

SE 4C03 Winter 2005
Firewall Design Principles

By: Kirk Crane

Firewall Design Principles

By: Kirk Crane 9810533

Introduction

Every network has a security policy that will specify what traffic is allowed to enter and leave the network. Most home computers have an open policy allowing all traffic to enter and leave. As more computers become connected to the Internet, the proliferation of viruses and hackers are a concern of even the casual Internet user. The generally accepted approach to securing your site is to take the following steps[1]:

- 1) Identify what you are trying to protect.
- 2) Determine what you are trying to protect it from.
- 3) Determine how likely the threats are.
- 4) Implement measures that will protect your assets in a cost-effective manner.
- 5) Review the process continuously and make improvements each time a weakness is found.

The most complex software will not protect your data if you have not identified what you are trying to protect and from whom. Although these steps are important to securing a network, only the implementation (step 4) of the security policy will be considered. This paper will discuss how a firewall can implement your security policy to prevent unauthorized access.

Architecture

There are generally four types of firewalls: Packet Filtering Firewalls, Circuit Level Gateways, Application Level Gateways, and Stateful Multilevel Inspection Firewalls[2]. These firewall designs are in increasing order of complexity and evolution.

Packet Filtering Firewalls

Packet filtering firewalls was the first firewall architecture and was developed by Cisco. As a packet enters the network, the IP header data is analyzed against a set of rules to determine if it should be allowed to enter the computer network[3]. These rules cannot contain any state information and as a result are relatively simple to produce and very efficient to process. Since these rules introduce little intelligence (keeping state and using higher level protocols), complex logging and policies cannot be implemented with this technique.

Circuit Level Gateways

Circuit level gateways were the second generation of firewalls and were developed at the AT&T Bell Labs. The higher level TCP header data was used to analyze the following information: session identifier, connection state (handshake, established, or closing), sequencing information, source and destination IP address, and the physical network interface. The name comes from circuit relays, which creates a direct connection between computers. The firewall will analyze a new connection and apply the security policies against it. If the connection is allowed, the remaining packets that arrive within the same connection are allowed without further tests, just as a direct connection would work. This added intelligence can help eliminate connections from

only certain hosts without excessive overhead. However, protocols other than TCP have no advantages over first generation firewall architectures.

Application Level Gateways

Application level gateways were the third generation firewalls. This architecture employs a special server to act as a proxy for a host on the secured network to talk to the outside network[3]. The proxy server will respond as though it was the actual host outside the network. Before it sends a request to the host outside the network it will apply the security policy to the request and determine if it is authorized. Since the proxy server works with the application level protocol, it can implement very complex rules. Incoming requests are handled in much the same way. Instead of a direct connection, the outside computer must communicate with the proxy server in order for its packets to be passed along to the inside host. The proxy server will allow no direct connection between a trusted and outside host. Also, a proxy server has a great deal of intelligence since it can inspect application level protocols, allowing it to produce detailed logs and applying complex rules. The overhead involved in a proxy server is significant and can easily be a bottleneck for communication. Also, some applications must be configured to work with proxies and therefore complicates setup of a client computer.

Stateful Multilevel Inspection Firewalls

The latest generation is stateful multilevel inspection firewalls. These firewalls use and record the state information from the IP, transport and application level[3]. This translates into being a very complex and powerful tool to use and offers the greatest amount of security. For example, if a hacker were to intercept packets and change the IP address to an outside host he has control over, the state of previous packets would be used to detect the inconsistency and the appropriate action (drop/ignore remaining packets in the connection and log) could be taken. This architecture will have considerable overhead since certain state information needs to be tracked and compared. Also, most operating systems either do not contain a stateful firewall or have just released one (and probably require more testing).

Network Design

A network with few hosts, providing no public services is likely going to have one firewall installed on the router (see Figure 1 in Appendix A). This type of setup is referred to as a single layer architecture[4]. For such a setup a good policy to employ would be to deny all incoming connections (since you are not providing any services) and limit the outgoing connections to the applications you use. This can be accomplished with a packet filter firewall and be very effective and easy to setup. All other computers in the private network will then be separated from the outside network by the firewall.

A more complex network setup will likely employ the concept of a demilitarized zone (DMZ) in a multi-layer architecture[4]. When services such as HTTP or FTP are required, it is best to keep these machines protected, but still accessible to the outside network. Since these machines will accept connections from outside networks they cannot be considered trusted and hence are part of the DMZ network. A firewall should

still be placed between the DMZ and the outside network to reduce the chance for compromise and only allow incoming connections to pass through for the services being offered. The other machines in the internal network that do not provide services to the outside network should be isolated from the DMZ by using another firewall with a more strict security policy (see Figure 2 in Appendix A). These two firewalls should be connected in series and use different software implementations to provide additional security—two firewalls must be penetrated using different techniques in order to obtain access to the internal network. The second firewall for the private network should not let any incoming connections come through. In the scenario of a computer in the DMZ being compromised, all computers in the private network would still be protected by the second firewall.

A firewall can also be placed on a single computer. The firewall could be setup to ensure that any malicious software will be blocked from connecting outside of the computer (unless it tunnels through popular ports such as HTTP). Another added benefit would be to log activity and also to ensure the host is secure from other hosts that may have been compromised on the same network.

Most technologies used for networking were designed in a time when security was not a main concern. As security became a major concern, these unsecured technologies became the base for implementing secure technologies on top of. New technology to replace the old has been slow to catch on, but has kept security as a top concern. When these technologies (IPv6/IPSEC) become required for communication, technology such as the firewall will be reduced to what its initial intentions were and additions such as detecting spoofing will be handled at the protocol level instead of by the firewall.

Conclusion

A firewall is just another piece in the network security puzzle. To secure a network, a security policy must be devised to outline what you are trying to protect and from what threats. The policy can then be implemented using any existing technologies. A firewall is placed between all external networks to separate it from the internal network, creating a secure boundary around your network. The policy for the internal network is implemented using rules in the firewall software. The firewall will monitor all traffic entering and leaving the network by analyzing the header information within the packets. When a packet fails to pass a specific rule, it will not be permitted to continue to its destination.

To offer a public services network externally, requires a more complex multi-layer architecture. The most secure approach is to create a DMZ between two firewalls connected in series. These firewalls will have different security policies and should have different software to diversify the defenses.

Appendix A




Figure 1: Single layer architecture




Figure 2. Multi-layer architecture: Dual Firewall with DMZ network design

References:

[1] M. Fites, P. Kratz, and A. Brebner, "Control and Security of Computer Information Systems", Computer Science Press, 1998.

[2] Vicomsoft, 2003. Firewall White Paper-What different types of firewalls are there? http://www.firewall-software.com/firewall_faqs/types_of_firewall.html (March, 24, 2005)

[3] University of Georgia, 2005. Firewall Technology, <http://www.infosec.uga.edu/firewall.html> (March, 26, 2005)

[4] CERT, 2004, "Design the Firewall System", A practice from the CERT Security Improvement Modules, (March, 24, 2005)