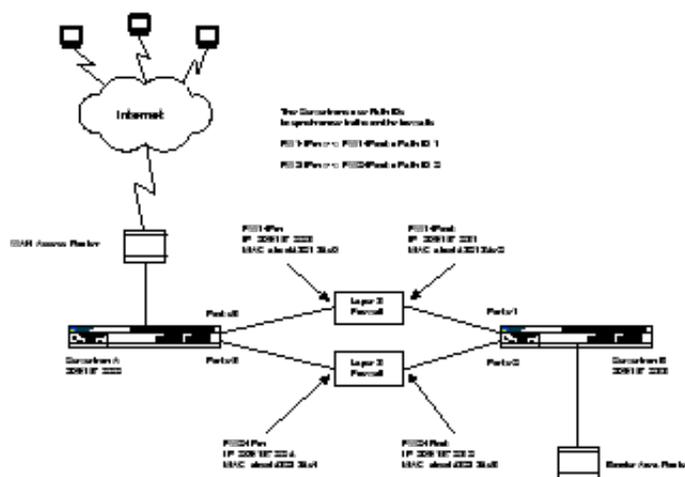
operating system independence and for a heterogeneous server environment. Since there would be only a single device, upgrades and maintenance would be much easier also. Zero downtime for maintenance could be achieved by using redundant hardware load balancing devices, which showcases the flexibility of this solution.

Although configuring a redundant unit can eliminate certain disadvantages with the hardware solution, it does not alleviate all of them. Hardware solutions are limited in the information they have about the servers. Information like CPU usage and memory utilization is difficult, if not impossible, to obtain without software counterparts loaded on each server. Since all traffic to and from the servers go through a single piece of hardware, it could be considered a single point of failure, unless a redundant unit is properly configured.

Even though both the software and hardware solutions have independent disadvantages, they do offer several advantages when implemented. Load balancers overcome the limited traffic load capacity on a firewall machine. They also make it easier to add firewall servers to the environment without adding complexity. If configured correctly, they can eliminate a single point of failure at the firewall level. Load balancers ensure the highest degree of availability by implementing redundancy and fault tolerance. Finally, both of these options offer an effective growth plan, as long as the solution chosen fits the environment. Load balancers allow any single or multiple servers to be taken out of the server pool, for routine maintenance or upgrades, with no interruption of service. These load balancers can also be configured to work with both asynchronous and synchronous firewalls, by maintaining session and client tables independently for each server. Most firewall load balancing solutions will also work well with NAT (Network Address Translation) firewalls, firewalls that use dynamic or static routing and even layer 2 firewalls, although there may be special design or configuration considerations for each of these types.

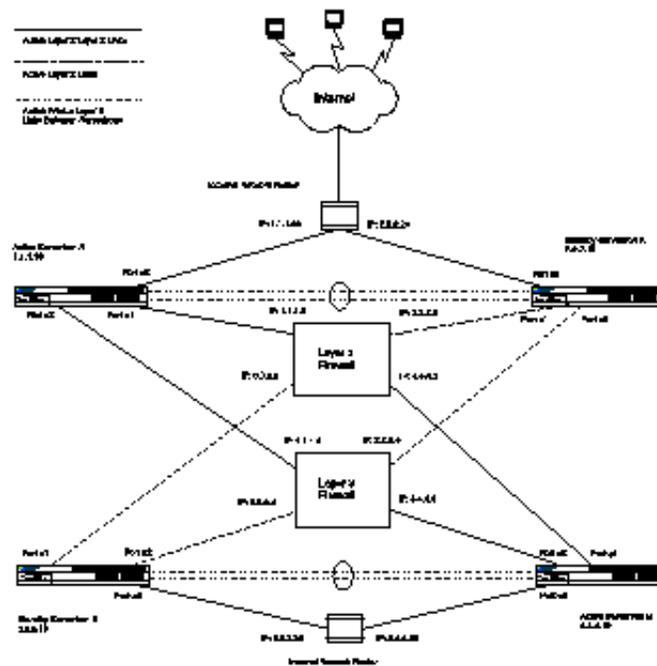
There are a few disadvantages to implementing firewall load balancing solutions. Although these solutions offer high availability and redundancy, they do require extra layers of hardware or software leading to a more complex environment. Also, there is no standard of communication between the various load balancer vendors, making it difficult for different solutions to work together. Finally, careful thought should be given to the placement of the firewall servers and load balancing solution, as a flaw in design or misconfiguration could lead to security problems. Even though these disadvantages could cause serious problems, there are many solutions to choose from to minimize these problems.

Although the examples shown within are specifically exhibiting the hardware solutions, they can be easily adapted for the software solutions. The idea is to show an example configuration for the purpose of explaining the benefits, not specifics of configuration. The basic configuration does not implement load balancer redundancy, but it does demonstrate the advantages.



Basic Configuration, provided by Foundry Networks

The high availability example displays the paths incoming and outgoing packets may travel in this environment.



High Availability Configuration, provided by Foundry Networks

Firewall load balancing servers can increase the security of a network by allowing upgrades and maintenance to be a less painful process, and providing constant uptime. Some load balancing solutions have built-in anti-DOS-attack measures, or use encrypted protocols or proprietary checksums. Administrators no longer have to worry when the firewall becomes overloaded with traffic, as the load balancers will distribute the load across all the servers. Load balancers also have a higher throughput and performance capacity.

It is very important to determine the right solution for a specific environment, especially because there are so many different choices available. In the high availability environment, all single points of failure need to be eliminated. Firewall load balancing solutions should be implemented for the purpose of eliminating these failures, as well as for the benefit of increasing redundancy and fault tolerance in the network.

Bibliography

Foundry Networks. Chapter 13 Configuring Firewall Load Balancing.

URL: http://www.foundrynet.com/techdocs/SI/Foundry_ServerIron_Firewall_Load_Balancing.html

Radware. Building Bullet Proof Internet/Intranet Sites with IP Load Balancing.

URL: <http://www.radware.com/content/support/whitepaper/bp.htm>

Radware. High Availability Security Solutions.

URL: <http://www.radware.com/archive/pdfs/whitepapers/highaval.pdf>

Radware. Technical Application Note 1055 - Introduction to FireProof.

URL: <http://www.radware.com/archive/pdfs/whitepapers/app1055.pdf>

Radware. What to Look for in IP Load Balancers.

URL: <http://www.radware.com/archive/pdfs/whitepapers/look.pdf>

Dan Strom, SANS Reading Room. The Packet Filter: A Basic Network Security Tool.

September 25, 2000. URL: http://www.sans.org/infosecFAQ/packet_filter.htm

Gregory Yerxa, Network Computing. Firewall & Load-Balancer: Perfect Union?

February 7, 2000. URL: <http://www.networkcomputing.com/1102/1102ws1.html>