The Structure of Firewalls

Abhinay Kampasi IEEE Member, PICT abhinaykampasi@ieee.org
Devendra Desai IEEE Member, PICT d_devendra@hotmaiil.com
Sapna Bafna IEEE Member, PICT sapnabafna@ieee.org

Abstract

A firewall is a combination of hardware and software used to protect one network, the secure corporate network, from another network, the supposedly insecure network. The secure network is referred to as the trusted network. The insecure network is the untrusted network. The objective of the firewall is to control access to the trusted network. It allows only authorized traffic to enter the network. A firewall is part of an overall security policy, which should include tools and procedures such as system security and audits. Firewall gets its name from fire security systems in real life, which protects people from fire, simultaneously providing them with a safe exit.

Index Terms

Internet security, malicious attacks, trust, packet filters, proxy server, socks server, incident handling, disaster recovery.

1. Introduction

1.1 Internet Security

The Internet has made large amount 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.

This understanding can be achieved by knowing a thing or two about server security, which refers to controlling access to the web server. Access needs to be restricted in two ways. Only documents in certain directories should be available to the client and for certain documents, access should be limited to certain users. This is typically implemented with user-id and password security. It can be supported by the server software or by the application software.

1.2 Necessity

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 or a software program 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. In simple words firewall means segmenting a network into different physical subnetworks, that limit the damage that could spread from one subnet to another just like firedoors or firewalls.

Some firewalls permit only e-mail traffic through them, thereby protecting the network against any attacks other than attacks against the e-mail service. Other firewalls provide less strict protections, and block services that are known to be problems.

2. Why would One want a firewall?

The Internet, like any other society, is plagued with the kind of jerks who enjoy the electronic equivalent of writing on other people's walls with spraypaint, tearing their mailboxes off, or just sitting in the street blowing their car horns. Some people try to get real work done over the Internet, and others have sensitive or proprietary data they must protect. Usually, a firewall's purpose is to keep the jerks out of your network while still letting you get your job done.

Many traditional-style corporations and data centers have computing security policies and practices that must be adhered to. In a case where a company's policies dictate how data must be protected, a firewall is very important, since it is the embodiment of the corporate policy. Frequently, the hardest part of hooking to the Internet, if you're a large company, is not justifying the expense or effort,

but convincing management that it's safe to do so. A firewall provides not only real security--it often plays an important role as a security blanket for management.

Lastly, a firewall can act as your corporate ‘‘ambassador” to the Internet. Many corporations use their firewall systems as a place to store public information about corporate products and services, files to download, bug-fixes, and so forth. Several of these systems have become important parts of the Internet service structure (e.g.: UUnet.uu.net, whitehouse.gov, gatekeeper.dec.com) and have reflected well on their organizational sponsors.

3. Types of Attack

3.1 Social Engineering

Example 1: Inexperienced user is tricked into changing password.

Example 2: Attacker masquerades as administrator and asks for password for some reason or gives user new password and tells them to change it.

3.2 Impersonation

Any attack where the attacker captures valid user-id and password and reuses them to gain access to system.

Example: A user uses Telnet program to connect to system from remote site and an attacker with network sniffer such as tcpdump or nitsniff etc. captures the login session. The attacker is later able to login to system with captured user-id and password.

3.3 Exploits

These are attacks that seek to exploit a hole in a piece of software.

Example: The UNIX sendmail program runs with system privileges. Sending a message with the "To" and "From" fields completed as shown has given root access to the sender :-

To : mrinvisible@nonexistnat.com
From "| /bin/sed ’1,//d’ | sh"

3.4 Transitive Trust

Transitive trust attackers take advantage of the trust models used by remote services.
Example: Many networks use ".rhost" files so that users can log in from "trusted" hosts without giving a password. An attacker who gains access to a host and scans for exported file systems using a remote procedure call is able to build a trust model of the network. This is one of the attack strategies that the 1988 Internet "Worm" Virus used to propagate itself.

3.5 Data Driven

Data driven attacks take the form of Viruses and Trojan Horses.

