NEAR EAST UNIVERSITY

Faculty Of Engineering

Depatment Of Computer Engineering

INTERNET SECURITY

GRADUATION PROJECT
COM400

Student: Serkan ŞEKER(990550)
Supervisor: Asist.Prof.Dr.Firudun MURADOV

Lefkoşa 2004

ACKNOWLEDGEMENTS

First, I want to thanks my family, especially my parents. Without their endless support and love for me, I would never achieve my current position. I wish my mother and my father lives happily always.

I want to thanks Asist.Prof.Dr.Firudun Muradov to be my supervisor. Under his guidence, I succesfully overcome many difficulties and learn a lot about internet security , In each discussion, he explained my questions patiently, and I felt my quick progress from his advices. He always helps me a lot either in my study or my life. I asked him many questions in my subject and he always answered my questions quickly and in detail

I also want to thank to my teachers in Near East University

ABSTRACT

Internet's uncontrolled growth, and the increasing dependence on it by many commercial organizations, has made it a very critical but also a very vulnerable infrastructure.We start with an overview of security requirements such as privacy, integrty, authentication, threats such assniffing, spoofing, denial of services and vulnerabilities such as virus, Trojan Horse. We then separate the three main components of security:Data security, System security and Transmission/Network Security and present useful defence mechanisms used at each layer.

At the heart of most defence mechanisms are Cryptographic protocols, which are security related interactions for reaching agreement between two or more principals, for example for authentication and key distribution. Such protocols are usually described by alternate transmission and receipt of (encrypted) messages in pre-defined format and sequence between principals. We conclude by introducing several logics that have been proposed to model cryptographic protocols and corresponding proof techniques used to prove properties of the protocol such as vulnerability to attack by interception or spoofing of messages.

The role of a security policy is to ensure that each of the four fundamental components that make up computer security, Authentication, Access Control, Integrity and Confidentiality are adequately addressed.

The integration of firewall security policy in network security policy should be fulfilled by integrating the security management in a network management system. The desired architecture of an active firewall component can be inferred from the goals of the firewall systems.

TABLE OF CONTENTS

INTRODUCTION

Firewalls have gained great fame recently as the ultimate in Internet Security. This project describes a comprehensive introduction to Internet Security, and present firewalls as the primary means by which organisations can manage the risks associated with connecting their network to the Internet. It includes brief introduction to Internet Security and to the TCP/IP protocols, and uses the TCP/IP model as a framework to discuss security weaknesses in the protocols that make up the Internet suite. Security risks are discussed in terms of generic computer security and in terms of Internet specific attacks. The nature and role of a Network Security Policy is presented as an organisation's central strategy for countering the risks (a firewall's function being to enforce this policy).

An extensive description of Firewalls includes both the theory of how they function, and description of the common architectures

Chapter 1 presents an Introduction to Internet Security and the four of it, that areAuthentication, Access Control, Integrity and Confidentiality.

Chapter 2 provides a brief introduction to the TCP/IP Protocols. The TCP/IP layer model is described as a framework to show security weaknesses in the protocols.

Chapter 3 describes generic Computer Security Risks and specific Internet Attacks. It also provides Security Analysis Tools to test our network.

Chapter 4 analysis at the nature and role of a Network Security Policy, and gives that possible solution may be Firewall design.

Chapter 5 provides an introduction to Firewalls, and presents Firewall theory and their components such as Packet filter, Application gateway and Security management.

Chapter 6 presents Firewall Architectures. It also shows us that the best solution is to built a firewall with combinations of components according to problems that we have in our site.

Chapter 7 describes the most important attacks on Firewall components and the most important attacks that are caused by concept errors.

Chapter 8 describes how to configure some internet services to run with a firewall, including IIOP configuration.

CHAPTER ONE:INTRODUCTION TO INTERNET SECURITY

Any one responsible for the security of a trusted network will be concerned when connecting it to a distrusted network. In the case of connections to the Internet this concern may be based largely on anecdotal evidence gleaned from widespread media coverage of security breaches. A closer inspection of the facts and statistics behind some of the media coverage will, however, only serve to deepen that concern. For example, the US National Computer Security Agency (NCSA) asserts that most attacks to computer systems go undetected and unreported, citing attacks made against 9000 Department of Defence computers by the US Defence Information Systems Agency (DISA). These attacks had an 88 per cent success rate and went undetected by more than 95 per cent of the target organisations. Only 5 percent of the 5 per cent that detected an attack, mere 22 sites, reacted to it.

NCSA also quote the FBI as reporting that in more than 80 percent of FBI investigated computer crimes, unauthorised access was gained through the Internet. There is amazement in the computer security industry at the level of ignorance to the problem. To understand the risks often involves a steep learning curve and they have few real parallels in everyday life, for example nobody worries that a burglar will be able to trick their front door into opening by posting cryptic messages through the letterbox.

When there is a good "hacker" story to report the press goes into frenzy, but the general level of awareness is still surprisingly low. For example the Sunday Times which prides itself on providing accurate coverage of IT issues that claimed that most businesses worry too much about Internet security. The article goes on to explain that encryption is all that is needed to be completely secure. The article focuses purely on privacy of communication and completely misses the possibility of an attack originating from the Internet.

Despite fears about security, organisations are increasingly coming to regard a presence on the Internet as an important part of their strategic planning. Security concerns will not be allowed to prevent organisations from exploiting the commercial opportunities the Internet is perceived to offer. As a result organisations have to find ways to manage the security issue. This ties growth in the Internet security market directly to growth in the Internet.

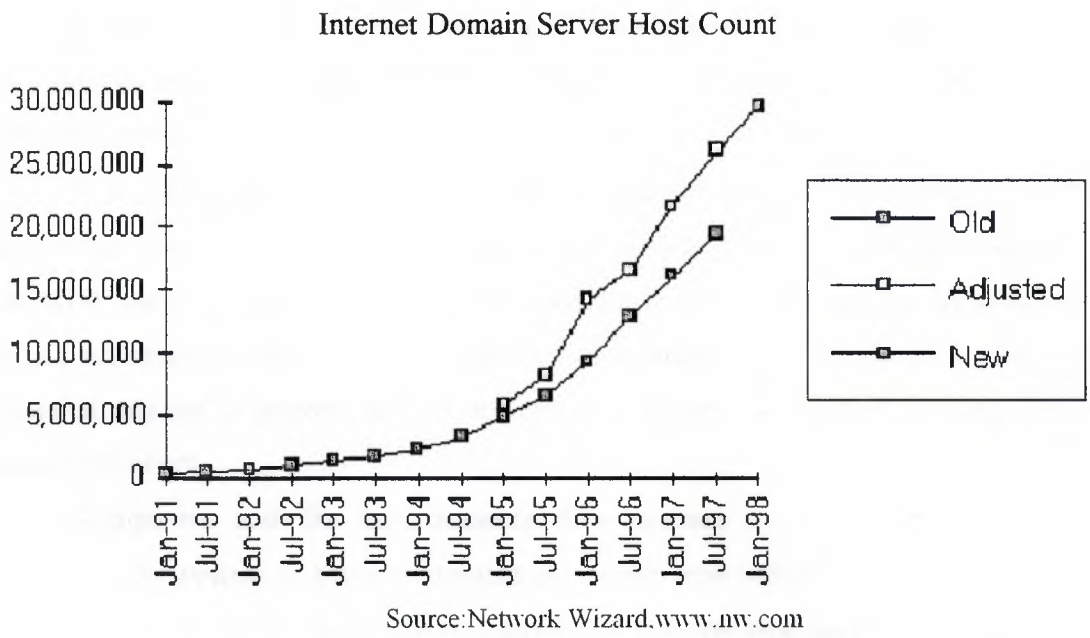

Figure 1.1 Growth of the Internet

The compound annual growth rate (CAGR) of the Internet firewall market between 1995 and 2000 is projected to be 174% driven by rapid growth of both the Internet (see Figure 1.1), and Intranets .

Given that approximately 40% of the fortune 500 companies using the Internet have still to install a firewall and that the Internet continues to double annually, it is little surprise that the security auditing business is booming. Organisations are finding that they do not have the in-house skills or knowledge necessary to assess either the current situation or the potential risks, and are wrestling with what level of security they require. The rest of this chapter investigates what is meant by the term Network security - often the starting point when an organisation calls in an external consultant .

1.1 What is Internet Security?

The hardware, software and information that constitute computer systems is increasingly mission critical. Protecting them can be as important as protecting other valuable resources, such as money, buildings, or employees. The purpose of computer security is to protect computer resources through the selection and application of appropriate safeguards.Internet security protects computer resources against the risks and threats that arise as a result of a connection to the Internet.

Computer security supports the organisation's mission by protecting its physical and financial resources, reputation, legal position, employees, and other tangible and intangible assets.

If a system has external users then its owners have a responsibility to share appropriate knowledge about the existence and general extent of security measures so that other users can be confident that the system is adequately secure. In addition to sharing information about security, organisation managers "should act in a timely, co-ordinated manner to prevent and to respond to breaches of security" to help prevent damage to others.

Computers and the environments they operate in are extremely dynamic. Changes in the system or the environment can create new vulnerabilities and it is almost inevitable that a system's users and operators will discover new ways to intentionally or unintentionally bypass or subvert security. It is therefore necessary to reassess the security of computer systems regularly to provide effective computer security.

Providing effective computer security requires a comprehensive approach that considers a variety of areas both within and outside of the computer security field and that extends throughout the entire information life cycle.

There are three general areas of concern when a trusted network is attached to a distrusted network:

1. That inappropriate material will deliberately, or inadvertently, be passed to and from the distrusted network;
2. That unauthorised users will be able to gain access to the trusted network from the distrusted network;
3. That the operations of the trusted network may be disrupted as a result of attack from the distrusted network.

The computer and network security measures that are taken by an organisation are intended to minimise the potential for these to occur by means of the four fundamental components that make up computer network security:

1. Identification and Authentication
2. Access Control
3. Integrity
4. Confidentiality

1.1.1 Identification and Authentication

The first component of computer security is authentication, or ensuring that users and computers are who they claim to be by establishing proof of identity. This is usually based on one, or a combination, of something you are (a biometric i.e., such characteristics as a voice pattern, handwriting or a fingerprint), something you know (a secret i.e., a password, Personal Identification Number (PIN), or cryptographic key), or something you have (a token i.e., an credit card or a smart card). For example acquaintances can authenticate your identity (to a point) based on your physical features[1]. Banks authenticate you based on something you have such as your credit card, and something you know, often your mother's maiden name. One-time passwords are passwords that can only be used once then expire, are generally based on something you have. An example of this type of authentication is the one-time pads[2] used by the intelligence services during the Second World War.

The lack of strong authentication has inhibited the development of electronic commerce. It is still necessary for contracts, legal documents and official letters to be produced on paper. Strong authentication is then, a key requirement if the Internet is to be used for electronic commerce.

