

# Considerations In Developing Firewall Selection Criteria



Adepteck Systems, Inc.

## Table of Contents

Introduction .....	1
Firewall's Function .....	1
Firewall Selection Considerations.....	1
Firewall Types.....	2
Packet Filter.....	2
Stateful Inspection.....	2
Application Proxy .....	3
Firewall Design Issues .....	3
Firewall Platform.....	4
Logging and Reporting.....	5
Virtual Private Networks.....	5
Authentication Methods .....	6
Proxy Features.....	6
Management/Configuration .....	7
Network Address Translation.....	7
High Availability/Load Balancing .....	8
Performance .....	8
Security Philosophy.....	9

## **Introduction**

Papers and articles about firewall cover the complete breadth of the subject with one major exception, firewall selection criteria. Papers and articles on this subject invariably leave this as an exercise for the reader. This is sufficient for experienced firewall and network administrators. However, this does nothing to help the intermediate to novice administrators. This paper will give some ideas on what information should be looked at in the evaluation process.

Firewall selection criteria fall into two major categories: one involves information on how a firewall is designed, built and operates while the other involves the network and environment that the firewall will operate in.

This paper deals with developing questions and selection criteria from the first category involving how a firewall is designed, built and operates. This paper does not deal with any questions or criteria that involve the environment that the firewall will be deployed in.

## **Firewall's Function**

Firewalls have one purpose: provide access control at the border of a security domain. For a firewall to accomplish its purpose, it must meet the following requirements: inspect packets and network connections, apply access control rules, authenticate network or user connections, and log actions. For a firewall to function effectively, all network traffic entering or leaving a security domain must pass through the firewall. This makes putting other software on the firewall very appealing because of its location within the network infrastructure; but avoid this temptation. Adding software that is not needed for the primary purpose of controlling access to a security domain greatly increases the risk of compromise, and thus resulting in failure to achieve its primary purpose.

## **Firewall Selection Considerations**

When considering a firewall product there are many areas of concern that should factor into the selection criteria. The following list covers the major categories that should be investigated during the selection process:

- Firewall type
- Firewall design issues
- Firewall platform
- Logging and reporting
- Virtual private network capabilities
- Authentication mechanisms
- Proxy features
- Management and configuration
- Network address translation features

- High availability and load balancing
- Performance

Suggested areas of interest in all of the above categories will be detailed in the rest of this paper.

## **Firewall Types**

Firewalls are based on one of three fundamental approaches: packet filtering, stateful inspection, and application proxies. Today, most firewalls use one approach for most of their functionality but employ elements from another approach to strengthen the overall product resulting in a hybrid. As an example, most hybrid firewalls based on stateful inspection technology, still contain the strengths and suffer most of the weaknesses of this technology. The primary advantage that this example hybrid provides, involves the addition of intelligent proxies for a very few of the most popular protocols.

### **Packet Filter**

Packet filter firewalls are the most basic type of firewall technology. They are essentially routers with access control lists, which allow them to be fast and flexible. However, these firewalls look primarily at layer 3 of the OSI model, which gives them their speed. The firewalls do gain flexibility by operating at layer 3 in being able to provide some access control for layer 3 protocols other than IP. The major problem with operating this low in the OSI model is the inability to access information of the higher layers, which denies the use of this information for logging or decision making. These firewalls suffer a number of disadvantages listed below:

- Inability to protect against vulnerabilities in application layer protocols.
- Limited logging information available.
- Limited number of authentication methods.
- Must allow all high ports from 1024 to 16384 in so that the return traffic to client ports.
- Vulnerable to problems within the TCP/IP specification and protocol stack.
- More susceptible to improper configuration due to limited number of variables.

The only true packet filter firewalls currently sold are routers being deployed as a firewall. All other current firewalls have more advanced features incorporated. The most appropriate deployment for packet filter firewalls are areas requiring very high speed but where the weaknesses listed above are not important.

### **Stateful Inspection**

Stateful inspection firewalls are a superset of the functionality of packet filter firewalls. These firewalls are fully aware of layer 4 of the OSI model, which gives them complete access to the information contained in the TCP/IP protocol stack and specification. The firewall now has the information to work on sessions that are kept track of using state tables. State tables give the firewall the ability to only allow connections initiated from the internal network to come in from the outside on the client port range of 1024 to 16384. This additional information also allows the capability of doing more advanced network address translation. This superset functionality only works with TCP/IP traffic.