Example: An attacker can e-mail the victim a postscript file with hidden file operations in it. If the victim displays the file on his workstation with a postscript interpreter (such as Ghostscript), the postscript interpreter will execute the file operations. These may perform actions such as adding the attacker's host name to the victim's ".rhosts" file allowing the attacker to gain access to the victim's computer.

3.6 Infrastructure

Infrastructure attacks include DNS Spoofing, ICMP Bombing and Source Routing.

Example: ICMP Bombing. ICMP (Internet Control Message Protocol) is used to re-route traffic on the fly and by routers to notify a host when a destination system or network is unreachable. An attacker can use widely available tools such as "icmpbomb" or "nuke" to send ICMP "host unreachable" packets to a target system effectively knocking the network off the Internet.

Most firewalls and routers can screen ICMP traffic. However ICMP is used for legitimate purposes such as Ping and screening ICMP messages in routers can cause network problems. Firewalls that are a single point of connectivity correctly interpret ICMP without letting it through.
Firewalls can block and log all source routed packets and tools like TCP wrappers can detect source routed packets and trigger alarms. Many routers can block source routed packets.

3.7 Magic

These are attacks that nobody has thought of yet. Such attacks if and when discovered will be full of surprises.

Example: (Possible) Racing Authentication, where an attacker is able to sniff packets as a legitimate user logs in with SecurID or other similar authentication token. The attacker mirrors the user's keystrokes and takes a guess at last digit of SecurID code, thereby winning the "race" with the user to login. If the attack is successful (an average of 1 in 10 should be) then the attacker is granted access, and the user probably just thinks they have made a typing error.

3.8 Denial of Service Attack

A denial of service attack seeks to deny use of resources to legitimate users. This type of attack can be achieved in a multitude of ways, for example by corrupting routing tables etc. causing messages to be re-routed, by overloading resources with junk messages, by damaging stored data, by locking user accounts, and so on.

Example 1 - The attacker ICMP bombs router off the network.

Example 2 - The attacker floods network link with garbage packets.

Example 3 - The attacker floods mail hub with junk mail (or many users send many messages to one address).

4. Firewall Architecture

There are lots of ways to structure your network to protect your systems using a firewall. If you have a dedicated connection to the Internet through a router, you could plug the router directly into your firewall system. Or, you could go through a hub to provide for full access servers outside your firewall.

4.1 Dial-up Architecture

You may be using a dialup service like an ISDN line. In this case you might use a third network card to provide provide a filtered DMZ. This gives you full control over your Internet services and still separates them from your regular network.

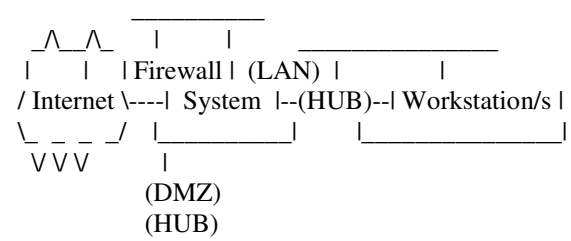

4.2 Single Router Architecture

If there is a router or cable modem between you and the Internet. If you own the router you could setup some hard filter rules in the router. If this router is owned by your ISP so you may not have the needed controls. You can ask your ISP to put in filters.

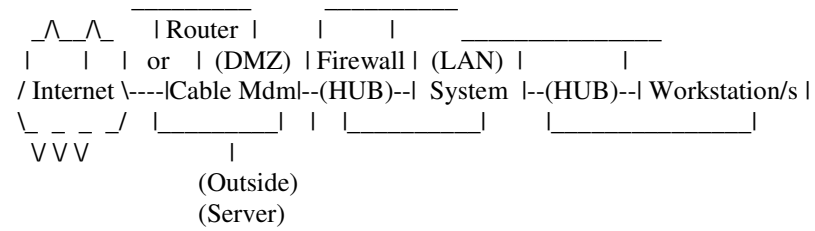

4.3 Firewall with Proxy Server

If you need to monitor where users of your network are going and your network is small, you can intergrate a proxy server into your firewall. ISP's some times do this to create interest in list of their users to resell to marketing agencies.

