

White Paper for Firewall

Firewall

The Internet has made large amounts of information available to the average computer user at home, in business and in education. For many people, having access to this information is no longer just an advantage, it is essential. Yet connecting a private network to the Internet can expose critical or confidential data to malicious attack from anywhere in the world. Users who connect their computers to the Internet must be aware of these dangers, their implications and how to protect their data and their critical systems. Firewalls can protect both individual computers and corporate networks from hostile intrusion from the Internet, but must be understood to be used correctly.


What is a firewall?

A firewall protects networked computers from intentional hostile intrusion that could compromise confidentiality or result in data corruption or denial of service. It may be a hardware device (see Figure 1) or a software program (see Figure 2) running on a secure host computer. In either case, it must have at least two network interfaces, one for the network it is intended to protect, and one for the network it is exposed to.

A firewall sits at the junction point or gateway between the two networks, usually a private network and a public network such as the Internet. The earliest firewalls were simply routers. The term firewall comes from the fact that by segmenting a network into different physical sub networks, they limited the damage that could spread from one subnet to another just like firedoors or firewalls.


Figure 1: Hardware Firewall

Hardware firewall providing protection to a Local Network



2: Computer with Firewall Software

Computer running firewall software to provide protection



What does a firewall do?

A firewall examines all traffic routed between the two networks to see if it meets certain criteria. If it does, it is routed between the networks, otherwise it is stopped. A firewall filters both inbound and outbound traffic. It can also manage public access to private networked resources such as host applications. It can be used to log all attempts to enter the private network and trigger alarms when hostile or unauthorized entry is attempted. Firewalls can filter packets based on their source and destination addresses and port numbers. This is known as address filtering. Firewalls can also filter specific types of network traffic. This is also known as protocol filtering because the decision to forward or reject traffic is dependant upon the protocol used, for example HTTP, ftp or telnet. Firewalls can also filter traffic by packet attribute or state.

What can't a firewall do?

A firewall cannot prevent individual users with modems from dialing into or out of the network, bypassing the firewall altogether. Employee misconduct or carelessness cannot be controlled by firewalls. Policies involving the use and misuse of passwords and user accounts must be strictly enforced. These are management issues that should be raised during the planning of any security policy but that cannot be solved with firewalls alone.

Who needs a firewall?

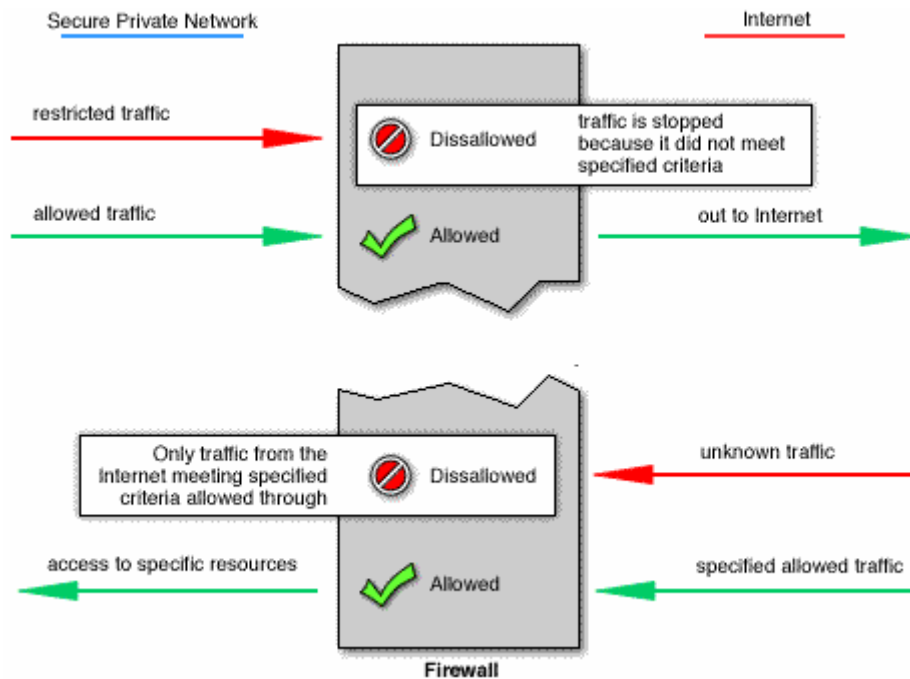
Anyone who is responsible for a private network that is connected to a public network needs firewall protection. Furthermore, anyone who connects so much as a single computer to the Internet via modem should have personal firewall software. Many dial-up Internet users believe that anonymity will protect them. They feel that no malicious intruder would be motivated to break into