## **Application Proxy**

Application proxy firewalls work by completely rewriting all packets that traverse the firewall using specific proxy instances for each protocol. These firewalls provide two types of proxies: intelligent proxies and generic proxies. The intelligent proxies can decode and understand the application protocols while generic proxies only pass the application data through the firewall untouched. Access to information in OSI model layer 3 through layer 7 and not requiring a layer 3 route between the inside and outside interfaces of the firewall are the primary advantages of an application proxy firewall. Access to the higher layers of the OSI model allows the firewall to make decisions with more knowledge than available to other firewall types that may only access layer 3 or layer 4. Application proxies prevent attacks that rely on problems within the TCP/IP specification and protocol stack, by completely rewriting all packets that traverse them.

The disadvantages of application proxy firewalls are speed and flexibility. Since the proxies decode and analyze the entire packet, this requires much more overhead than other firewall types. Firewall vendors usually provide intelligent proxies for only the ten to twenty most popular protocols, which limits the firewalls flexibility. Protocols for which an intelligent proxy does not exist are supported via a generic proxy, which provides only the advantage of completely rewriting packets over stateful inspection.

## **Firewall Design Issues**

The design of a firewall is critical to understanding what environment the firewall may be suitable for. Fundamental design issues include:

- Does the firewall fail closed or open?
- Does the firewall completely rewrite packets traversing the firewall?
- How does the firewall handle fragments in relation to the application of filters?
- What is the default state out of the box?

These issues reflect what tradeoffs a firewall vendor makes with respect to speed and security.

Firewalls are either designed to fail closed or fail open. A firewall that fails closed will stop passing traffic if a critical component fails, while a firewall that fails open, will still pass traffic even when critical components of the firewall have failed. The most secure design philosophy is failing closed; however, in some environments this may not be acceptable. A good example is with logging; a firewall that fails closed will stop passing traffic if logging ceases, while a firewall that fails open, will keep passing traffic.

The next major firewall design consideration involves rewriting packets. Firewalls either completely recreate a new packet on the outbound side of the firewall for all incoming packets allowed through or rewrite only a portion of the packet. Firewalls that completely rewrite a packet are generally immune from future undiscovered TCP/IP specification and protocol stack vulnerabilities. Firewalls that only rewrite a portion of a packet are generally vulnerable to all future TCP/IP specification and protocol stack vulnerabilities. Application proxy firewalls completely rewrite all packets that traverse the firewall, while packet filter and stateful inspection firewalls do not.

Another firewall design consideration is the handling of fragmentation in relation to firewall rules. Most firewalls that are not based on application proxy technology try to meet the requirements of an IP router as specified in RFC 1812, which does not allow routers to defragment any packets except those destined for the router. This presents a serious problem for the firewall. The firewall needs to defragment packets to access the information to make an access control decision but an RFC-1812-compliant firewall may not pass on the defragmented packet to the destination host. Only within the last year or so have the majority of firewalls on the market succeeded in addressing this problem adequately.

The firewall's default state gives good insight into the security philosophy of the vendor. Firewalls should not pass any traffic unless the implementor configures the firewall to allow traffic. Some firewalls have an implicit rule set upon installation that allows internal traffic outbound.

## **Firewall Platform**

Current firewalls are implemented on a wide variety of platforms. These vary from blackbox network appliances to application software running on general-purpose operating systems. Understanding the benefits and limitations of the platform that a given firewall product operates on is vital to assessing and understanding a firewalls capabilities. Firewalls implemented as blackbox network appliances are simple to set up and maintain but usually suffer from limited capabilities. On the other hand, firewalls implemented on top of general-purpose operating systems are limited to the operating systems constraints and the implementor's skill in configuring the system.

Firewall operating systems come in several varieties: Microsoft based operating systems, standard UNIX operating systems, and firewall vendor-proprietary operating system. Each has its own strengths and weaknesses. Microsoft and standard UNIX operating systems have the major problem of proper operating system installation and configuration for a firewall environment. Firewall systems must operate with a minimum operating system installation, which entails removing all non-essential operating system modules and options. Operating systems that function with a smaller set of capabilities are more suitable for a firewall platform than those that require a larger set of capabilities to operate. Additionally the implementor must configure the operating system for maximum firewall performance and operating system security, which usually requires a high degree of skill with the operating system. Firewall vendor-proprietary operating systems come with a minimal set of functionality and with the operating systems configured for maximum firewall performance and security.

Some understanding of the hardware architecture that a firewall will help the evaluator qualify the performance claims of firewall vendors. Firewalls that use a standard 33 MHz 32-bit PCI bus are not capable of supporting more than two 100Mb interface on the same bus when PCI latency factors are considered. Systems that use a 66MHz 64-bit PCI bus are not fully capable of supporting 2 1000Mb interfaces on the same bus.