Authentication is an important part of everyday life. Letters are printed on headed paper and signed by the author. Digital signatures fulfil a similar requirement; although they are much more trustworthy as they are based on mathematical encryption algorithms and attest to the contents of a message as well as its author. Digital signatures are based on public key (asymmetric) encryption. Whitfield Diffie and Martin Hellman in order to solve the key management problem that exists with secret key or symmetric encryption introduced the concept of public key encryption in 1976[3]. Asymmetric cryptography uses key pairs, one key in the key pair is called the public key and the other is called the private key. Either key can be used to encrypt the message, but once encrypted only the other key in the pair can be used to decrypt it. It is immediately apparent that two scenarios are possible, one where the private key is used to encrypt the message and hence the public key is used to decrypt it, and vice versa. By encrypting the message using the receiver's public key, the sender is assured that only the receiver can decrypt it confidentially. To digitally sign a message the sender passes the message through a hashing algorithm[4] to produce the message digest, which he then encrypts with his private key. The output is called a digital signature and is attached to, and sent with, the message. In order to verify the signature the receiver also passes the

message through the same hashing algorithm to re-create the message digest, and then decrypts the sender's digital signature using the sender's public key. If the message did not originate from the sender, or if its contents were altered, then the two digests will not match.

Under normal circumstances the private key is kept secret by the individual, but the public key is distributed as required. There is no need for the sender and receiver to share a secret key; however, asymmetric encryption key management still requires public keys to be distributed in an authenticated or trustworthy manner. One means of achieving this is to use a certification authority. The main attribute of a certification authority is that it is trusted by a group of users to create certificates on their behalf. The certification authority verifies a user's public key by digitally signing it. This creates a certified public key, referred to as a certificates. The certification authority's digital signature attests that the public key is valid, and guarantees that it cannot be altered in any way.

One such certification authority is VeriSign Incorporated who began issuing key pairs and certificates in late April 1996 and has trademarked the term "Digital ID". Security aware applications are required to make use of certificates and secure e-mail tools and browsers for the World Wide Web are now becoming available. When the certificate details have been installed in the client software (i.e. browser) they are automatically provided along with the client's requests allowing the server to authenticate you.

Identification and Authentication is a critical building block of computer security since it is the basis for most types of access control and for establishing user accountability.

1.1.2 Access Control

Access is the ability to do something with a computer resource (i.e., use, change, or view). Access controls are the means by which the ability is explicitly enabled or restricted in some way (usually through physical and system-based controls).

Access control often requires that the system be able to identify and differentiate users. For example, access control is often based on least privilege, which refers to the granting to users of only those accesses required to perform their duties. User accountability requires the linking of activities on a computer system to specific individuals and, therefore, requires the system to identify users.

Access controls provide technical means of controlling what information users can utilise, the programs they can run, and the modifications they can make. Computer based access controls are called logical access controls. Logical access controls can prescribe not only who or what (i.e., in the case of a process) is to have access to a specific system resource but also the type of access that is permitted. These controls may be built into the operating system, may be incorporated into applications programs or major utilities (i.e., database management systems or communications systems), or may be implemented through add-on security packages. Logical Access controls may be implemented internally to the computer system being protected or may be implemented in external devices.

Logical access controls can help to protect:

- operating systems and other system software from unauthorised modification or manipulation (and thereby help ensure the system's integrity and availability)
- integrity and availability of information by restricting the number of users and processes with access
- confidential information from being disclosed to unauthorised individuals. The concept of access modes is fundamental to access control. Common Access modes, which can be used in both, operating systems and applications, include the following:
- Read
- Write
- Delete
- Create
- Execute

In deciding whether to permit someone to use a system resource, logical Access controls examine whether the user is authorised for the type of access requested based on access criteria such as:

- Identity - It is probably fair to say that the majority of access controls are based upon the identity of the user (either human or process), which is usually established through identification and authentication.
- Roles - Access to information may also be controlled by the job assignment or function (i.e., the role) of the user who is seeking access. Examples of roles include data entry clerk, purchase officer, project leader and programmer. Access rights are grouped by role name, and the use of resources is restricted to individuals authorised to assume

the associated role. An individual may be authorised for more than one role, but may be required to act in only a single role at a time. Changing roles may require logging out and then in again, or entering a role changing command. The use of roles can be very effective means of providing access control.

- Location - Access to particular system resources may also be based upon physical or logical location for example, users can be restricted based upon network addresses (i.e., users from sites within a given organisation may be permitted greater access than those from outside).
- Time - Time of day or day of week restrictions are common access limitations. For example, use of confidential personnel files may be allowed only during normal working hours and denied at all other times.
- Transaction - Another approach to access control can be used by organisations handling transactions (i.e., account inquiries). Phone calls may first be answered by a computer that requests that callers key in their account number and perhaps a PIN. Some routine transactions can then be made directly, but more complex ones may require human intervention. In such cases, the computer, which already knows the account number, can grant a clerk, for example, access to a particular account for the duration of the transaction. When completed, the access authorisation is terminated. This means that users have no choice in which accounts they have access to, which can reduce the potential for mischief. It also prevents users from casually browsing through accounts thereby improving confidentiality.
- Service Constraints - Service constraints refer to those restrictions that depend upon the parameters that may arise during use of the application or that are preestablished by the resource owner/manager. For example, a particular software package may only be licensed by the organisation for five users at a time. Access would be denied for a sixth user, even if the user were otherwise authorised to use the application. Access may also be selectively permitted based on the type of service requested. For example, users of computers on a network may be permitted to exchange electronic mail but may not be allowed to log into other computers.
- External Access Controls - External access controls are means of controlling interactions between the system and outside people, systems, and services. External access controls use a wide variety of methods; often including firewalls as will be discussed in later chapters.

1.1.3 Integrity

Integrity is the degree to which something is free from corruption, i.e. whether or not something has been damaged, altered, added or removed. In addition to improving authentication, digital signatures also improve the level of confidence in the integrity of a message as discussed earlier in this chapter. However, Integrity does not apply only to messages. The integrity of files and applications is also very important. One of the most common means of gaining unauthorised access to a computer system is to install altered copies of operating system programs that provide access to the intruder when they are executed[6]. It is important therefore that the integrity of operating system components can be verified. Attackers themselves understand this well, as is illustrated by which describes how an attacker who, while being monitoring began, immediately upon discovering that one of his back door programs had been removed, to compare copies of other files he had replaced with the originals that he had stored elsewhere.

The integrity of anti virus software should also be verified regularly. Most packages on the market perform a self-verification of their integrity. The problem with this is that rogue software would presumably not be designed to point out that it differed from the original. In cases such as this verification of integrity should be independent in order to be trustworthy.

In some cases the integrity of data files is also often assumed to be verified by the application software. While the application software will generally notify the user of damage or corruption to a file it will not generally report that Company A has been removed from a list of companies tendering for a major contract for example. Again integrity needs to be verified independently.

Both message digests and digital signatures can attest to the integrity of files in all of these cases. The point is that in order to be trusted independent verification is required.

1.1.4 Confidentiality

Confidentiality is the degree to which the privacy or secrecy of something can be trusted. The confidentiality of most paper-based communication is entrusted to envelopes. Most messages transmitted over the Internet cannot claim even this level of confidentiality, being more akin to postcards. The lack of privacy (or confidentiality) on the Internet applies equally to files transferred over it, and information moving to and from World Wide Web clients and servers. E-mail, File Transfer and World Wide Web

applications accounted for approximately half of the bytes transferred on the Internet backbone in 1994. Regardless of what the data was, the vast majority of this traffic was transmitted without any regard for its confidentiality.

Initiatives to correct this state of affairs have been underway for some time and are likely to come to fruition in 1996. For example, Web Browsers that are able to use certificates and therefore make use of the Privacy Enhanced Mail (PEM) standard, and Secure Multipurpose Internet Mail Extensions (S/MIME) standards.

CHAPTER TWO: TCP/IP PROTOCOL

The name TCP/IP refers to an entire suite of data communications protocols. The suite gets its name from two of the protocols that belong to it: Transmission Control Protocol and the Internet Protocol. Although there are many others protocols in the suite, TCP and IP are certainly two of the most important.

The popularity of the TCP/IP protocols on the Internet did not grow rapidly just because the protocols were there, or because military agencies mandated their use. They met an important need (worldwide data communication) at the right time, and they had several important features that allowed them to meet this need.These are:

- Open protocol standards, freely available and developed independently from any specific computer hardware or operating system.
- Independence from specific physical network hardware.
- A common addressing scheme that allows any TCP/IP device to uniquely address any other device in the entire network, even if the network is as large as the world-wide Internet.
- Standardised high-level protocols for consistent, widely available user devices.

2.1 History of TCP/IP

In 1969 the US Defence Advanced Research Projects Agency (DARPA) began funding a project to develop a high speed, packet switching communications network to link its research centres and laboratories. The system became known as ARPANET and was one of the first communications systems to utilise a layered architecture, preceding the ISO OSI reference model by almost a decade. Although the ARPANET was a general success, its first generation protocols were expensive to implement, slow and prone to crashes (both of the individual stacks and of the networks themselves). In 1974 a new set of core protocols was proposed by Vinton Cerf and Robert Kahn [Cerf74]. This proposal was the base for the development of the Internet Protocol (IP) and the Transmission Control Protocols (TCP). Over a period of three years the ARPANET hosts migrated to use these protocols. TCP/IP was better suited to inter-networking than he Xerox Networking System (XNS) protocol stack, the other major protocol stack available at that time, for two reasons. Firstly it utilised a defined routing hierarchy that allowed large inter-networks to be managed in a structured way, and secondly, its addresses were centrally administered so duplicates could only be the result of error.

DARPA funded the integration of the TCP/IP protocols into the University of California's Berkeley Software Distribution (BSD) version of UNIX. Version 4.2 of the BSD UNIX released in September 1983 was the first to include the TCP/IP protocols in the generic operating system, and this was eventually carried over into commercial versions of UNIX. SUN Microsystems later published their Open Network Computing (ONC) Standards, better known as the Network Filing System (NFS). NFS is designed to utilise a TCP/IP stack and has since been widely licensed.

TCP/IP has two major shortcomings at present. The first is that address space is limited and will eventually run out. The second is that there are a number of security weaknesses. To explore the latter it is necessary to understand the layered architecture of TCP/IP.

2.2 TCP/IP Architecture

TCP/IP has a layered architecture characterised by increasing abstraction as we move up the layers. Entities within each layer provide services to those in the layer above it, and request services from those in the layer below it. The TCP/IP model consists of five layers[1]:

1. The *Application Layer* provides the application program or process. Examples of Internet application layer protocols are File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Network News Transfer Protocol (NNTP), and Hypertext Transfer Protocol (HTTP). An application engages network services from the TCP or UDP transport layers through one of several APIs such as Berkeley Sockets or the Transport Layer Interface.
2. The *Transport Layer* provides two types of service to the *Application Layer*. The first is a connection-oriented, full duplex service provided by the Transmission Control Protocol (TCP). The second is a connectionless service provided by the User Datagram Protocol (UDP).
3. The *Network Layer* is responsible for moving data between communicating endpoints. Unless the sending and receiving hosts are on the same network this job involves routing - determining and delivering the data along the best inter-network path. The primary protocol in use at the network layer is the connectionless Internet Protocol (IP). The basic unit of communication at this layer is called the IP datagram (sometimes referred to as an IP Packet).