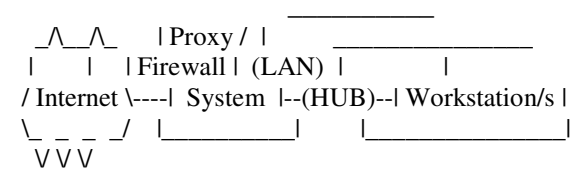

You can put the proxy server on your LAN as will. In this case the firewall should have rules to only allow the proxy server to connect to the Internet for the services it is providing. This way the users can get to the Internet only through the proxy.

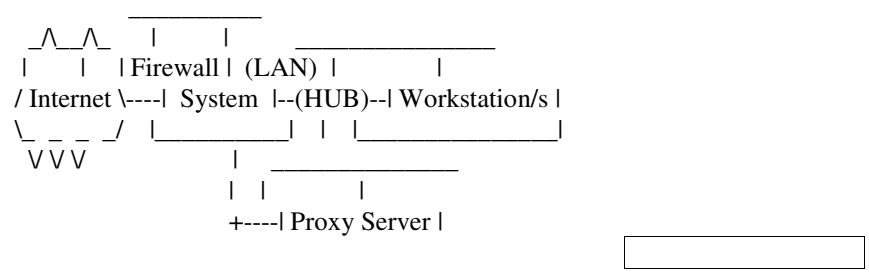

4.4 Redundant Internet Configuration

If you are going to run a service like YAHOO or maybe SlashDot you may want to make your system by using redundant routers and firewalls.

By using a round-robin DNS techniques to provide access to multiple web servers from one URL and multiple ISP's, routers and firewalls using High Availability techniques you can create a 100% uptime service.

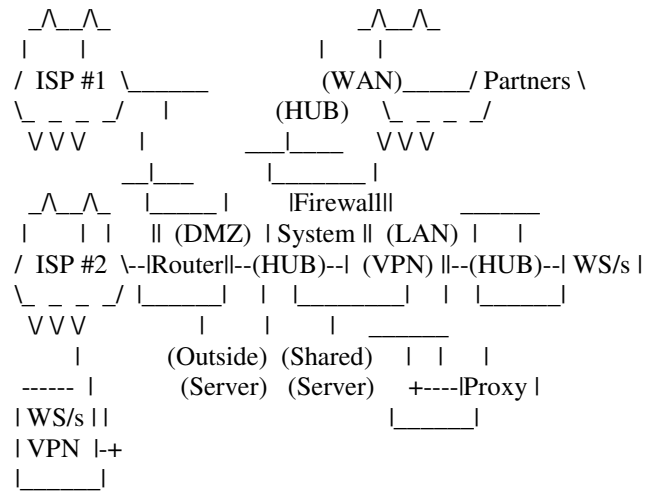

5. *Firewall Components*

5.1 Packet Filters

A packet filtering system selectively routes packets between internal and external hosts according to rules that reflect the organizations network security policy. Packet filtering may occur in a router, in a bridge, or on an individual host and operates at the network layer.

The type of router used in a packet filtering firewall is called a screening router. It is configured with rules to block or filter protocols and addresses and are installed at the external network gateway. Internal users usually have direct access to the Internet while all or most access to site systems from the Internet is blocked. However, the router could allow selective access to systems and services, depending on the policy. Inherently dangerous services such as NIS, NFS, and X Windows are usually blocked .

The screening router passes or rejects an IP packet based on information contained on the packet's header. The main information used is :-

IP Source and Destination Address - By filtering packets on the IP source and destination address the screening router is able to effectively block access to or from any site or host that is not trusted.

TCP or UDP source and destination port - The screening router makes use of the TCP "well known ports" to permit, deny, or re-route access to particular Internet services. For example many firewalls block all inward traffic except for email by rejecting all externally sourced packets bound for any port other than port 25, the Simple Mail Transfer Protocol port. The screening router can also route all World Wide Web traffic to a particular host.