## Logging and Reporting

Logging is a critical function of a firewall. For an access control device to be useful it must be able to log all activity in which it is involved. The more information that a log entry contains, the more useful that entry is for reporting and investigative purposes. Organizations can benefit from logging both successful and unsuccessful actions.

Since logging is such a critical function of a firewall, an organization needs to understand how and what is logged for any firewalls under consideration. Some firewall products use a proprietary log mechanism and a binary format while others use industry standard logging methods such as syslog or NT Event Log. The key features to understand about a logging mechanism are the following:

- Does the mechanism log all messages sent to the logger?
- How does the mechanism and firewall handle a disk full state for the log partition?
- Does the firewall continue to operate if logging ceases to function?
- Can the firewall logs be exported to COTS reporting tools?
- Under what circumstances could log entries be lost?
- Can log entries be written directly to a logging host?
- What happens if the logging host is not available?

Industry standard logging mechanisms are not always appropriate for a firewall logging mechanism.

Syslog is a widely deployed industry standard logging mechanism that is one common method used by firewalls to provide logging. This is unfortunate because syslog has several undesirable features some of which are:

- Syslog only allows 1024 byte log entries.
- Syslog uses UDP with zero error checking for network communication problems.
- Syslog does not guarantee that a log entry gets recorded.
- Syslog will throw away packets under high load scenarios.

These limitations make syslog entirely unsuitable as a logging mechanism for a firewall.

## Virtual Private Networks

Virtual private networks are widely deployed to provide communication security between organizations, within organizations, and between organizations and remote employees. Virtual private networks provide communication security between two endpoints by encrypting the communication channel, but do not provide access control for this channel. This requires that virtual private networks work with firewalls, which provide the access control needed for a secure perimeter. Questions to consider about the virtual private network capabilities of a firewall are:

- Does the firewall support virtual private networks using IPSEC?
- Does the firewall support hardware acceleration for the encryption of the virtual private network traffic?
- With which other vendors' IPSEC implementations has the firewall been certified to interoperate?

Firewalls are not always the most appropriate end point for a virtual private network. However, all virtual private network traffic should terminate outside of at least the innermost firewall of an organization.

## **Authentication Methods**

An important ability of a firewall is authenticating the traffic that traverses the firewall. The different firewall technologies allow different authentication methods. Usually the higher up the OSI model the firewall is able to operate the more authentication options it is able to provide. There are many types of authentication that a firewall may provide some of which are:

- IP address and port
- Local username and password
- SecureID
- Defender
- Safeword
- Radius
- NT domain
- LDAP
- TACACS+

These authentication methods provide varying degrees of confidence from the basic and insecure IP address based up to SecureID, Defender, and Safeword. These methods also vary in the ease of management. Several methods allow firewall authentication to be integrated into the organizations centralized Directory Services. An important consideration for an organization is which proxies or protocols allow which authentication methods. Lastly, many firewalls authentication methods provide the ability to limits access via time of day, week, or month on a per user or rule basis.

## **Proxy Features**

In addition to application proxy firewalls, most hybrid firewalls implement at least a few application proxies. These intelligent proxies allow the firewall to make access control decisions within the application layer. Even the best application proxy firewalls only have intelligent proxies implemented for the most popular protocols. Typically an application proxy firewall will have intelligent proxies for HTTP, SSL, FTP, Telnet, SMTP, Finger, SNMP, SQL, Ping, SUNRPC, Rlogin, and RealAudio. Hybrid firewalls based on other technologies will only implement intelligent proxies for a small subset of these protocols.

The capabilities and versions of protocols supported by intelligent proxies vary between vendors and protocols. Discovering exactly which application and protocol versions are supported is important but not always easily accomplished. Some desired proxy features among various protocols are the ability to control these characteristics:

- URL length
- MIME types allowed
- Active content

- File/transmission/message size
- Commands

Limiting URL length in application proxies protecting web servers is a highly effective method of preventing buffer overflow attacks. Intelligent proxies with more capabilities are able to provide more access control options to the organization implementing the firewall.

## **Management/Configuration**

Ease of management and configuration are the most thoroughly covered aspects of firewalls in virtually all reviews. It is also usually weighted the most strongly when these reviews provide weighting or ranking. While management and configuration issues are very important, a firewalls primary purpose is access control. The mechanism to implement the access control and the quality of the implementation of that mechanism is the most critical factors in evaluating a firewall. That being said, firewall management and configuration issues are important to properly implementing and operating a firewall.