4. The *Data Link Layer* controls the link between 2 nodes in a network by controlling the flow of data to and from the physical layer, the detection and correction of errors in the data, and sequencing.

5. The *Physical Layer* is the lowest layer of the model. This layer sends and receives a stream of bits across the physical medium that connects the systems.

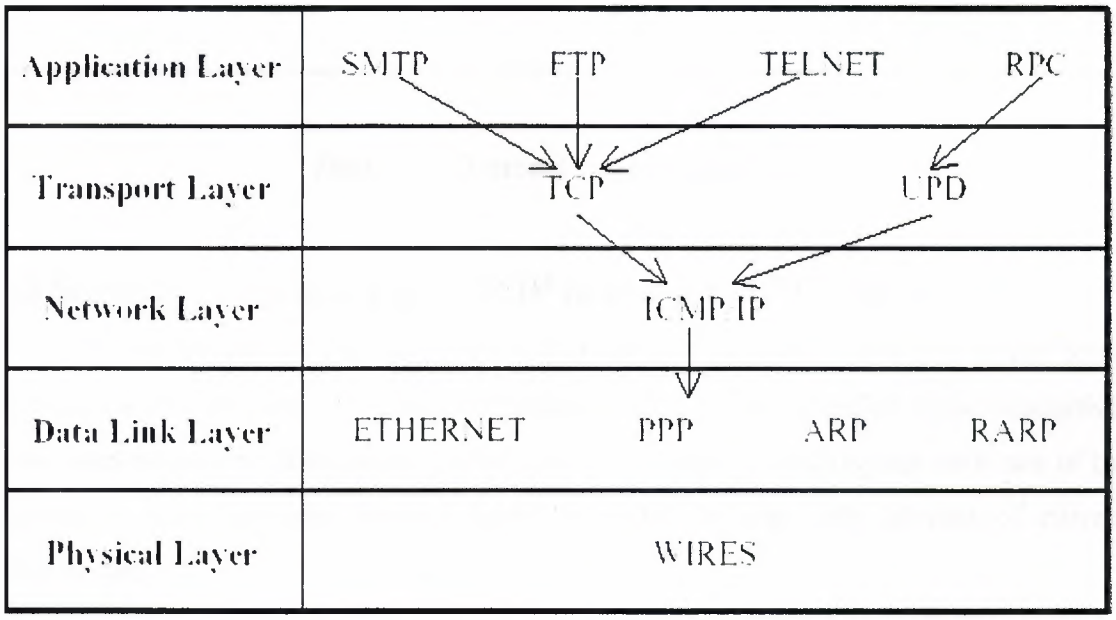

Figure 2.1 TCP/IP Layer Model

A message is transferred across the network by passing down the layers on the sending host and back up them on the receiving host. The message is packaged in a new envelope by each layer before it passes it down (see figure 2.2). After transmission cross the network the message is passed up the layers with each layer removing its respective envelope. The original message is finally passed to the application layer on the receiving host. Armed with an understanding of the layered nature of Internet communications we are now equipped to explore some of the security issues intrinsic to each of the layers.

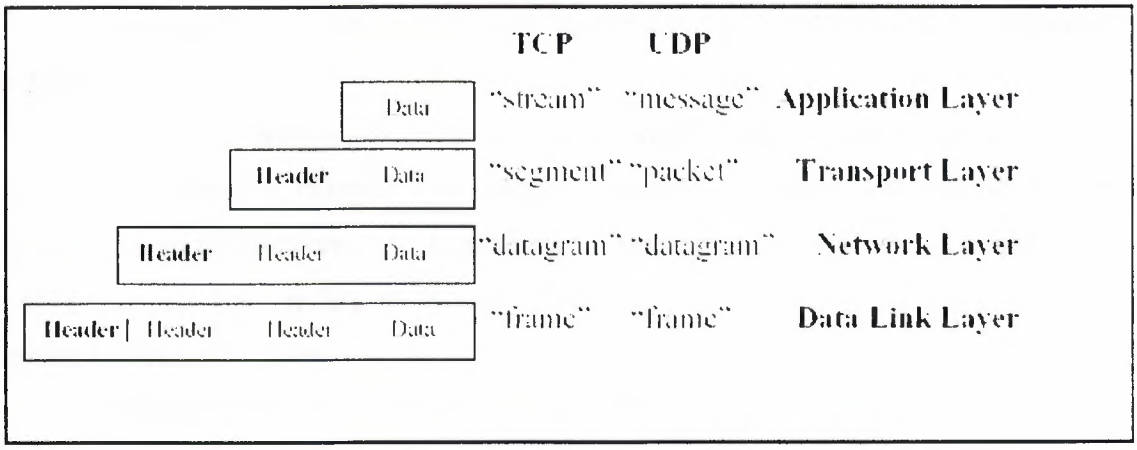

Figure 2.2 Protocol Enveloping Model

2.3 Security Weaknesses in TCP/IP in context of TCP/IP Model

There are security weaknesses at each layer of the model and an attacker could exploit any one of them. The analysis presented here is not intended to be exhaustive. New weaknesses are discovered regularly and the reader should register with one of the Internet security advisory services such as CERT to stay fully abreast of current weaknesses.

2.3.1 Network Layer

Security issues at the *Network Layer* are related to end-to-end delivery of a datagram or IP Packet. The issues include Network Snooping, Message Replay, Message Alteration, essage Delay and Message Denial.

Network Snooping - An earlier analogy likened messages travelling over the Internet to messages written on a post card. Just as everyone who handles a postcard can choose whether or not to read the message, so any system on a network with a shared transmission medium can read every network datagram, whether addressed to it or not. nooping or Sniffing is a passive attack; i.e. the attacker observes network traffic but does not disturb it.

Sniffer software is readily available [Tard95], and is used legitimately as a networktroubleshooting tool. However sniffer software is also widely used by attackers o collect account names and passwords. Sniffing works by placing a system's network interface software into promiscuous mode. All packets are then passed to the sniffer software, which can then display or record the information.

Message Replay - An attacker can sometimes replay traffic recorded with sniffer software to attack a computer.

Message Alteration - The importance of Integrity was discussed in chapter 1.1.3 (page 11). There is currently no widely implemented means of guaranteeing the integrity of an IP datagram. An attacker who modifies the contents of a datagram can also recalculate and update its header checksum, and the datagram recipient will be unable to detect the change.

Message Delay and Denial - An attacker can delay or deny IP messages by changing the screening and routing rules used by routers, or by overwhelming one of the end systems with large amounts of network traffic. Routers and hosts have to iscard incoming packets if they have no remaining buffer space. If the packets contain UDP datagrams they are lost forever. If the packets contain TCP segments the reliable transport mechanism eventually recovers by requesting retransmission of the lost data.

Authentication at the *network layer* is concerned with identifying computer systems rather than computer users. The identity of a computer system on the Internet is its IP address. There is no widespread means of authenticating an IP address and this results in attacks such as Address Masquerading and Address Spoofing.

Address Masquerading occurs when an attacker configures his network interface with the same address as another system. This can gain the attacker access to resources intended for the true owner of the IP address since access to some services, such as NFS, is only contingent upon the use of a "correct" network address[2]. Address masquerading is limited to machines on the same network.

Address spoofing attacks a weakness of TCP but its net effect is at the network layer. It is a sophisticated attack on the three-way handshake that is used to establish a reliable transport connection between two systems. An address spoofing attack exploits several weaknesses

- The trust relationship between two computers.
- The predictability of TCP's Initial Sequence Number (ISN) which plays an important part in the ordering of all subsequent exchanges in the conversation.
- The weak (IP address based) authentication used by some commands.

Routing Attacks - Most routing protocols are susceptible to false route update messages as they do not use secure authentication mechanisms. IP also supports a source routing

option that allows an attacker to specify the routing path packets travel along to their destination.

Tunnelling - IP can be encapsulated within TCP, a technique known as tunnelling. Such tunnels can then be used to bypass the firewall. The extent of damage done depends upon how routing information is propagated.

2.3.2 Transport layer

TCP segments and UDP datagrams are transmitted across the network as IP datagrams and so many of the security weaknesses already discussed apply, specifically the authenticity, integrity and confidentiality of TCP and UDP messages is not guaranteed. If means of solving these problems at the *network layer* were to be found, then these issues in the *Transport layer* protocols would similarly be resolved. There are however some weaknesses introduced in the transport layer itself that allow attacks, such as UDP packet storm denial of service attacks [CA-96:01] and TCP Session Hijacking [CA-95:01].

UDP packet storm denial of service attacks - [CA-96:01] reports of programs that launch denial of service attacks by creating a "UDP packet storm" either on a system or between two systems. An attack on one host causes that host to perform poorly. An attack between two hosts can cause extreme network congestion in addition to adversely affecting both hosts' performance. When a connection is established between two UDP services, and each of them produces output, a very high number of packets can be produced. This can be used to effect a denial of service attack on the machines providing the UDP services. Anyone with network connectivity can launch such an attack as no account access is needed.

Session Hijacking - Having gained root access to a system, an attacker can use a type of tool to dynamically modify the UNIX kernel. This modification allows the attacker to hijack existing terminal and login connections from any user on the system. An attacker can bypass one-time passwords and other strong authentication schemes by taking over the connection after the authentication is complete. An attacker can gain access to remote sites by hijacking the connection after the user has completed the authentication to the remote location.

2.3.3 Application Layer

Application layer security issues include all of the issues discussed so far. However even if a totally secure "pipe" between two systems could be provided, it would not assist in the authentication of a remote user, or in preventing attacks targeted on application layer protocols such as SMTP, FTP or HTTP. These are discussed in the next section.

2.4 Internet Services and Security

After receiving incoming data from IP, the transport protocol (TCP or UDP) passes it to the correct application process. Application processes (also called network services) are dentified by 16 bit values called port numbers. TCP ports are defined by non-negative numbers in the range 0 to 65 635. "Well known" TCP port numbers (i.e.: ftp - 20, telnet -23 ...) are the first 1024 ports (0-1023), which are managed and assigned by the Internet Assigned Numbers Authority (IANA). The first header word of each TCP segment or UDP packet contains the source port number identifying the application that sent the data and the destination port number identifying the application that is to receive the data.

2.4.1 Domain Name System (DNS)

DNS is a distributed database system used to match host names with IP addresses. A host normally requests the IP address of a given domain name by sending a UDP message to the DNS server which responds with the IP address or with nformation about another DNS server. [Ches94] discusses several weaknesses in DNS software, concluding that machines should use address-based authentication, which, lthough weak, is far better than name based authentication. [Ches94] also notes that DNS contains a wealth of information about a site that can be useful to an attacker.

2.4.2 Simple Mail Transfer Protocol (SMTP)

While SMTP is fairly innocuous[3], the most common implementation of it, the sendmail program, is a security nightmare. Sendmail violates the principle of least privilege, discussed in chapter 4.1.1 (page 40) as it runs with root privilege. Privileged programs should be as small and modular as possible. Sendmail consists of tens of thousands of lines of C code. The content of mail messages can also pose dangers.