their computer. Dial up users who have been victims of malicious attacks and who have lost entire days of work, perhaps having to reinstall their operating system, know that this is not true. Irresponsible pranksters can use automated robots to scan random IP addresses and attack whenever the opportunity presents itself.

How does a firewall work?

There are two access denial methodologies used by firewalls. A firewall may allow all traffic through unless it meets certain criteria, or it may deny all traffic unless it meets certain criteria (see figure 3). The type of criteria used to determine whether traffic should be allowed through varies from one type of firewall to another. Firewalls may be concerned with the type of traffic, or with source or destination addresses and ports. They may also use complex rule bases that analyze the application data to determine if the traffic should be allowed through. How a firewall determines what traffic to let through depends on which network layer it operates at. A discussion on network layers and architecture follows.

Figure 3: Basic Firewall Operation



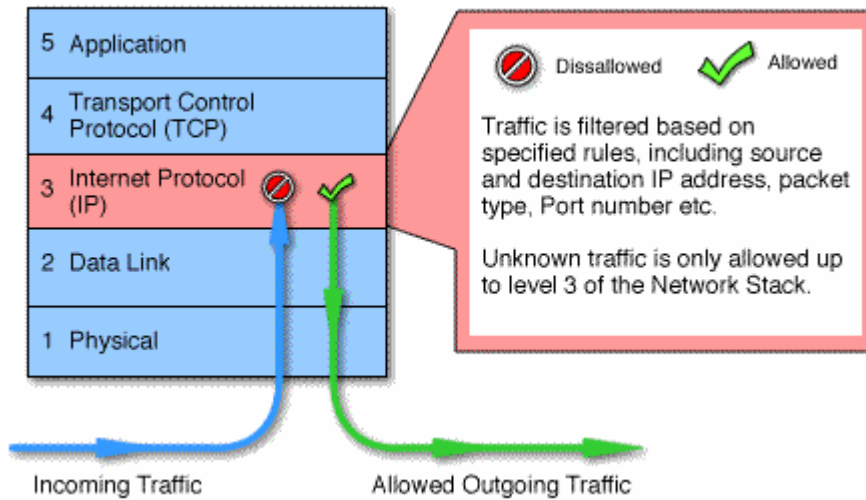
What different types of firewalls are there?

Firewalls fall into four broad categories: packet filters, circuit level gateways, application level gateways and stateful multilayer inspection firewalls.

Packet filtering firewalls work at the network level of the OSI model, or the IP layer of TCP/IP. They are usually part of a router. A router is a device that receives packets from one network and forwards them to another network. In a packet filtering firewall each packet is compared to a set of criteria before it is forwarded. Depending on the packet and the criteria, the firewall can drop the

