



SENTINEL RIDGE CONSULTING GROUP INC.

“YOUR COMPUTER SECURITY EXPERT”

A Basic Primer On Firewalls

Ian Parker, GCUX, GCFW, CCSA, CSSE
President
Sentinel Ridge Consulting Group, Inc.

1 Introduction

An enterprise firewall prevents unauthorized connections to or from a computer network. Firewalls are typically found at the perimeter of an organization's network, directly behind the router that connects the organization to the Internet. However, firewalls may also be used internally to segment critical parts of the network. For example, a firewall may be used to restrict access to the servers that run the organization's financial applications.

This article explores the various firewall technologies on the market today, discusses the pros and cons of these technologies, presents some purchasing considerations and provides guidance on maintaining a firewall once it has been deployed.

Future articles will discuss concepts such as firewall virtualization and unified threat management.

2 Firewall Technologies

This section describes the main firewall technologies in use today. From a purchasing standpoint, customers should realize that most advanced firewalls combine two or more of these technologies.

2.1 *Packet Filtering*

Packet filtering is the oldest and most basic firewall technology. A packet filtering firewall operates like a router with an access control list. Packet filtering treats each packet as a discrete entity and decisions on whether to allow or block the packet are made accordingly. These decisions take into account such attributes as the source and destination IP addresses, network or transport protocol, source and destination ports, the firewall interface on which the packet appears and the direction of the packet.

In a perfect world, packet filtering technology would be sufficient to maintain the security of an organization's network. However, the ease in which TCP/IP packets can be manipulated (generally referred to as packet crafting) makes packet filtering alone susceptible to all kinds of funny business, such as spoofed source addresses, odd packet fragmentation and manipulation of TCP flags.

For example, packet filtering alone is unable to determine whether a packet with the ACK bit set is part of an existing connection or is in fact a malicious packet that has had the ACK bit set manually to fool the firewall into passing the packet.

Despite its drawbacks, packet filtering is still very useful for simple policy enforcement, such as blocking all NetBIOS connection attempts into an organization's network, as such decisions do not depend on the state of the connection. However, a perimeter firewall will often delegate such filtering rules to the border router, thereby freeing up the firewall to concentrate on more complex policy enforcement.

2.2 Stateful Inspection

Stateful inspection technology attempts to overcome the fundamental limitations of packet filtering by tracking the state of every packet that crosses the firewall. This is accomplished by saving various attributes of a packet into memory resident state tables. In this way, the firewall can determine with a high degree of accuracy whether a particular packet belongs to an existing connection or is a new connection attempt, authorized or otherwise.

Although simple in concept, stateful inspection can be tricky to perform for certain complex protocols. A further drawback to the technology is the memory required to maintain the various state tables in cases where thousands of simultaneous connections are crossing the firewall. Finally, it is important to realize that the inspection is still occurring primarily at the network and transport layers of the packet. What happens at the application layer, where more and more attacks are directed, is still largely invisible.

2.3 Stateful Protocol Analysis

Stateful protocol analysis takes the stateful inspection concept a step further by inspecting the packet at the application layer also, a feature some vendors refer to as deep packet inspection. Traditionally, this kind of inspection has been the domain of intrusion detection systems and, in fact, stateful protocol analysis is nothing more than a marriage of firewall and intrusion detection technology. Currently, dedicated intrusion detection or prevention systems offer more comprehensive protection than firewalls with deep packet inspection capabilities but this will probably change as the latter technology matures.

Stateful protocol analysis can be used to block malware contained within the data portion of a packet, such as a buffer overflow attack. It can also be used for very fine-grained policy enforcement, such as preventing certain types of e-mail attachments from entering the network, blocking ActiveX or Java content in Web pages or disallowing certain types of commands within an application, such as FTP Put commands.

This technology can also be used to ensure that certain protocols are used in compliance with RFC standards. However, this capability must be used with caution, as software vendors interpret certain standards differently and don't always follow the RFC's exactly in their protocol implementations.

Stateful protocol analysis requires even more memory than stateful inspection and certain types of analysis can also be very processor intensive.

2.4 Application-Proxy Gateways

The technologies discussed so far are alike in the sense that, in all cases, a direct connection is made between a source and destination host and the firewall merely inspects the traffic as it passes through it. An application-proxy gateway, otherwise known simply as a proxy firewall, operates in a fundamentally different way. A proxy firewall contains one or more agents that act as intermediaries between the source and destination hosts. All connections from either side are made to the firewall and never directly to the other host.

Proponents of proxy firewalls argue that this lack of direct connectivity between the hosts provides an inherently higher level of security than either packet filtering or stateful inspection.

While this is arguably true, it is offset somewhat by the fact that, as connections are being made to the firewall, the firewall itself, if not properly protected, can be attacked. As I discuss later in this article, it is always necessary to harden the operating system of the host on which the firewall software is running. However, this is especially true in the case of proxy firewalls.

Proxy firewalls suffer from three fundamental weaknesses. These weaknesses have served to lessen the popularity of this technology over the years.

The first weakness arises from the fact that every application that needs to communicate through the firewall needs to be “proxy-aware”. In other words, the client side of the application needs to be able to communicate with a proxy server as an intermediary, rather than the target host. Related to this requirement is the fact that the firewall needs to have an agent for every required protocol. These limitations at best make scalability an issue and may even render the firewall unusable at some point.