Automatic execution of messages, the ability to mail executable programs and the ability to mail postscript files are all dangers with MIME

2.4.3 Telnet

Telnet provides simple terminal access to a host computer. The user is normally authenticated based on user name and password. Both of these are transmitted in plain text over the network however, and are therefore susceptible to capture. Recommends that strong authentication be used when establishing a telnet session. Strong authentication does not protect the rest of the session from attackers however, who can hijack it or capture packets using techniques already discussed. Symmetric encryption of telnet sessions presents the problem of a distrusted computer being provided with the key. One solution is to use asymmetric encryption for key management of telnet sessions.

2.4.4 The Network Time Protocol (NTP)

NTP is used to synchronise a host's clock with reference clocks to within 10ms or less. This is extremely useful when analysing an attack using the log files of different computers involved. Describes NTP as "an absolute, non-negotiable requirement" because of the importance of synchronised log files.

NTP itself can be the target of various attacks aimed at altering the host time. More recent versions of NTP support cryptographic authentication, but as this is provided on a hop-by-hop basis, a host upstream may still be attacked.

2.4.5 Finger and Whois

Two standard protocols, finger and whois, are commonly used to look up information about individuals .

The finger protocol in particular presents security problems, not in its implementation, but in the information it provides. Farmer and Venema call finger "one of the ost dangerous services, because it is so useful for investigating a potential target".

As firewalls do not generally allow users to login there is little point in providing the finger service. However the wider point is still noteworthy, i.e. that services that provide information about a network's users are of immense value to an attacker.

2.4.6 Remote Procedure Call Based Protocols

There are several remote procedure call protocols known as RPCs. The most popular, Sun RPC, was developed by Sun Microsystems [Sun88, Sun90] and is generally referred to as simply "RPC".

RPC is layered on top of TCP or UDP. However it is used as a general purpose transport protocol in the same way as TCP and UDP by a variety of application protocols such as Network File System (NFS) and Network Information Service (NIS). NFS and NIS are vulnerable services from a security point of view; for example an attacker with access to an NFS server can probably read any file on the system. When an RPC based service such as an NFS server starts it allocates itself a random TCP or UDP port (some use TCP, some use UDP, and some use both). It then contacts an RPC service called the portmapper. The portmapper, which is allocated to the well-known port 111, registers its RPC service number and the particular port(s) it is using at present.

An RPC based client that wishes to contact a particular RPC based server on a ost first contacts the host's portmapper server. The client passes the server the RPC service number of the server it wishes to access. If the server is running, the portmapper responds with the TCP or UDP port number it registered. The client is then able to communicate directly with the server. RPC servers are susceptible to many different attacks, for example:

- RPC based services are vulnerable to denial of service attacks, as the call to deregister a service is not well authenticated .
- The portmapper does not authenticate requests to provide information about servers and is therefore a security risk.
- If access to the portmapper is blocked by a screening router, an attacker can bypass the portmapper and simply probe each of the TCP and UDP ports until he receives a response indicating an RPC based server is active.
- The most serious problem with the portmapper though, is its ability to issue indirect calls . To avoid the overhead associated with the additional round trip necessary to identify a server's port number, the remote client can ask the portmapper to forward the call to the server. However the server is not able to tell that the request did not originate locally, and is therefore unable to assess the level of trust that should be afforded to the call.

Network Information Service - One of the most dangerous RPC applications according to [Ches94] is the Network Information Service (NIS). NIS is used to distribute a variety of important databases, including the password file, the host address table, and the public and private key databases used for secure RPC .

The Network File System - The Network File System (NFS) allows computers to mount file systems that are physically attached to other computers. There are some serious security problems associated with NFS, for example:

- NFS clients are allowed to read, change or delete files without having to log onto the server or enter a password.
- NFS has very weak client authentication.
- If not properly configured NFS can allow *any* other host to simply mount its file system.

Given the scope for abuse, it is not surprising that most of the literature advises extreme caution before allowing RPC based services through a firewall.

2.4.7 File Transfer Protocols

TFTP - The trivial file transfer protocol (TFTP) is a simple UDP based file transfer mechanism. TFTP is often used to boot diskless workstations. The protocol has no authentication. [Ches94] recommends that TFTP be disabled if it is not absolutely needed because its simplicity makes it very useful to attackers.

FTP – The file transfer protocol is one of the most widely and heavily used Internet applications. FTP can be used to transfer both ASCII and binary files. Separate channels are used for commands and data transfer. Anonymous FTP allows external users to retrieve files from a restricted area without prior arrangement or authorisation. By convention users log in with the userid "anonymous" uses this service. Some sites request that the user's electronic mail address be used as the password.

The FTP daemon runs with extremely high privilege levels. Reports.There have been several bugs in the daemon, which have opened disastrous security holes. Most organisations will require an anonymous FTP repository somewhere.

2.4.8 The "r" Commands

The "r" commands (rlogin, rsh and rexd) can be used for remote terminal access and remote execution of programs. These programs are used in a trusted environment to

allow users remote access without the need to re-authenticate themselves . The host that the user is connecting to trusts the host they are connecting from to have correctly authenticated the user.

The trusted host model is inappropriate for use on distrusted networks. Nevertheless the "r" commands are still widely used and have been involved in some of the higher profile attacks of recent times.

2.4.9 World Wide Web (WWW)

Most commentators agree that the World Wide Web is responsible for the explosive growth the Internet has seen in the last two years. To many organisations, the World Wide Web *is* the Internet. The commercial implications of the World Wide Web are staggering, and so, therefore, are the implications of inadequate security. Information based industries such as banking and insurance are likely to be transformed by Web developments. Sales, marketing and post-sales support in other industries will similarly be affected.

Developments are occurring more rapidly than the IETF committees can process them. For example version 1.0 of the HTTP is the de-facto Standard WWW protocol. However it has still to be sanctioned by the IETF, version 1.1 is now an Internet draft , and there are several HTTP enhancements and extensions at Internet draft stage .

Despite the fact that information sent to and from web browsers is currently visible to others, organisations are increasingly building links to corporate databases into their World Wide Web Pages.

The incidence of attacks on web sites is increasing rapidly and WWW security is a significant cause for concern as the HTML specification allows protocols other than HTTP to be used (i.e., FTP, TELNET, RLOGIN). HTML may therefore be used to bypass the filters normally applied to those protocols by a firewall. Using an HTTP proxy, which filters the relevant protocols as required , can rectify this. Other problems include:

- unexpected input values can cause actions which were not intended by the author;
- special characters may allow unauthorised access to the host;
- unexpectedly large input may cause a buffer overflow resulting in inappropriate actions;
- the potential for data driven attack especially for Trojan horses;

• Authentication/Confidentiality/Integrity Issues - Mutual authentication and protection from message stream modification (i.e., to support electronic commerce).

Potential Solutions:

• Type enforcement - In 1985, the US government published the Trusted Computer Security Evaluation Criteria, more commonly known as the Orange Book, which offers a range of ratings for secure systems. The ratings start with D, for systems with no security, and go to A1, which requires formal methods to verify security.

One of the key elements of the Orange Book was mandatory access control, where all the resources in the computer (such as users, files, services, and programs) are labelled with a security level, or sensitivity. The label identifies the degree of sensitivity of each resource, such as Unclassified, Confidential, Secret, and Top Secret. Labels effectively assign the data on the system to separate classifications.

Many initial Orange Book implementations were too restrictive. The DoD funded seven years of R&D to create a flexible implementation, resulting in the typeenforcement security model. Type enforcement is a security mechanism based on least privilege that controls how users, programs and data interact. Type enforcement works by grouping all the processes of the system into classes based on least privilege. Each process group is called a domain. In a similar manner, the files on the system are grouped into classes called types. The Domain Definition Table describes each domain's access rights for each type. The table cannot be changed while the system is running. Type enforcement also creates what is termed an "assured pipeline" to organise data flow between programs - to assure that information moves securely, type enforcement controls the data each program can read and write. Each program can only read from the stage in front of it, and write to the next stage, of the pipeline. No stage of the pipeline can be skipped.

• Digest Authentication - Digest based authentication uses a challenge-response paradigm in which the server issues a unique challenge string to the client. The client concatenates his password to this string and computes a one way hash of the result and transmits the result of this back to the server. The server also concatenates the user's password with the original message and generates a digest. If this digest matches the one returned by the user then the client has been authenticated without his password having been transmitted. This scheme requires the server to know the user's password;

however, the scheme proposed by has a nice refinement, which means that the server only needs to know the digest of the user's password.

• S-HTTP - The secure hypertext transfer protocol is compatible with HTTP but incorporates security extensions that support sender authentication, message integrity and confidentiality, and non-repudiation of origin.

• SSL - The Secure Sockets Layer developed at Netscape Communications Corporation takes a radically different approach to S-HTTP. Rather than enhancing World Wide Web security by extending the HTTP application protocol, SSL creates channel security between the application layer protocol, HTTP, and the transport layer protocol, TCP. SSL is backed by, amongst others, IBM, Microsoft and SpyGlass, all of whom are incorporating SSL in client server applications. Netscape have submitted SSL 3.0 in Internet Draft form to the Internet Engineering Task Force (IETF). SSL is far from flawless. Opponents have expressed concern about the weakness of the encryption borne out when an SSL 2.0 key was cracked in 1995 by two Berkeley graduate students.

• PCT - Microsoft's Private Communications Technology Protocol [Simo96] is a superset of Netscape's SSL protocol and is intended to address the perceived shortcomings of SSL. PCT will spawn a second key specifically for authentication. Microsoft also intends to develop a more robust random number generator, used to seed the encryption key, as this is also considered to be a weakness in SSL. PCT is intended to be backward compatible with SSL 2.0.

2.4.10 The X Window System

The X Window System was developed as part of Project Athena. Version 11 of the X Window System, commonly referred to as X11, was released in September 1987. With Release 2 of X11 in March 1988, control of X passed from MIT to the X consortium, an association of computer manufacturers who support the X protocol.

The X Window System is a network oriented windowing system. It uses the network for communicating I/O between the windowing display software and applications. An application need not be running on the same system that opens the display. The program that controls each display is the X server and the applications are the X clients[6].Applications that have connected to an X11 server can "do all sorts of things". For example the screen contents can be printed and key presses can be both detected and generated. This allows an attacker to read passwords as they are being

typed. The protection mechanisms, for example the so called "magic cookie" system, built into X11 are generally considered to be inadequate . Some research has suggested that it is feasible to allow X11 securely over the Internet using an application gateway, however most of the literature recommends extreme caution when considering this .

2.4.11 Internet Inter – ORB (IIOP)

The IIOP protocol was developed by The Object Management Group (OMG) to implement CORBA solutions over the World Wide Web. IIOP enables browsers and servers to exchange integers, arrays, and more complex objects, unlike HTTP, which only supports transmission of text.

IIOP is a specialisation of the abstract General Inter-ORB Protocol (GIOP). GIOP provides a standard set of message formats, data representations and connection management rules which are fully described in The Common Object Request Broker:

Architecture and Specification, version 2.0. IIOP essentially maps GIOP to work over TCP/IP and is meant to serve as a 'common language' which all CORBA 2.0 ORBs can use for intercommunication. ORB support for IIOP is a requirement for CORBA 2.0 compliance.