A screening router is also able to base routing decisions on information not found in the packet header, for example the source and destination interfaces.

5.2 Application Level Gateways

Application level gateways are specialized application or server programs that run on a firewall host. These programs provide a safety barrier between the internal user and the Internet. Instead of connecting to the Internet directly with, say, a World Wide Web browser, the internal user connects to the application level gateway instead. The application level gateway then establishes the connection with the required world wide web server on the Internet and acts as a go-between for the session.

Application gateways operate at the application layer and can therefore provide access controls at the application protocol level and can handle store and forward as well as interactive traffic.

The main disadvantage of application level gateways is that they require special purpose code to provide each service that is relayed. However, this means that they therefore implement a policy of "deny

everything unless explicitly permitted" by default, which is often advantageous from a security perspective.

Application level proxies understand the application protocol and are therefore able to control the session based on the operations being requested. For example an application level proxy is able to block FTP PUT commands whilst permitting FTP GET commands.

The custom application acts as a "proxy" between the client and the server. Because all data between the client and the server is routed through the application proxy it is able to both control the session and provide detailed logging. This ability to log and control all incoming and outgoing traffic is one of the main advantages of application level gateway.

5.3 Circuit Level Gateway

Another type of application level gateway is called the circuit level gateway. Circuit-level proxies do not interpret the application protocols but they authenticate the user before establishing the circuits. They relay packets between the two communicating end-points but are not able to do any additional processing or filtering based on the protocol.

The advantage of circuit level gateways is that they provide services for a wide range of different protocols however they require special client software that has had system calls replaced with secure equivalents from a library such as Socks [Kobl92]. This re-introduces the problem that host based security does not scale well. As the size of the network increases the task of managing secure clients becomes increasingly time consuming and prone to error.

In general application level proxies use modified procedures and circuit level gateways use modified clients.

The simplest firewall architecture utilises a dual homed host. A dual-homed host is a computer that has separate network connections to two networks, as illustrated in figure 3. Such a host could act as a router between the two networks, however, this routing function is disabled when dual-homed hosts are used in firewall architectures.

Because the routing function is disabled the host isolates the two networks from each other whilst retaining the ability to see traffic on both networks. Systems inside the internal network can communicate with the dual homed host via one network interface, and systems on the Internet via the other, however these systems cannot communicate with each other directly.

In a dual homed host architecture the dual homed host itself is critical to the network's security. Such hosts are often referred to as Bastion Hosts in the firewall literature.

A dual homed host can only provide services by proxying them. Where proxies are not available a screened host or screened subnet architecture provide extra options for providing new and/or untrusted services.

6 Example

Firewall, Realplayer and Real Server:

A firewall is used to prevent unauthorized access to a network. A network can be made up of a company's local area networks, wide area networks, and the Internet, or it can be just an Internet Service Provider preventing inappropriate access to the files of its customers.

The firewall's role is to ensure that all communication between an organization's network and the Internet, in both directions, conforms to the organization's security policies.

In general, firewalls permit one-way access to the Internet. Because RealServer and RealPlayer need to establish two-way communication to stream and receive media content, firewalls may reject RealPlayer's attempt to establish this connection, and the RealPlayer's request for a clip will "bounce" off the firewall.

RealNetworks designed both RealPlayer and RealServer to work with a firewall while still protecting the company's internal networks. By making a few quick changes to your firewall, RealServer, RealPlayer, or a combination of the three, you can still use the security advantages of a firewall while enjoying streaming media.

As discussed above, a firewall's main security feature requires that the firewall block two-way communication--but RealPlayer and RealServer need two-way communication. The effect of firewalls on RealPlayer and RealServer is shown in this section.

When no firewall exists between a RealPlayer and RealServer, the RealPlayer first establishes a two-way TCP connection to the RealServer. RealServer uses this connection initially as a means of sending information to the Player about the streamed media, such as the name, length, and copyright of the clip. The Player uses the connection to send commands to RealServer when features such as the "play" and "stop" buttons are activated.