Effective firewall management is a complex issue, which tends to affect large organizations more than small ones. Large organizations need to manage multiple firewalls with varying rule sets. These organizations also need to automate as many of the routine maintenance tasks as possible. Lastly, the management of the firewalls needs to be able to integrate into the organizations systems and security management infrastructure.

Some questions that need to be answered to properly evaluate the management capabilities of a firewall are:

- Does the firewall management solution use a centralized management server?
- Can multiple firewall administrators concurrently access the firewall management server?
- Does the firewall management solution provide multiple levels of access (ex. Read only, log viewing, full access)?
- Can the firewall management system push the same rule set to multiple firewalls?
- What type of remote management mechanism does the firewall vendor provide?
- Does the remote management software allow real-time viewing of log files?
- Can you backup all relevant firewall configuration files? Is the backup done through the command line, a GUI, or both? Is there a way to automate this?

Evaluators will need to develop more specific criteria as appropriate for their environment.

## **Network Address Translation**

Network address translation is an important ability of a firewall. Most organizations' networks today use RFC1918 private address space both to conserve their registered IP numbers and to hide the addressing scheme in use behind the firewall. There are three methods of address translation: static address translation, hiding network address translation and port address translation. Static address translation involves setting up one to one translations between an internal address and an external address. Static address

translation has the limitation of requiring large numbers of external addresses. Hiding address translation involves translating all internal addresses to one external address. Hiding address translation is usually limited to only using the external address of the firewall and does not provide the ability to internal services available externally. Port address translation involves using one external address for all internal addresses but also provides the ability to translate individual ports differently. Port address translation provides the strengths of hiding address translation with the additional ability to make internal services available externally.

## **High Availability/Load Balancing**

High availability and load balancing are advanced features that are gaining more importance every day as firewalls are deployed in environments that require constant uptime and large amounts of network bandwidth. No matter the product or technology, these features require a thorough understanding of the implementation by the vendors and the environment that the products will be deployed in.

High availability can be provided using either a hot standby model or by having multiple fully functional systems that take over the load of any system that fails. The first model provides only high availability while the second is a more robust and scalable architecture that also provides capacity. It is very important to understand exactly how the high availability feature works in detail. Does the fail over mechanism detect the failure of a single service or must a larger scale failure occur? How this feature works and is configured will dramatically impact the behavior and usefulness of the system. Any high availability mechanism implemented must be regularly tested to verify correct operation. Senior management gets very upset when a single system failure takes the infrastructure down after they have paid for high availability.

Load balancing allows organizations to scale firewall solutions to situations where a single firewall can't handle the communications requirements. Most firewall vendors support their products when working with IP redirection technology, which is the usual implementation of firewall load balancing. Depending on the IP redirector and firewall products involved, this allows load balancing and high availability on a per service basis. Additional benefits to this technology include the ability to introduce changes to production infrastructure without impacting service.

## **Performance**

Firewall performance is a difficult issue to gauge accurately. The ruleset implemented, the network environment it resides in, the hardware it is running on, and the configuration of any underlying operating system will heavily influence the performance a firewall product. This make it difficult for any standardized testing to provide meaning numbers for a particular implementation. Most firewall testing performed for firewall reviews does not involve the firewalls using network address translation because network address translation has a dramatic adverse impact on the performance of packet and stateful inspection firewalls. Application proxy firewalls are not significantly impacted because the packets are always completely rewritten. Large organizations are strongly encouraged

to bring firewalls in house for testing in an environment similar to the intended implementation environment.

## **Security Philosophy**

The previous discussion was developed using the security philosophy described here. The most important concepts are least privilege and keep it simple. The philosophy of the author is when trade offs between security, speed, and functionality must be made that you should err on the side of security.

The concept of least privilege is probably the most important in all of information security. Only communications and services required for the organization to function should pass through the firewall. The organization should implement network address translation between internal and external networks to hide the addressing scheme used internally. Using RFC1918 addresses is the most preferable, because this provides an extra hurdle for attackers to overcome.

All implementations of a security system should be kept as simple as possible. The more complex added to a security system the more difficult to ensure that vulnerabilities and error are not introduced. The practical result of this philosophy is that firewalls should only provide access control. All other tasks should reside on their own dedicated servers. Web and mail filtering and screening software should not reside on a firewall. These tasks are complex and require a lot of computing overhead which is more appropriately implemented on their own systems.