An object accessible via IIOP is identified by an interoperable object reference (IOR). Since the format of normal object reference is not prescribed by the OMG, the format of an IOR includes an ORB's internal object reference as well as an internet host address and a port number. An IOR is managed internally by the interoperating ORBs– it is not necessary for an application programmer to know the structure of an IOR.

CHAPTER THREE:COMPUTER SECURITY RISK and ATTACKS

3.1 Generic Risks

Three generic risks when connecting a trusted network with one that cannot be trusted.

3.1.1 Intrusion

Intrusion occurs when attacker gains access to the system and is able to use it and modify it in the same way as a legitimate user. In some cases rigorous password protection can protect against this type of attack, with accounts locking after three failed Access attempts etc. However policies need to be geared against social engineering attacks as well, where an attacker uses ploys such as posing as a senior manager and demanding an immediate password change to allow very important and urgent work to continue. Some attacks in this category will exploit weaknesses in operating system security and will not require the attacker to knock at the door; the door opens itself for them.

3.1.2 Industrial Espionage and Information theft

Industrial espionage is on the rise, there are currently 122 countries actively engaged in industrial and economic espionage to the benefit of their respective states. A study in 1992 sponsored by the American Society for Industrial Security (ASIS) found that proprietary business information theft had increased 260 percent since 1985 . The data indicated 30 percent of the reported losses in 1991 and 1992 had foreign involvement. he study also found that 58 percent of thefts were perpetrated by current or former employees. The three most damaging types of stolen information were pricing information, manufacturing process information, and product development and specification information. Other types of information stolen included customer lists, basic research, sales data, personnel data, compensation data, cost data, proposals, and strategic plans

Traditional warfare may even be giving way to "Information Warfare". The implications that the failure of the communications infrastructure would have for technology dependent Western society have prompted both Britain and the USA to develop formal Information Warfare Policies. Information warfare represents a global challenge that faces all late industrial and information age nation states. It also

represents the cheapest way for less developed nation states and religious or political movements to anonymously and grievously attack major nations and industrial corporations .

3.1.3 Denial of Service

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.)[1]

There is little that a network administrator can do to prevent denial of service attacks as an attacker can always attack upstream of the point of connection to the Internet and disrupt service. This is one of the reasons that people are wary of using the Internet for mission critical or time critical connectivity.

Malicious Code:

Malicious code can be thought of as an indirect denial of service attack. Most users are now familiar with the threat posed by viruses, worms, Trojan horses and genetic algorithms. However new forms of malicious code are appearing all the time. A new type of virus attacks documents rather than programs using the advanced features in desktop productivity tools such as word processors.

Currently the two high-risk areas for infection with malicious code are when downloading files or in binary attachments to mail messages. However new Technologies are being developed that extend World Wide Web viewers by downloading and executing software on the client rather than the server. Such programs are known as Java applets and greatly increase the risk of infection from malicious code.

3.2 Types of Attack

In a presentation titled "A Taxonomy of Internet Attacks - What you can expect to see" Marcus Ranum described eight types of attack from the Internet.

3.2.1 Social Engineering

An attack based on deceiving users or administrators at the target site. For example, telephoning users or operators pretending to be an authorised user, to attempt to gain illicit access to systems.

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.

People generally like being helpful and co-operative and attackers exploit this ruthlessly. Social Engineering is very hard to protect against as it is essentially hitting a "soft" target and requires "soft" means of addressing it such as staff education, clear policies and mechanisms for reporting problems.

3.2.2 Impersonation

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

Example 1 : 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.

Example 2 : the attacker writes a shell script to present a false login session to the user. The user enters his correct user-id and password, which the script records before initiating a real login session to allow the user to login. The user thinks he has entered his password incorrectly and is none the wiser.

Impersonation attacks are primarily sniffer and spoofing attacks , with attackers seeking to capture passwords. It is a mistake to dismiss attacks on passwords as being of little danger.

3.2.3 Exploits

These are attacks that seek to exploit a hole in a piece of software. Most of CERT's advisories fall into this category. For 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"

Exploits succeed because badly written software is the norm, security is generally added as afterthought, too many programs run with excessive privilege violating the least privilege principle, and few programs use the operating systems underlying security features .

3.2.4 Transitive Trust

Transitive trust attacks take advantage of the trust models used by remote services, such as the "r" commands discussed in chapter 2.4.8 (page 25).

Example 1 : 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. The attacker then compromises a user account on one of the remote computers to gain a foothold on an entirely new network. This is one of the attack strategies that the 1988 Internet "Worm" Virus used to propagate itself .

3.2.5 Data Driven

Data driven attacks take the form of Viruses and Trojan Horses. For example an attacker can email 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.

The World Wide Web is currently particularly vulnerable to data driven attacks. The emergence of languages such as Java that will run code on the client computer present attackers with significant new potential for this type of attack. A firewall can help to screen out some data driven attacks. Some firewall vendors are incorporating anti-virus software into their products, and some are able to control executable files. However firewalls in general provide little protection from data driven attacks.

3.2.6 Infrastructure

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

Example 1: 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.

3.2.7 Magic

These are attacks that nobody has thought of yet. Such attacks if and when discovered will be full of surprises. An illustrative (and possible) example is 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.

3.2.8 Combination attacks

Attackers are likely to use a combination of the above methods when seeking to gain unauthorised access or to deny service etc.

Example 1 : The attacker tells a new user who is using IRC (Internet Relay Chat) to obtain a utility program that will help them to use system better. This phase of the attack can be categorised as Social Engineering. The user downloads the program and runs it causing all his messages to be deleted, and exposing the password file to the attacker. This phase of the attack can be categorised as Data-Driven.

3.3 Types of Attackers

There are many ways to categorise attackers; we can't really do justice to the many variants of attackers we've seen over the years, and any quick summary of this

kind necessarily presents a rather stereotyped view. But the most important attackers are:

- Joyriders,
- Vandals,
- Score keepers,
- Spies.

3.4 Security Analysis Tools

There are several tools that will probe a computer to test for known vulnerabilities:

- TAMU, is a collection of very useful tools. Some can be used to build your own firewall, others can detect attack signatures. The Tiger scripts can be used to assess the security of your own machines.
- COPS , is another popular auditing package along the lines of Tiger scripts.
- Tripwire, is a package that evaluates a system and checks for altered files and the like.
- ISS, is a network vulnerability auditing package, along the lines of TAMU and our network sweep programs. It can be used to probe entire networks for vulnerabilities.
- SATAN, is another network vulnerability auditing package.
- SPI, combines the functionality of programs such as COPS and tripwire. It also attempts to track important security patches on a per-platform base.

CHAPTER FOUR:NETWORK SECURITY POLICY

Much of the literature on firewalls concentrates on diagramming the numerous possible configurations of routers, host systems, interfaces, and sub-nets. It is imperative, however, not to lose sight of the broad definition of a firewall as a part of security policy .

The role of a security policy is to ensure that each of the four fundamental components that make up computer security, Authentication, Access Control, Integrity and Confidentiality are adequately addressed.

Typical questions that need to be answered when developing a network security policy are:

- What resources are we trying to protect?
- Which people do we need to protect the resources from?
- How likely are the threats?
- How important is the resource?
- What measures can be implemented to protect the resource?
- How cost effectively and in what time frame can these be implemented?
- Who authorises users?

These questions should be revisited periodically, as network security is very dynamic.

The security policy identifies the threats that need to be protected against and defines the level of protection required. The security policy will itself contain several different policies, for example a Network Service Access Policy and System Specific Policies and will be based on a security strategy.

4.1 Security Strategies

4.1.1 Least Privilege

The principle of least privilege is to grant only those privileges that are required. Systems that allow permission to be granted or revoked by operation providing finegrain control are well suited to this. There is generally an overhead in terms of increased system maintenance.

Adopting a least privilege strategy limits exposure to attacks and limits the damage that can be done when an attack is successful. Many of the common security

problems on the Internet can be viewed as failures to follow the principle of least privilege.

4.1.2 Defence in Depth

The defence in depth strategy is summed up by the term "Belt and Braces", i.e. use as many security mechanisms as you can and arrange them so that they back each other up. One of the problems with firewall systems is that they provide an all or nothing solution to security. If the firewall is breached (and this has happened) the internal network is a soft target. This was noted by Bellovin who coined the term "Hard on the outside, soft and chewy on the inside" to describe it. Some firewalls however do implement the principle of defence in depth using techniques such as Type Enforcement.

An important aspect of defence in depth that is often overlooked is the need to avoid common mode failures. For example if an attacker can exploit a security weakness in brand X's router then there is little point in having two of them. However brand Y's router may not have the same weakness and therefore the principle of defence in depth is met.

This principle is important in the context of firewalls as many of the products commercially available are variations of Trusted Information Systems Gauntlet or their tool kit.

4.1.3 Choke Point

A choke point is a single point through which all-incoming and outgoing network traffic is funneled. As all traffic passes through a choke point it is the natural place to focus monitoring and control efforts such as Internet firewalls. It is also the natural place at which to break the connection with the external network if necessary.

Choke points are often criticised as an all eggs in one basket solution. This concern can be addressed by building some redundancy into the choke point. The key point is that the choke point provides control.

The largest threat to a choke point strategy is if an attacker is able to bypass the choke point. As Firewalls generally act as choke points this is a significant issue, especially given the ease with which SLIP or PPP connections to Internet Service providers can be established. As choke points can experience high levels of network traffic it is important to ensure that there is sufficient bandwidth available at the choke

point to prevent a network traffic bottleneck. Any monitoring and logging software should also be able to cope with the level of network traffic.

4.1.4 Fail Safe Stance

If a system is going to fail, it should be designed to fail into a safe state. This principle is particularly important in the design of Internet firewalls. Packet filters and application level gateways, both of which are discussed in the next chapter, should fail in such a way that traffic to and from the Internet is stopped.

4.1.5 Security through Obscurity

This strategy is based on the hope that if you keep a low profile, would be attackers will not find you, and if they do, they will pass you by. Many companies do not publish the telephone numbers of their dial-in modems, only divulging the numbers on a need to know basis. While this is a sensible precaution it is a poor basis for long term security.

Information tends to leak out, and attackers are often skilled at eliciting information from staff using social engineering techniques.

Many organisations assume that an attacker won't be interested in them, and that they are therefore unlikely to be the target of an attack. The rationale behind this stance assumes that a site is targeted because the attacker is interested in the information stored on it.

Attacks generally involve several computers and a multitude of accounts. An attacker may capture accounts and gain unauthorised access to several systems before reaching his real target. A site can be compromised for no other reason than to provide a staging post for attacks on other sites, and to the attacker, it means little more than another IP address.

4.1.6 Simplicity

Software is complex. As the size of a piece of software grows it becomes increasingly difficult to test all eventualities. Complex code will probably have unknown loopholes that an attacker can exploit. These loopholes may be convoluted but that will not prevent an attacker from trying to exploit them, some of the exploit attacks against sendmail have been extremely intricate.

Simplicity is an important factor in sound network defences. Application level

gateway network Security systems should have all extraneous functionality removed and should be kept as small and simple as possible.