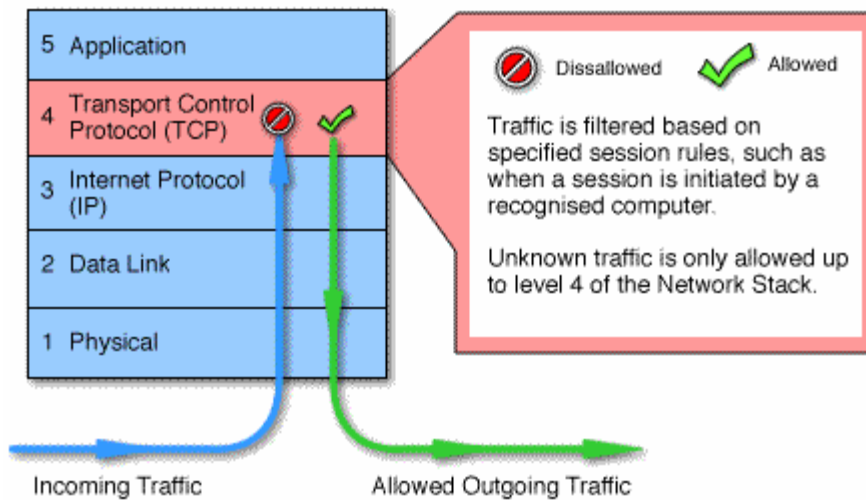
packet, forward it or send a message to the originator. Rules can include source and destination IP address, source and destination port number and protocol used. The advantage of packet filtering firewalls is their low cost and low impact on network performance. Most routers support packet filtering. Even if other firewalls are used, implementing packet filtering at the router level affords an initial degree of security at a low network layer. This type of firewall only works at the network layer however and does not support sophisticated rule based models (see Figure 4). Network Address Translation (NAT) routers offer the advantages of packet filtering firewalls but can also hide the IP addresses of computers behind the firewall, and offer a level of circuit-based filtering.

Figure 4: Packet Filtering Firewall



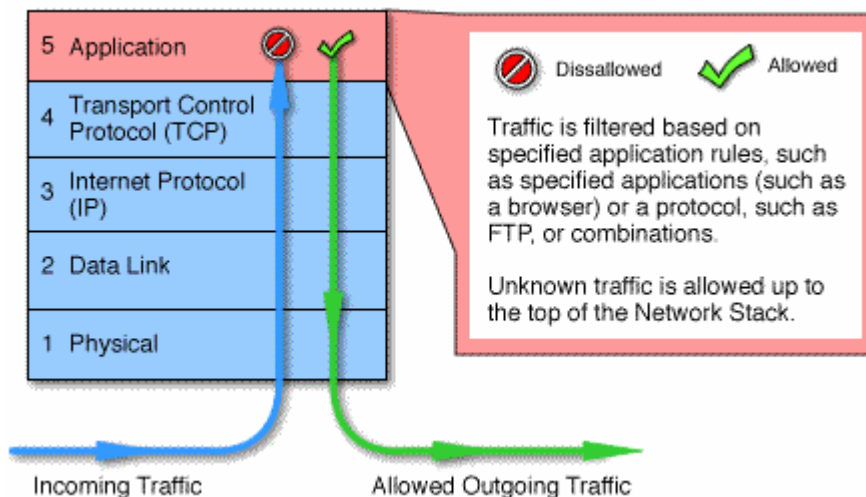
Circuit level gateways work at the session layer of the OSI model, or the TCP layer of TCP/IP. They monitor TCP handshaking between packets to determine whether a requested session is legitimate. Information passed to remote computer through a circuit level gateway appears to have originated from the gateway. This is useful for hiding information about protected networks. Circuit level gateways are relatively inexpensive and have the advantage of hiding information about the private network they protect. On the other hand, they do not filter individual packets.

Figure 5: Circuit level Gateway



Application level gateways, also called proxies, are similar to circuit-level gateways except that they are application specific. They can filter packets at the application layer of the OSI model. Incoming or outgoing packets cannot access services for which there is no proxy. In plain terms, an application level gateway that is configured to be a web proxy will not allow any ftp, gopher, telnet or other traffic through. Because they examine packets at application layer, they can filter application specific commands such as http:post and get, etc. This cannot be accomplished with either packet filtering firewalls or circuit level neither of which know anything about the application level information. Application level gateways can also be used to log user activity and logins. They offer a high level of security, but have a significant impact on network performance. This is because of context switches that slow down network access dramatically. They are not transparent to end users and require manual configuration of each client computer. (See Figure 6)