7. Types of Firewalls

There are two types of firewalls.
1. Filtering Firewalls - that block selected network packets.
2. Proxy Servers (sometimes called firewalls) - that make network connections for you.

7.1 Packet Filtering Firewalls

Packet Filtering is the type of firewall built into the Linux kernel.

A filtering firewall works at the network level. Data is only allowed to leave the system if the firewall rules allow it. As packets arrive they are filtered by their type, source address, destination address, and port information contained in each packet.

Many network routers have the ability to perform some firewall services. Filtering firewalls can be thought of as a type of router. Because of this you need a deep understanding of IP packet structure to work with one.

Because very little data is analyzed and logged, filtering firewalls take less CPU and create less latency in your network.

Filtering firewalls do not provide for password controls. User can not identify themselves. The only identity a user has is the IP number assigned to their workstation. This can be a problem if you are going to use DHCP (Dynamic IP assignments). This is because rules are based on IP numbers you will have to adjust the rules as new IP numbers are assigned. I don't know how to automate this process.

Filtering firewalls are more transparent to the user. The user does not have to setup rules in their applications to use the Internet. With most proxy servers this is not true.

7.2 Proxy Servers

Proxies are mostly used to control, or monitor, outbound traffic. Some application proxies cache the requested data. This lowers bandwidth requirements and decreases the access the same data for the next user. It also gives unquestionable evidence of what was transferred.

There are two types of proxy servers.
1. Application Proxies - that do the work for you.
2. SOCKS Proxies - that cross wire ports.

7.2.1. Application Proxy

The best example is a person telneting to another computer and then telneting from there to the outside world. With a application proxy server the process is automated. As you telnet to the outside world the client send you to the proxy first. The proxy then connects to the server you requested (the outside world) and returns the data to you.

Because proxy servers are handling all the communications, they can log everything they (you) do. For HTTP (web) proxies this includes very URL they you see. For FTP proxies this includes every file you download. They can even filter out "inappropriate" words from the sites you visit or scan for viruses.

Application proxy servers can authenticate users. Before a connection to the outside is made, the server can ask the user to login first. To a web user this would make every site look like it required a login.

7.2.2 SOCKS Proxy

A SOCKS server is a lot like an old switch board. It simply cross wires your connection through the system to another outside connection.

Most SOCKS server only work with TCP type connections. And like filtering firewalls they don't provide for user authentication. They can however record where each user connected to.

8. How does a firewall work?

It is easy to let your network get out of hand. Keep control of every connection. It only takes a user with a modem to compromise your LAN

A firewall is a set of related programs, located at a network gateway server that protects the resources of a private network from users from other networks. (The term also implies the security policy that is used with the programs.) An enterprise with an intranet allows its workers access to the wider Internet, installs a firewall to prevent outsiders from accessing its own private data resources and for controlling outside resources, its own users have access to.

Basically, a firewall, working closely with a router program, examines each network packet to determine whether to forward it toward its destination. A firewall also includes or works with a proxy server that makes network requests on behalf of workstation users. A firewall is often installed in a specially designated computer separate from the rest of the network so that no incoming request can get directly at private network resources.

There are a number of firewall screening methods. A simple one is to screen requests to make sure they come from acceptable domain name and Internet Protocol addresses. For mobile users, firewalls allow remote access in to the private network by the use of secure logon procedures and authentication certificates.

A number of companies make firewall products. Features include logging and reporting, automatic alarms at given thresholds of attack, and a graphical user interface for controlling the firewall.

8.1 Understanding Firewalls

A firewall is a structure intended to keep a fire from spreading. Building have firewalls made of brick walls completely dividing sections of the building. In a car a firewall is the metal wall separating the engine and passenger compartments.