4.2 Host Based Security

Host based security is probably the most common computer security model in current use. The major problem with the host based security model is that it does not scale well. The major impediment to effective host security in modern computing environments is the complexity and diversity of those environments . Even if all hosts are identical, the sheer number of them at some sites makes securing each of them difficult. Effectively implementing and maintaining host security takes a significant amount of time and effort, and is a complex task. While the host security model might be appropriate for small sites, and while all sites should implement some level of host security, it is not cost effective for larger sites, requiring too many restrictions, and too many people .

4.3 Network Based Security

Network security is designed to address the problems identified with host security. The network security model concentrates on controlling network access to hosts and services rather than on securing the hosts themselves.

Network security approaches include building firewalls to protect trusted Networks from distrusted networks, utilising strong authentication techniques, and using encryption to protect the confidentiality and integrity of data as it passed across the network.

4.4 Security Policy

A general security policy contains the security requirements of the enterprise. Every new security system should translate a security policy into reality, depending on the existing requirements.

Based on this policy, one can determine the security requirements of an entire network and define a network security policy. This can be divided further into subnetwork policies. The last step in this hierarchy will be the policy of the active security components in our network, i.e. the firewall security policy, which contains the security requirements of the firewall system. This mapped on rules, which are needed to guarantee the security of the protected network.

A security policy can be defined to include the following three aspects:

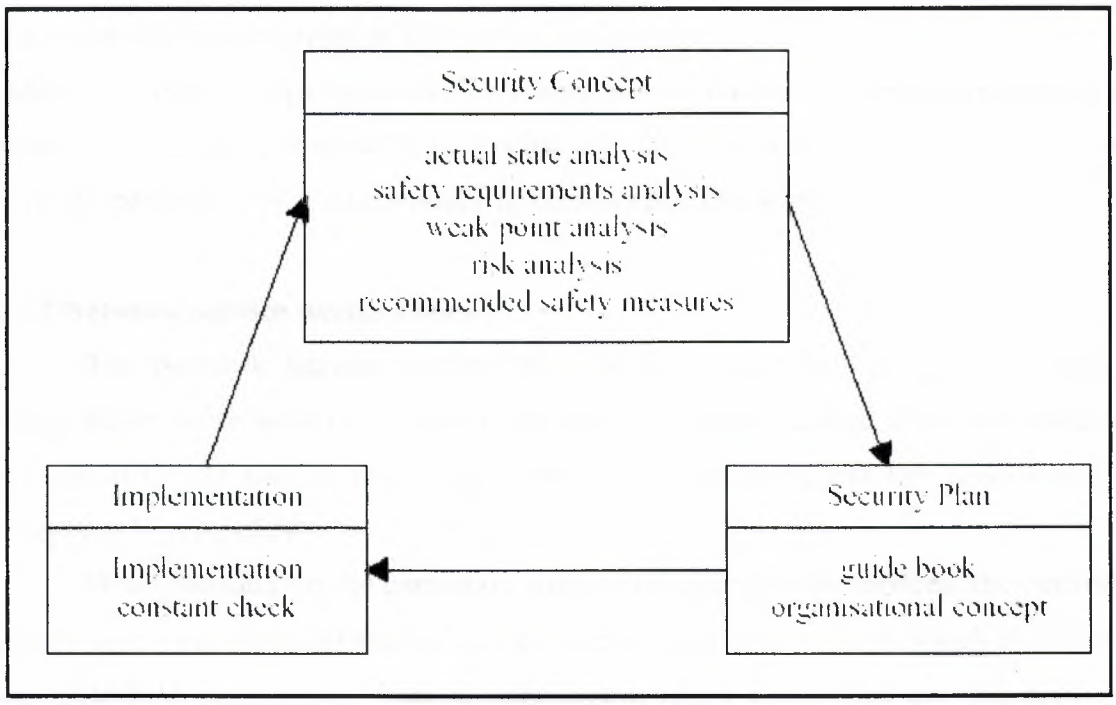

Figure 4.1 Security Policy Cycle

- the concept, based on the actual state analysis, the safety requirements analysis, the weak point analysis, the risk analysis and the recommended safety measures.
- the security plan, consisting of the guide book and the organisational concept.
- the implementation, also containing a constant check (supervision).

The cycle shown in Figure 4.1 shows the dynamic characteristics of such a security policy.

4.4.1 Site Security Policy

The Site Security Policy is an overall policy regarding the protection of the organisation's information resources. This includes everything from document shredders to virus scanners, and remote access to floppy disk tracking. At the highest level, the overall organisational policy might state:

- Information is vital to the economic well being of the organisation.

• Every cost-effective effort will be made to ensure the confidentiality, integrity, authenticity, availability and utility of information.

• Protecting the confidentiality, integrity, and availability of information resources is a priority for all employees at all levels of the company.

Below this come site-specific policies covering physical access to the property, general access to information systems, and specific access to services on those systems. The firewall's network service access policy is formulated at this level.

4.4.2 Network Service Access Policy

The Network Service Access Policy is a higher-level, issue-specific policy which defines those services that will be allowed or explicitly denied from the restricted network, plus the way in which these services will be used, and the conditions for exceptions to this policy.

While focusing on the restriction and use of internetwork services, the network service access policy should also include all other outside network access such as dial-in and SLIP/PPP connections. This is important because restrictions on one network service access can lead users to try others. For example, if restricting access to the Internet via a gateway prevents Web browsing, users are likely to create dial-up PPP connections in order to obtain this service. Since these are non-sanctioned, ad hoc connections, they are likely to be improperly secured while at the same time opening the network to attack.

For a firewall to be successful, the network service access policy should be drafted before the firewall is implemented. The policy must be realistic and sound. A realistic policy is one that provides a balance between protecting the network from known risks while still providing users reasonable access to network resources. If a firewall system denies or restricts services, it usually requires the strength of the network service Access policy to prevent the firewall's access controls from being modified or circumvented on an ad hoc basis. Only a sound, management backed policy can provide this defence against internal resistance. Here are the typical network service access policies that a firewall implements:

• Allow no access to a site from the Internet, but allow access from the site to the Internet; or, in contrast,

• Allow some access from the Internet, but only to selected systems such as information servers and e-mail servers.

Firewalls often implement network service access policies that allow some users access from the Internet to selected internal hosts, but this access would be granted only if necessary and only if it could be combined with advanced authentication.

4.4.3 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.

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 nonstandard 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.

4.4.4 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 organisation are authorised to run a particular program. Access controls are used by the system to implement (or enforce) this policy

4.4.5 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.

4.4.6 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:

1. Identify the mission or business critical functions.
2. Identify the resources that support the critical functions.
3. Anticipate potential contingencies or disasters.

4. Select contingency planning strategies.
5. Implement the contingency strategies.
6. Test and revise the strategy.

CHAPTER FIVE:FIREWALL THEORY ANDCOMPONENTS

Once a security strategy and network security policy have been decided means of implementation are required. The generic term "Firewall" is increasingly being used to describe the combination of hardware, software and management activities that are used to effect the policy.

5.1 What is an Internet Firewall ?

Internet firewalls are 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 as an approach to security. He uses the term Firewall to mean the strategies and policies. He uses the term Firewall System to refer to the hardware and software elements that implement the policy.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 .
- Only authorised traffic, as defined by the security policy, is allowed to pass through it.
- The system itself is immune to penetration

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

Typically, the two networks in question are an organisation's internal network (trusted) and the Internet (distrusted). 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. There are good reasons, discussed them in chapter 5.2, for using firewalls when connecting any trusted network with a less trusted network internal or external.

5.2 What can a firewall do ?

A firewall can enforce security policy. The firewall is the means by which the network access security policy is implemented. Internet services considered to be insecure can be restricted and access to or from certain hosts can be restricted.

A firewall can *log activity effectively*. A firewall can *limit your exposure to the distrusted network by controlling/restricting access to/from it to the level defined in the security policy*. This includes controlling what users use the Internet for.

A firewall can *be a focus for security decisions - a choke point*. All traffic to or from the Internet must pass through it. By focusing defences on this point they can reduce internal system security overhead since they allow an organisation to concentrate security efforts on a limited number of machines.

5.3 What can't a firewall do ?

While 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 introduceInternet 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.

5.4 Firewall Components

The main firewall components are as follows:

a) Components, which actively interfere with the communication between the trusted and the distrusted network. There are two types of active components:

- packet filters,
- application gateways.

b) Security Management Component, which is needed for the administration of active firewall components.

The most important objectives of every firewall systems are:

- access control at different levels (network level, user level),
- control at the application layer,
- user rights administration,
- isolation of certain services,
- proof backup and analysis of the log,
- alarm facilities,
- concealment of the structure of the internal network,
- structuring networks, defining security domains with different security necessities,
- confidentiality,
- resistance of the firewall against attacks. There must be possibilities of accounting and network address translations (NAT).

The desired architecture of an active firewall component can be inferred from the goals of the firewall systems. A good proposal for an active firewall component architecture is given by N. Pohlmann[Pohl97]. Such an architecture should have a modular structure. The following figure shows this architecture:

In practise, unfortunately not all the goals can be fulfilled by the available products. We need a Criteria to select a firewall product. The most important selection criteria are:

- the security of the firewall platform,
- the simplicity of administration,
- the transparency of usage of Internet services to increase the acceptance,
- the address translation,
- the server security,
- the authentication for external users,
- the usage of encrypting systems,
- the log and alarm facilities.

We need to understand how the active firewall components work and what advantages or disadvantages they have for selecting the right firewall product.