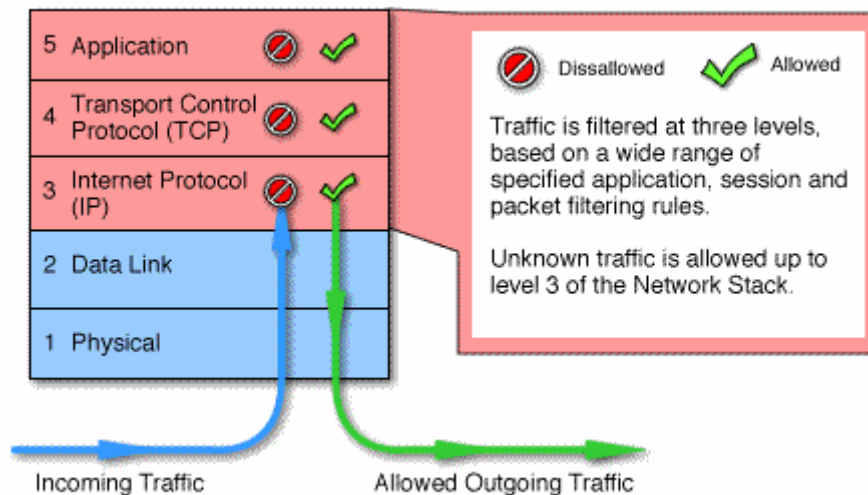
Figure 6: Application level Gateway



Stateful multilayer inspection firewalls combine the aspects of the other three types of firewalls. They filter packets at the network layer, determine whether session packets are legitimate and evaluate contents of packets at the application layer. They allow direct connection between client and host, alleviating the problem caused by the lack of transparency of application level gateways.

They rely on algorithms to recognize and process application layer data instead of running application specific proxies. Stateful multilayer inspection firewalls offer a high level of security, good performance and transparency to end users. They are expensive however, and due to their complexity are potentially less secure than simpler types of firewalls if not administered by highly competent personnel. (See Figure 7)

Figure 7: Stateful Multilayer Inspection Firewall



How do I implement a firewall?

We suggest you approach the task of implementing a firewall by going through the following steps:

Determine the access denial methodology to use.

It is recommended you begin with the methodology that denies all access by default. In other words, start with a gateway that routes no traffic and is effectively a brick wall with no doors in it.

Determine inbound access policy.

If all of your Internet traffic originates on the LAN this may be quite simple. A straightforward NAT router will block all inbound traffic that is not in response to requests originating from within the LAN. As previously mentioned, the true IP addresses of hosts behind the

firewall are never revealed to the outside world, making intrusion extremely difficult. Indeed, local host IP addresses in this type of configuration are usually non-public addresses, making it impossible to route traffic to them from the Internet. Packets coming in from the Internet in response to requests from local hosts are addressed to dynamically allocated port numbers on the public side of the NAT router. These change rapidly making it difficult or impossible for an intruder to make assumptions about which port numbers to use.

If your requirements involve secure access to LAN based services from Internet based hosts, then you will need to determine the criteria to be used in deciding when a packet originating from the Internet may be allowed into the LAN. The stricter the criteria, the more secure your network will be. Ideally you will know which public IP addresses on the Internet may originate inbound traffic. By limiting inbound traffic to packets originating from these hosts, you decrease the likelihood of hostile intrusion. You may also want to limit inbound traffic to certain protocol sets such as ftp or http. All of these techniques can be achieved with packet filtering on a NAT router. If you cannot know the IP addresses that may originate inbound traffic, and you cannot use protocol filtering then you will need more a more complex rule based model and this will involve a stateful multilayer inspection firewall.

Determine outbound access policy.