Internet firewalls are intended to keep the flames of Internet hell out of your private LAN. Or, to keep the members of your LAN pure and chaste by denying them access the all the evil Internet temptations. ;-)

The first computer firewall was a non-routing Unix host with connections to two different networks. One network card connected to the Internet and the other to the private LAN. To reach the Internet from the private network, you had to logon to the firewall (Unix) server. You then used the resources of the system to access the Internet. For example, you could use X-windows to run Netscape's browser on the firewall system and have the display on your work station. With the browser running on the firewall it has access to both networks.

This sort of dual homed system (a system with two network connections) is great if you can TRUST ALL of your users. You can simple setup a Linux system and give an account accounts on it to everyone needing Internet access. With this setup, the only computer on your private network that knows anything about the outside world is the firewall. No one can download to their personal workstations. They must first download a file to the firewall and then download the file from the firewall to their workstation.

BIG NOTE: 99% of all break-ins start with gaining account level access on the system being attacked. Because of this I don't recommend this type of firewall. It is also very limiting.

9. Design and Implementation Issues

What are some of the basic design decisions in a firewall?

There are a number of basic design issues that should be addressed by the lucky person who has been tasked with the responsibility of designing, specifying, and implementing or overseeing the installation of a firewall.

The first and most important decision reflects the policy of how your company or organization wants to operate the system: is the firewall in place explicitly to deny all services except those critical to the mission of connecting to the Net, or is the firewall in place to provide a metered and audited method of ''queuing'' access in a non-threatening manner? There are degrees of paranoia between these positions; the final stance of your firewall might be more the result of a political than an engineering decision.

The second is: what level of monitoring, redundancy, and control do you want? Having established the acceptable risk level (e.g., how paranoid you are) by resolving the first issue, you can form a checklist of what should be monitored, permitted, and denied. In other words, you start by figuring out your overall objectives, and then combine a needs analysis with a risk assessment, and sort the almost always conflicting requirements out into a laundry list that specifies what you plan to implement.

The third issue is financial. We can't address this one here in anything but vague terms, but it's important to try to quantify any proposed solutions in terms of how much it will cost either to buy or to implement. For example, a complete firewall product may cost between $100,000 at the high end, and free at the low end. The free option, of doing some fancy configuring on a Cisco or similar router will cost nothing but staff time and a few cups of coffee. Implementing a high end firewall from scratch might cost several man-months, which may equate to $30,000 worth of staff salary and benefits. The systems management overhead is also a consideration. Building a home-brew is fine, but it's important to build it so that it doesn't require constant (and expensive) attention. It's important, in other words, to evaluate firewalls not only in terms of what they cost now, but continuing costs such as support.

On the technical side, there are a couple of decisions to make, based on the fact that for all practical purposes what we are talking about is a static traffic routing service placed between the network service provider's router and your internal network. The traffic routing service may be implemented at an IP level via something like screening rules in a router, or at an application level via proxy gateways and services.

The decision to make is whether to place an exposed stripped-down machine on the outside network to run proxy services for telnet, FTP, news, etc., or whether to set up a screening router as a filter, permitting communication with one or more internal machines. There are pluses and minuses to

both approaches, with the proxy machine providing a greater level of audit and potentially security in return for increased cost in configuration and a decrease in the level of service that may be provided (since a proxy needs to be developed for each desired service). The old trade-off between ease-of-use and security comes back to haunt us with a vengeance.

10. Firewall Design Policy

The Firewall Design Policy is a lower-level policy which describes how the firewall will actually go about restricting the access and filtering the services as defined in the network service access policy.

The firewall design policy is specific to the firewall. It defines the rules used to implement the network service access policy. This policy must be designed in relation to, and with full awareness of, issues such as firewall capabilities and limitations, and the threats and vulnerabilities associated with TCP/IP. Firewalls generally implement one of two basic design policies:

Permit any service unless it is expressly denied; or
Deny any service unless it is expressly permitted.

A firewall that implements the first policy allows all services to pass into the site by default, with the exception of those services that the network service access policy has identified as disallowed. A firewall that implements the second policy denies all services by default, but then passes those services that have been identified as allowed. This second policy follows the classic access model used in all areas of information security.