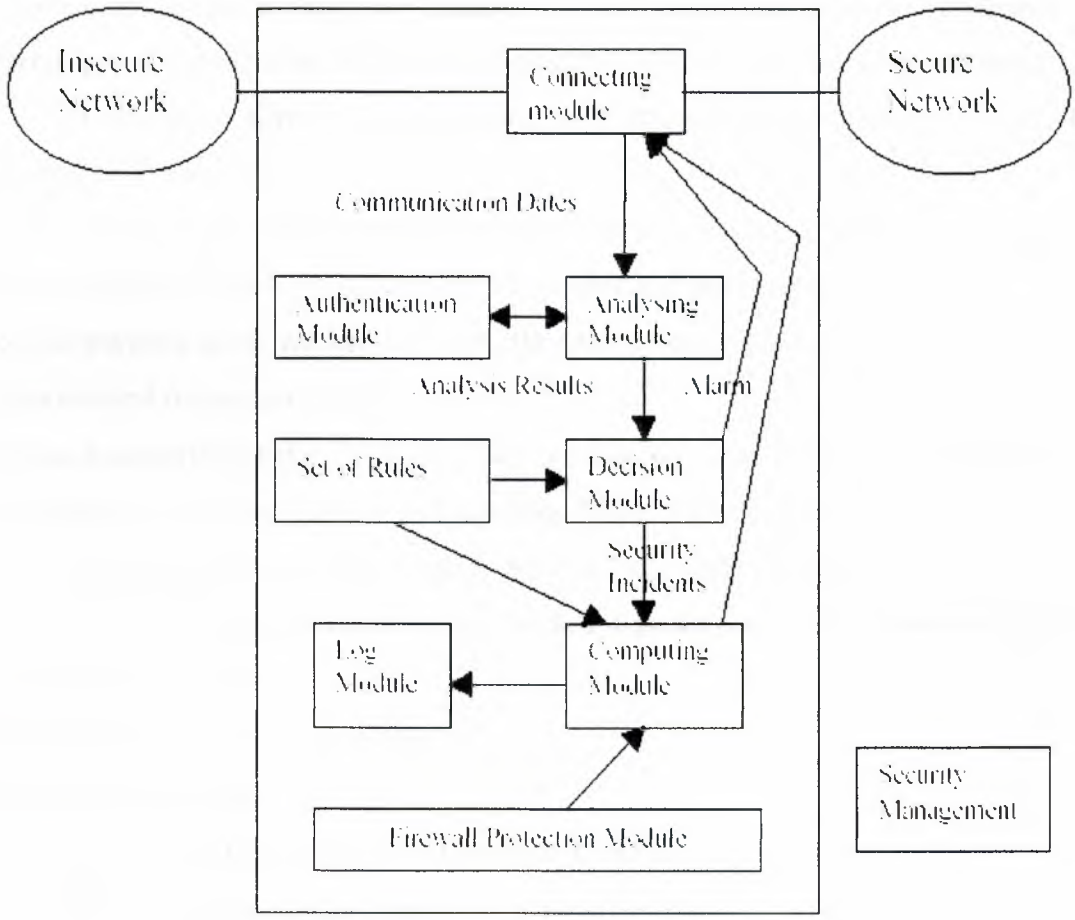

Figure 5.1 Active Firewall Component Accourding to Pohlman

5.4.1 Packet Filter

A *Packet Filter* is a network security mechanism that works by controlling what data can flow to and from a network and represents the simplest category of firewall components. It is implemented on the network layer and the transport layer in the TCP/IP protocol. The type of router used in a packet filtering firewall is known as a *screening router*. Packet Filtering lets you control (allow or disallow) data transfer based on:

- the address the data is (supposedly) coming from,
- the address the data is going to,
- the session and application protocols being used to transfer the data.

We can distinguish between packet filters and circuit relays. Circuit relays are a special form of packet filters. They consist of defining complete relations of services for

protecting them as who can use them and to whom, what parameters you must use... The main advantage is that every attack has few possibilities of success, because of internal limitations. And the big disadvantage is that they can influence the efficiency.

Furthermore, Circuit Relays cannot with attacks against applications, i.e., IP-Spoofing or Flooding.

For the packet analysis the following information should be used:

- it has to be proofed on which interface the packet had been received,
- on the network layer, we have to check the used protocol type (i.e., IP, ICMP, ...) and the source and destination address.
- on the transport layer (i.e., TCP or UDP), we have to check the source and destination port numbers, which can identify services (like Telnet, FTP, HTTP, ...).

This kind of filtering is sometimes called "Service Dependant Filtering",

- an additional information could be the check if a packet has been transmitted during a defined period of time.

Packet Filters have many advantages:

- they are transparent,
- it is simple to include extensions for new protocols and new services,
- they have low complexity and therefore high performance,
- one screening router can help protect an entire network,
- packet filtering doesn't require user knowledge or co-operation,
- packet filtering is widely available in many routers.

Although packet filtering provides many advantages, there are some disadvantages to using packet filtering as well:

- current filtering tools are not perfect,
- some protocols are not well suited to packet filtering,
- some policies can't readily be enforced by normal packet filtering routers.

But the main disadvantage is that data above the transport layer is not analysed, consequently there is no security for the application layer. The structure of the secured network cannot be concealed. Logging of communication is possible only up to layer 4.

You can set up packet filters incorporating them in routers (screening router) or implementing them as a dedicated firewall component. Router with packet filter facilities are cheap firewall components, but they have some disadvantages:

• there is little or no logging capability. It is often therefore difficult for an administrator to determine whether the router has been compromised or is under attack;

• packet filtering rules are difficult to test thoroughly, which may leave a site open to untested vulnerabilities;

• if complex filtering rules are required, the filtering rules may become unmanageable, and

• each host directly accessible from the Internet will require its own copy of advanced authentication measures.

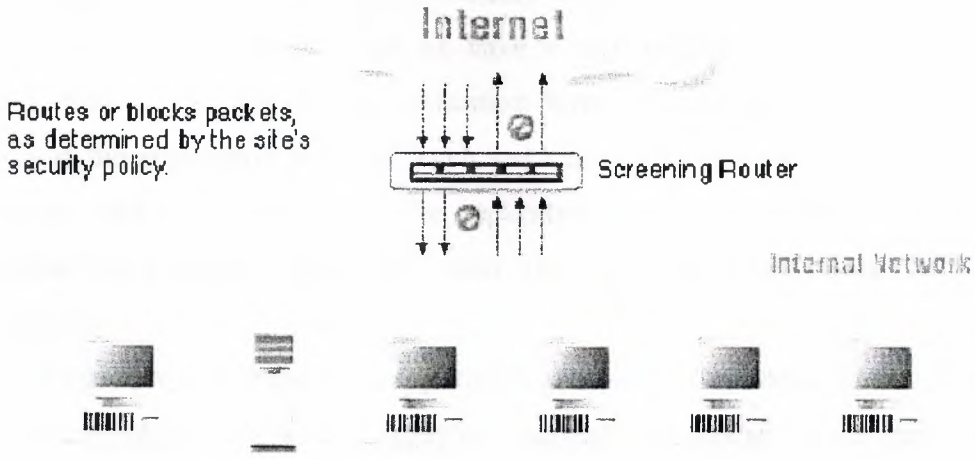

Figure 5.2 Packet Fitler(screening Router)

Packet filters as dedicated firewall components have advantages as:

• fulfil the design criteria for active firewall elements,

• delimitation between communication requirements and security requirements,

• can be managed by a security component.

But there are disadvantages too:

• the firewall component is more expensive than software updates in routers,

• decrease the availability of services to our network.

Packet filters are configured with an access control list. This list tells the firewall what IP address can be in a source and what can be in a destination. By filtering on a source and destination pair it is possible to limit specific clients to communicating with specific

servers. Since they can also look at TCP and UDP port numbers, we can restrict access down to the individual application running on a server.

Usually, packet filters are used as extensions for security solutions. They can be used for smaller networks with fewer security requirements. As long as we do not allow any access from the Internet to the Intranet, this type of firewall component offers basic protection. The definition of any user that can access the Intranet from the Internet, can result in an not acceptable firewall component. A better solution will be given in the next point by the application gateway concept.

5.4.2 Application Gateway

Application gateways are specialised application or server programs that run on a firewall host. These programs provide a safety barrier between the internal user and the Internet. They are necessary when we have a high security demand. An application gateway firewall operates at the application layer and can therefore provide access controls at the application protocol level[1] and can handle store and forward as well as interactive traffic. By operating at the application layer it is possible to close up the connection very tightly, and only open the connection up under well defined circumstances.

We use proxies to create an application gateway. Proxy services provide the user with intermittent services from a proxy of a server. The real server remains concealed from the user. A proxy is a software component that is responsible for a certain service. When we prohibit a certain service, we have to stop the appendant proxy on the application gateway.

There are two kinds of proxies: application level proxies and generic proxies, also called circuit level proxies. Application level proxies understand the application protocol and are therefore able to control the session based on the operations being requested.

The custom application acts as a "proxy" between the client and the server[2]. Because all data between the client and the server is routed through the application proxy it is able to control both, 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.

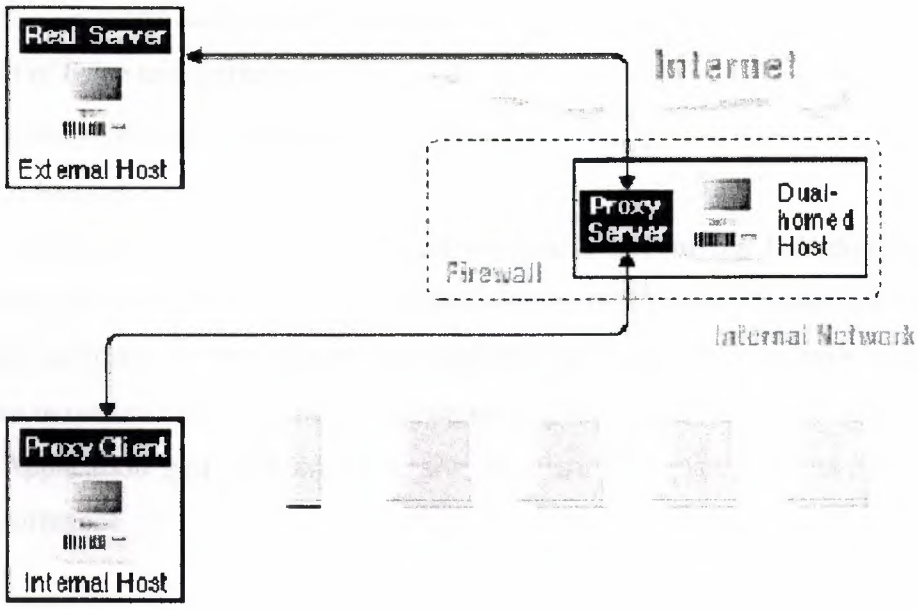

Figure 5.3 Application Level Gateaway(Dual-Homed Host)

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. This reintroduces 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.

Application gateways offer many advantages:

- the decoupling of services by proxies provide us with a high degree of security,
- services can be handled very simply, they can be switched on or off,
- they can offer additional security services, like encryption facilities, or other facilities which allow us to react to upcoming security breaches,
- simple accounting facilities,
- concealment of the internal network structure, using network address translation,

- challenge response authorisation possible,
- control of behaviour patterns,
- allow users to access Internet services `directly´,
- good at logging.

The flexibility of application level proxies is one of the few disadvantages of application gateways, as every new service needs a new proxy. Generic proxies do not have this problem. In the case of user oriented services we first have to identify or authorise ourselves before we have a transparent connection.

Application gateways are expensive in relation to packet filters and work with less performance.

5.4.3 Security Management Component

The security management component defines the rules for the active firewall components and evaluates the relevant logged security data. The computer on which the security management component is running must be resistant against attacks. A security breach could result in deactivation of the active firewall component via the attacked security management component.

The security management should offer at least the following security mechanisms:

- identification and authentication,
- auditing facilities,
- encryption of relevant data,
- task management and distribution.

The integration of our firewall security policy in our network security policy should be fulfilled by integrating the security management in a network management system.

CHAPTER 6:FIREWALL ARCHITECTURES

The packet filtering technologies that are used in screening routers provide an efficient and general way to control network traffic. They have the advantage that no changes are required to host or client applications because they operate at the transport and network layers. Application level gateways extend control of network traffic to the application layer, and have the advantage that because they can understand the application protocol they can implement a finer degree of control and provide detailed logs.

Firewalls bring these components together to provide extremely effective network based security control. To illustrate this, several "standard" Internet firewall architectures or configurations are presented.

6.1 Dual-Homed Host Architecture

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 6.1. 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.