The second weakness derives from the adverse performance impact of having a proxy agent mediate every connection.

Finally, in cases where the proxy firewall must intervene in a “non-transparent” manner, the user experience is negatively impacted.

Due to these weaknesses, proxy firewalls have become a rarity. Instead, the marketplace now favours proxy servers, which are devices that retain the proxy functionality but forgo the firewalled capability. Proxy servers will be discussed fully in a future article.

3 Purchasing Considerations

For most companies, the most effective security perimeter for Internet connectivity would include a stateful inspection firewall, an outbound proxy server and a border router that can perform basic packet filtering. If the organization maintains a demilitarized zone containing publicly accessible servers, a reverse proxy server can also be useful. A decision to perform inspection at the application layer, whether by means of a dedicated intrusion prevention device or a deep packet inspection firewall should be based in large part on the human resources available to both configure the device appropriately and to respond to alerts as they occur.

Beyond the issue of which firewall technologies to use, a number of questions need to be addressed before an organization can choose an appropriate firewall. The following discussion summarizes the major points. A security consultant can be a useful resource to guide the organization towards the most appropriate combination of hardware and software.

Modern day firewalls run on a bewildering variety of hardware platforms. Some organizations may prefer the alluring simplicity of a firewall appliance, while others may choose an open server, preferably one that conforms to the company's server purchasing standards.

Some firewall vendors provide appliances only, in which case the customer has no choice as to which operating system to use. Other firewall vendors offer software packages that can run under a wide variety of Windows, Unix and Linux operating systems. Despite claims by some that a Unix/Linux-based firewall is inherently more secure than one that runs on Windows, the truth is that either can make an excellent firewall platform, provided the appropriate hardening steps are undertaken. While this may be an easier task with Unix/Linux, there are many

excellent resources available to assist in hardening a Windows system. (The Center for Internet Security is one such resource. See cisecurity.org for details.)

Once the hardware and operating system platforms have been established, the next step is to determine the processor speed, memory and disk space requirements. A firewall designer who undersizes a firewall to the point where all traffic that crosses it slows to a crawl will quickly learn the difference between fame and notoriety.

Some firewall software can take advantage of multi-core processors, which can be advantageous, especially when deep packet inspection is being employed. As stated earlier, stateful inspection firewalls can have robust memory requirements. Organizations need to ensure that memory sizing takes account of both current and anticipated future connectivity requirements.

Probably the most overlooked resource requirement is that of disk capacity. If the organization has a policy that dictates all traffic through the firewall be logged, disk utilization can grow very quickly.

Many firewall vendors now offer solutions that guarantee high availability of the firewall infrastructure. These solutions can take the form of a hot standby firewall that automatically takes over in the event the production system fails or a load sharing setup, whereby two or more firewalls share the traffic load according to some pre-defined algorithm. These load sharing solutions also provide high availability in the event that one module fails. In most cases, stateful inspection firewalls are able to fail over to a backup system without losing existing connections, although the reliability of this mechanism depends to some degree on the tolerance of the application to any delays in response.

Firewall vendors may offer additional functionality, such as virtual private networking, anti-virus software and content management. The combination of these security technologies into a single product has given rise to a class of security devices known as Unified Threat Management (UTM) devices. I will discuss UTM more fully in a future article but it is important to realize that a UTM device is a compromise akin to the “Home Theater in a Box” concept, as no single vendor has the best of breed application across all security technologies.

4 Firewall Maintenance

As is the case with all security technologies, one cannot simply install a firewall on one's network and then walk away. Several factors contribute to making the care and feeding of a firewall a time consuming activity.

The first obvious factor is that, even if the perfect firewall ruleset is installed on day one, business requirements change over time and the firewall policy must change to reflect these new requirements. Unfortunately, all too often, a new application will be prepared for deployment and only when it is time for commissioning will it occur to the project team that the firewall rules will need to be updated. If firewall changes need to go through a change management process (as they should), project delays can occur while the security team determines what firewall changes are required and obtains the necessary approvals. The solution to this problem is to ensure that all projects involving the deployment of a new application (or the activation of new functionality within an existing application) consider possible firewall impacts early on in the project.

A second complication involves the deployment of commercial applications. In my experience, few vendors consider any firewall impacts when they deliver an application to a customer. In addition, when pressed on the issue, many vendors are unclear about what protocols and services the firewall needs to allow for the application to function properly. The usual response is something along the lines of "Turn on the application and see what traffic is dropped." Such an approach can lead to frustrating delays in deployment of a new application and can also lead to opening up the firewall to more traffic than the application really requires. Again, the solution is to make it a requirement early on in the project that the vendor specify exactly which protocols and services the application uses. There should also be a stipulation that site acceptance testing not be considered complete until the application works successfully within the firewall environment.

Firewall rulesets tend to grow over time, leading to configuration mistakes, poor performance and potential security holes. Often, security administrators are reluctant to remove rules that are no longer required for fear of breaking something. Programs exist that can analyze a firewall ruleset and optimize the order (placing the most used rules at the top), inform the administrator of risky rules and flag rules that have not been used for a long time. A security consultant can be a useful resource to audit a firewall ruleset and verify that it remains in compliance with any applicable regulatory requirements, such as Sarbanes-Oxley.

Ian Parker is the president of Sentinel Ridge Consulting Group, a firm specializing in the security of corporate computer networks. You can reach him at sentinelridgeconsulting.com.