The first policy is less desirable, since it offers more avenues for getting around the firewall. For example, users could access new services currently not denied by the policy (or even addressed by the policy). For example, they could run denied services at non-standard TCP/UDP ports that are not specifically denied by the policy. Certain services, such as X Windows, FTP, Archie, and RPC are difficult to filter. For this reason, they may be better accommodated by a firewall that implements the first policy. Also, while the second policy is stronger and safer, it is more restrictive for users; services such as those just mentioned may have to be blocked or heavily curtailed.

Certain firewalls can implement either design policy but one particular design, the dual-homed gateway, is inherently a "deny all" firewall.

Systems which require services which should not be passed through the firewall could be located on screened subnets separate from other site systems.

In other words, depending on security and flexibility requirements, certain types of firewalls are more appropriate than others, making it extremely important that policy is considered before implementing a firewall. Failure to do so could result in the firewall failing to meet expectations.

10.1 System Specific Policies

System-specific policy is often implemented through the use of access controls. For example, it may be a policy decision that only two individuals in an organization are authorized to run a particular program. Access controls are used by the system to implement (or enforce) this policy

10.2 Incident Handling

When a site that is not protected comes under sustained attack one of two things can happen. The site can rapidly develop a policy and defences or it can withdraw from the Internet. Internet security incidents, such as break-ins and service disruptions, have caused significant harm to several organisations' computing capabilities. Many organisations have an ad hoc response when initially confronted with an attack which can exacerbate the damage caused by the attack. For this reason it is often cost-effective to develop an in-house capability for the quick discovery of, and controlled response to, network security incidents.

The primary benefits of an incident handling capability are the ability to contain and repair damage resulting from network attacks. An incident handling capability also assists an organisation to prevent, or at least to minimise, damage from future incidents. Incidents can be studied internally to gain a better understanding of the organisation's vulnerabilities so that more effective safeguards can be implemented.

10.3 Disaster recovery

It is prudent to assume that an attack may fundamentally compromise an organisation, for example deleting large amounts of data. It is for such eventualities that organisations develop disaster recovery plans. The basic steps in establishing a disaster recovery plan are :-

Identify the mission or business critical functions.
Identify the resources that support the critical functions.
Anticipate potential contingencies or disasters.
Select contingency planning strategies.
Implement the contingency strategies.
Test and revise the strategy.

(1) Serial Line Internet Protocol - A means of using IP over serial (telephone) lines.
(2) Point to Point Protocol - A replacement for SLIP
(3) For example an electrical switch should always fail to the open (i.e. off) position and break any circuit.
(4) A detailed discussion of the development and role of a disaster recovery plan is beyond the scope of this report. Interested readers should consult.

Internet firewalls are a means of protecting networks by implementing access control to and from the Internet. In practice this is achieved by controlling the means of communication between the two networks, the TCP/IP suite of protocols. Firewall means the strategies and policies and the term Firewall System to refer to the hardware and software elements that implement the policy.

Note that in practice an Internet firewall is more like a moat around a castle than a firewall in a modern building.

A Firewall System is a collection of components that is placed between two networks and possesses the following properties :

All traffic from inside to outside, and vice-versa, must pass through it [Chap95].
Only authorized traffic, as defined by the security policy, is allowed to pass through it.
The system itself is immune to penetration [Ches94]

In other words a Firewall System is a mechanism used to protect a trusted network whilst it is connected to an untrusted network.

Typically, the two networks in question are an organization's internal network (trusted) and the Internet (untrusted). But there is nothing in the definition of a firewall that ties the concept to the Internet. Although the majority of firewalls are currently deployed between the Internet and internal networks, there are good reasons for using firewalls when connecting any trusted network with a less trusted network, be it internal or external.

11. Limitations

11.1 What can't a firewall do ?

Whilst firewalls provide good protection at the lower levels of the TCP/IP model, they provide almost no protection against higher level protocols.

Any data that is passed by the firewall still has the potential to cause problems which, were these to be exploited deliberately would be labelled as denial of service or data driven attacks. For example a firewall offers no protection against viruses contained in files transferred via ftp or as MIME attachment to an e-mail message.