If your users only need access to the web, a proxy server may give a high level of security with access granted selectively to appropriate users. As mentioned, however, this type of firewall requires manual configuration of each web browser on each machine. Outbound protocol filtering can also be transparently achieved with packet filtering and no sacrifice in security. If you are using a NAT router with no inbound mapping of traffic originating from the Internet, then you may allow LAN users to freely access all services on the Internet with no security compromise. Naturally, the risk of employees behaving irresponsibly with email or with external hosts is a management issue and must be dealt with as such.

Determine if dial-in or dial-out access is required.

Dial-in requires a secure remote access PPP server that should be placed outside the firewall. If dial-out access is required by certain users, individual dial-out computers must be made secure in such a way that hostile access to the LAN through the dial-out connection becomes impossible. The surest way to do this is to physically isolate the computer from the LAN. Alternatively, personal firewall software may be used to isolate the LAN network interface from the remote access interface.

Decide whether to buy a complete firewall product, have one implemented by a systems integrator or implement one yourself.

Once the above questions have been answered, it may be decided whether to buy a complete firewall product or to configure one from multipurpose routing or proxy software. This decision will depend as much on the availability of in-house expertise as on the complexity of the need. A satisfactory firewall may be built with little expertise if the requirements are straightforward. However, complex requirements will not necessarily entail recourse to external resources if the

system administrator has sufficient grasp of the elements. Indeed, as the complexity of the security model increases, so does the need for in-house expertise and autonomy.

Is a firewall sufficient to secure my network or do I need anything else?

The firewall is an integral part of any security program, but it is not a security program in and of itself. Security involves data integrity (has it been modified?), service or application integrity (is the service available, and is it performing to spec?), data confidentiality (has anyone seen it?) and authentication (are they really who they say they are?). Firewalls only address the issues of data integrity, confidentiality and authentication of data that is behind the firewall. Any data that transits outside the firewall is subject to factors out of the control of the firewall. It is therefore necessary for an organization to have a well planned and strictly implemented security program that includes but is not limited to firewall protection.

What is IP spoofing?

Many firewalls examine the source IP addresses of packets to determine if they are legitimate. A firewall may be instructed to allow traffic through if it comes from a specific trusted host. A malicious cracker would then try to gain entry by "spoofing" the source IP address of packets sent to the firewall. If the firewall thought that the packets originated from a trusted host, it may let them through unless other criteria failed to be met. Of course the cracker would need to know a good deal about the firewall's rule base to exploit this kind of weakness. This reinforces the principle that technology alone will not solve all security problems. Responsible management of information is essential. One of Courtney's laws sums it up: "There are management solutions to technical problems, but no technical solutions to management problems".

An effective measure against IP spoofing is the use of a Virtual Private Network (VPN) protocol such as IPSec. This methodology involves encryption of the data in the packet as well as the source address. The VPN software or firmware decrypts the packet and the source address and performs a checksum. If either the data or the source address has been tampered with, the packet will be dropped. Without access to the encryption keys, a potential intruder would be unable to penetrate the firewall.

Firewall related problems

Firewalls introduce problems of their own. Information security involves constraints, and users don't like this. It reminds them that Bad Things can and do happen. Firewalls restrict access to certain services. The vendors of information technology are constantly telling us "anything, anywhere, any time", and we believe them naively. Of course they forget to tell us we need to log in and out, to memorize our 27 different passwords, not to write them down on a sticky note on our computer screen and so on.

Firewalls can also constitute a traffic bottleneck. They concentrate security in one spot, aggravating the single point of failure phenomenon. The alternatives however are either no Internet access, or no security, neither of which are acceptable in most organizations.

Benefits of a firewall

Firewalls protect private local area networks from hostile intrusion from the Internet. Consequently, many LANs are now connected to the Internet where Internet connectivity would otherwise have been too great a risk.

Firewalls allow network administrators to offer access to specific types of Internet services to selected LAN users. This selectivity is an essential part of any information management program, and involves not only protecting private information assets, but also knowing who has access to what. Privileges can be granted according to job description and need rather than on an all-or-nothing basis.