A firewall can't protect against malicious insiders. A firewall cannot differentiate between hosts on the same side of a network therefore any Internet Host can spoof any other Internet Host and any internal host can spoof any other internal host.

A firewall can't protect against connections that don't go through it (i.e. backdoors). Firewalls can restrict the access to certain facilities and users will sometimes bypass the firewall to gain access to those facilities. A good example would be a firewall that didn't allow access to the World Wide Web. Users on that network may establish point to point connections with an Internet service provider over a normal telephone line and introduce Internet connectivity behind the firewall. This type of threat can only be addressed by management procedures which are embodied in the organisations security policies.

A firewall can't protect against completely new threats if the security strategy is different from "deny everything unless specifically permitted." Again this is dealt with within the security policy by basing it on just such a strategy.

11.2 What about viruses?

Firewalls can't protect very well against things like viruses. There are too many ways of encoding binary files for transfer over networks, and too many different architectures and viruses to try to search for them all. In other words, a firewall cannot replace security-consciousness on the part of your users. In general, a firewall cannot protect against a data-driven attack--attacks in which something is mailed or copied to an internal host where it is then executed. This form of attack has occurred in the past against various versions of *sendmail*, *ghostscript*, and scripting mail user agents like *OutLook*.

Organizations that are deeply concerned about viruses should implement organization-wide virus control measures. Rather than trying to screen viruses out at the firewall, make sure that every vulnerable desktop has virus scanning software that is run when the machine is rebooted. Blanketing your network with virus scanning software will protect against viruses that come in via floppy disks, modems, and Internet. Trying to block viruses at the firewall will only protect against viruses from the Internet--and the vast majority of viruses are caught via floppy disks.

Nevertheless, an increasing number of firewall vendors are offering ''virus detecting" firewalls. They're probably only useful for naive users exchanging Windows-on-Intel executable programs and malicious-macro-capable application documents. There are many firewall-based approaches for dealing with problems like the ''ILOVEYOU" worm and related attacks, but these are really oversimplified approaches that try to limit the damage of something that is so stupid it never should have occurred in the first place. Do not count on any protection from attackers with this feature.

A strong firewall is never a substitute for sensible software that recognizes the nature of what it's handling--untrusted data from an unauthenticated party--and behaves appropriately. Do not think that because ''everyone" is using that mailer or because the vendor is a gargantuan multinational company, you're safe. In fact, it isn't true that ''everyone" is using any mailer, and companies that specialize in turning technology invented elsewhere into something that's ''easy to use" without any expertise are more likely to produce software that can be fooled.

12. Conclusion:

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 intruder would be motivated to break into their computer. Dial up users who have been victims of hazardous 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.

In order to prevent such attacks a firewall should be used which is a group of systems that enforces an access control policy between two networks. The actual means by which this is accomplished varies widely, but in principle, the firewall can be thought of as a pair of mechanisms: one which exists to block traffic, and the other which exists to permit traffic. Some firewalls place a greater emphasis on blocking traffic, while others emphasize permitting traffic. Probably the most important thing to recognize about a firewall is that it implements an access control policy.

Firewalls are also important since they can provide a single ''choke point" where security and audit can be imposed. Unlike in a situation where a computer system is being attacked by someone dialing in with a modem, the firewall can act as an effective ''phone tap" and tracing tool. Firewalls provide an important logging and auditing function; often they provide summaries to the administrator

about what kinds and amount of traffic passed through it, how many attempts there were to break into it, etc.

13. References

1. Internet Security,IBM
2. Internet Encyclopedia
3. www.ieee.org
4. www.vicomsoft.com
5. www.vnunet.com
6. www.cisco.com
7. www.interhack.net
8. www.firewall.com
9. www.linuxsecurity.com
10. www.pctoday.com
11. www.enteract.com
12. www.robertgraham.com
13. www.linuxdoc.org
14. www.service.real.com
15. www.venus.soci.niu.edu