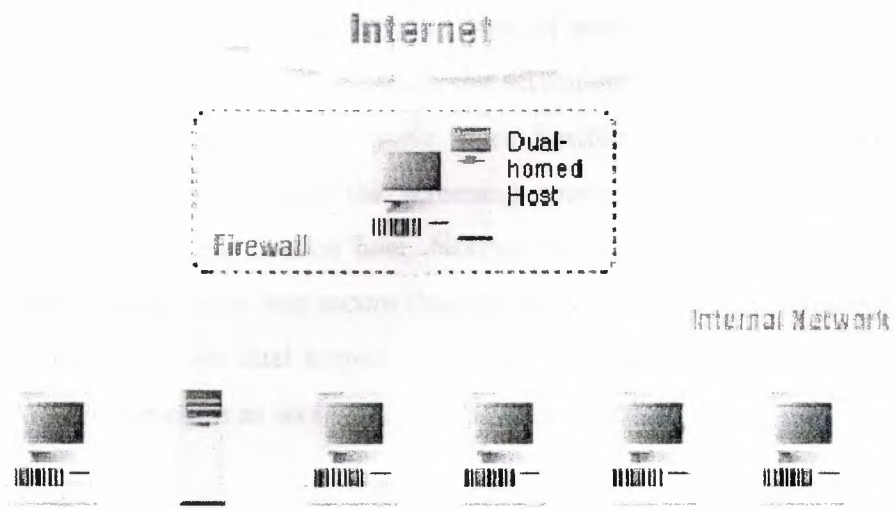

Figure 6.1 Dual-homed host architecture

Because the routing function is disabled the host isolates the two networks from each other while 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 distrusted services.

6.2 Screened Host Architecture

In this architecture, illustrated in figure 6.2, the primary security is provided by packet filtering and a bastion host sits on the internal network providing the required application.

The screening router's packet filtering rules are configured such that the bastion host is the only host accessible from the Internet. Connections to the Internet may be routed through an application proxy on the bastion host, or in some cases, allowed directly through the screening router, depending on the network security policy. Screened host architecture adds an additional layer of security to the dual homed host architecture, as an attacker has to first bypass the screening router, and then the bastion host. And screened host architecture, while more flexible than the dual homed host architecture, is less secure because the screening router is allowed to pass certain "trusted" services around the bastion host. Noting that the screened host architecture may, at first sight, appear to be less secure than the dual homed host architecture, states that this is misleading as the dual homed host itself may fail in some unexpected way, and that the two are therefore as secure as each other in practice.

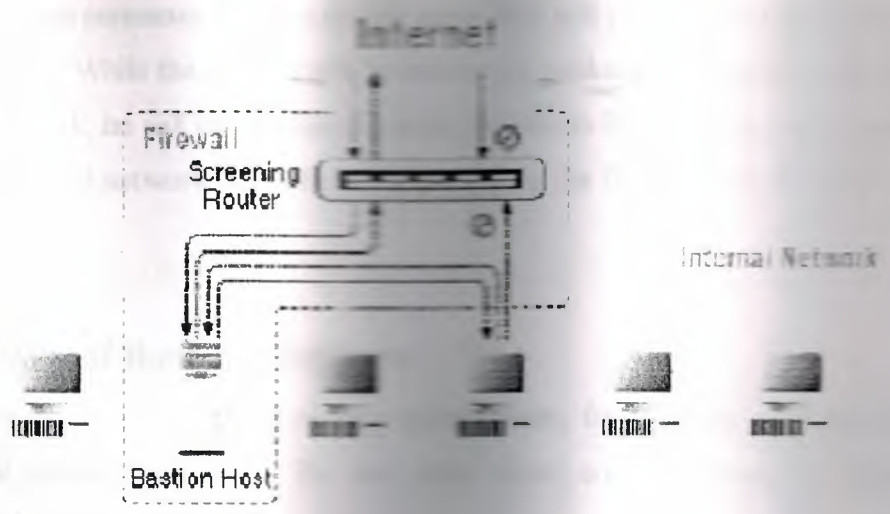

Figure 6.2 Screened Host architecture

6.3 Screened Subnet Architecture

With both the dual homed host and screened host architectures, the trusted network is vulnerable if the bastion host is compromised. The impact of the bastion host being compromised can be reduced by isolating it on a perimeter network (DMZ). The simplest way to provide a perimeter network is to add an additional screening router to the screened host architecture. This architecture, illustrated in figure 6.3, is called the screened subnet architecture. The bastion host is then located on the perimeter network between the two screening routers.

Figure 6.3 Screening Subnet Architecture

An attacker that successfully compromises the bastion host now will only be able to access the perimeter net. The trusted network is still protected by the internal screening router. While the attacker will be able to use packet sniffer software on the perimeter network, he will not be able to collect passwords for, or to examine sensitive files on, the trusted network unless these are passed via the DMZ, which is itself a security weakness.

6.4 Variations of these architectures

There are many variations of these architectures, for example providing internal and external demilitarised zones. The next table shows some common variations, and which are either corrects or dangerous:

Table 6.1 Variation of These Architectures

OK	Dangerous
• Multiple Bastion Hosts. For example: to separate external and internal services. • Multiple exterior Routers. • Multiple perimeter Networks. • To use Dual-Homed Host and Screened Subnets. • To merge the interior Router and the exterior Router.	• To merge the Bastion Host and the interior Routers. • Multiple interior Routers in the perimeter Network.

CHAPTER 7:ATTACS ON FIREWALL COMPONENTS

7.1 Types of Attacks on Firewall components

The main categories which endanger our systems are: the variety of network services, the loss of confidence and integrity, concept errors and the misuse of free available information. The most important attacks are caused by concept errors like:

7.1.1 IP Spoofing

Simulating trustworthy addresses we can reach the internal network without problems. Simple packet filters no protection against this problem. Usually application gateways either provide no protection.

7.1.2 ICMP attack

Using "redirect" packets to change routing tables. An other possibility is to perform an "denial of service" attack, falsifying "destination unreachable" or "time to live exceed" packets. ICMP packets can be filtered by packet filters.

7.1.3 Internet routing attack

This means that the source of the packet supplies the route it should follow to the destination. There's two flavors, loose and strict. Strict source routing supplies the list of routers the packet should follow. Loose source routing does the same, but it's ok for the packet to go through others in between the ones listed in the source route in the packet. Evaluating source routing information, attacker can learn something about the internal network. We can defend this attack by using only static routing. Dynamic routing should be turned off.

7.1.4 TCP SYN Flooding

TCP SYN Flooding consists of a tool that only implements one portion of the Sequence Number Guessing attack, with a completely different focus. TCP SYN Flooding causes servers to quit responding to requests to open new connections with clients – a denial of service attack. Denial of service attacks prevent people from using the affected system or networks. These attacks usually proceed by overloading the target in some fashion. For example, simply sending large ping packets can "fill up" a

site's connection to the Internet. Flooding is still a danger for application gateway and packet filters.

7.1.5 Snooping or Sniffing

Sniffing is a passive attack, the attacker observes network traffic but does not disturb it.

7.1.6 IP Splicing/hijacking

An attack whereby an active, established, session is intercepted and co-opted by the attacker. IP Splicing attacks may occur after an authentication has been made, permitting the attacker to assume the role of an already authorised user. Primary protections against IP Splicing rely on encryption at the session or network layer.

7.1.7 Trojan horse

A software entity that appears to do something normal but which, in fact, contains a trapdoor or attack program.

7.1.8 Data driven attack

A form of attack in which the attack is encoded in innocuous seeming data which is executed by a user or other software to implement an attack is a concern since it may get through the firewall in data form and launch an attack against a system behind the firewall.

7.1.9 Virus

A program that "infects" computer files (usually other executable programs) by inserting in those files copies of itself. This is usually done in such a manner that copies will be executed when the file is loaded into memory, allowing them to infect still other files, and so on. Viruses often have damaging side effects, sometimes intentionally, sometimes not. A virus that propagates itself across computer networks is sometimes referred to as a "Worm", especially if it is composed of many separate segments distributed across the network.

CHAPTER 8:CONFIGURING SOME INTERNET SERVICES

This chapter describes some Internet Services: how they work, what their packet filtering and proxying characteristics are, what their security implications are with respect to firewalls, and how to make them work with a firewall.

8.1 File Transfer Protocol (FTP)

There are two types of FTP access: user FTP and anonymous FTP. The first requires an account on the server and the second is for people who do not have an account.

FTP uses two separate TCP connections: one to carry commands and results, called the *command channel* and the other to carry any actual files and directory listing transferred, the *data channel*. On the server, the command channel uses the port 21, and the data channel uses the port 20. The client uses ports above 1023 for both. To start an FTP session, a client first allocates two TCP ports, each of them with a port number above 1024. One opens the command channel and then issues FTP's PORT command to tell the server the number of the second port for the data channel. This data channel connection open complicates things for sites that are attempting to do "start of connection" packet filtering to ensure that all TCP connections are initiated from the inside, because external FTP servers will attempt to initiate data connections to internal clients, in response to command connections opened from those internal clients.

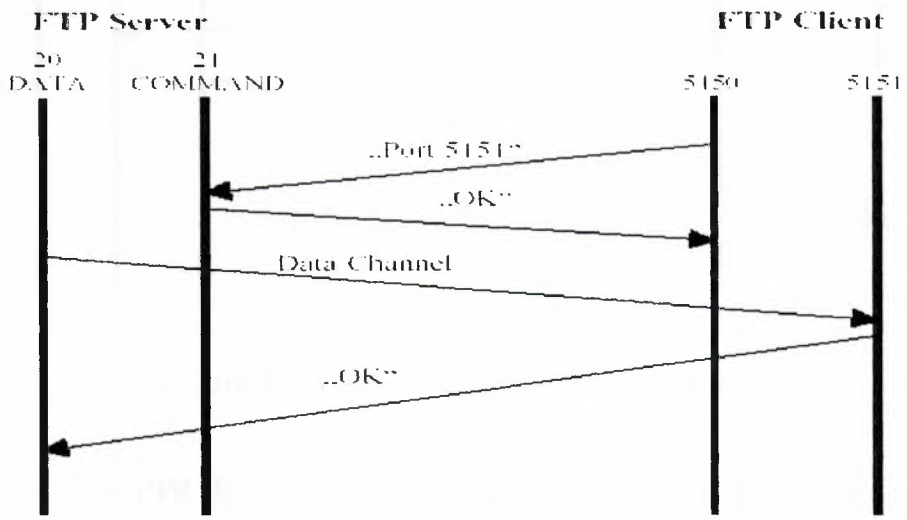

Figure 8.1 A Normal mode connection

Most FTP servers and many FTP clients support an alternative mode that allows the client to open both the command and the data channels to the server. This mode is called "passive mode" or "PASV mode". You can avoid start of connection filtering problems, because all connections will be opened from the inside, by the clients.

To use passive mode, an FTP client allocates two TCP ports, the first port to contact the FTP server. Instead of issuing the PORT command to tell the server the client's second port, the client issues the PASV command. This causes the server to allocate a second port of its own for the data channel (for architectural reasons, servers use random ports above 1023 for this, not port 20 as in normal mode; you couldn't have two servers on the same machine simultaneously listening for incoming PASV-mode data connections on port 20), and tell the client the number of that port.

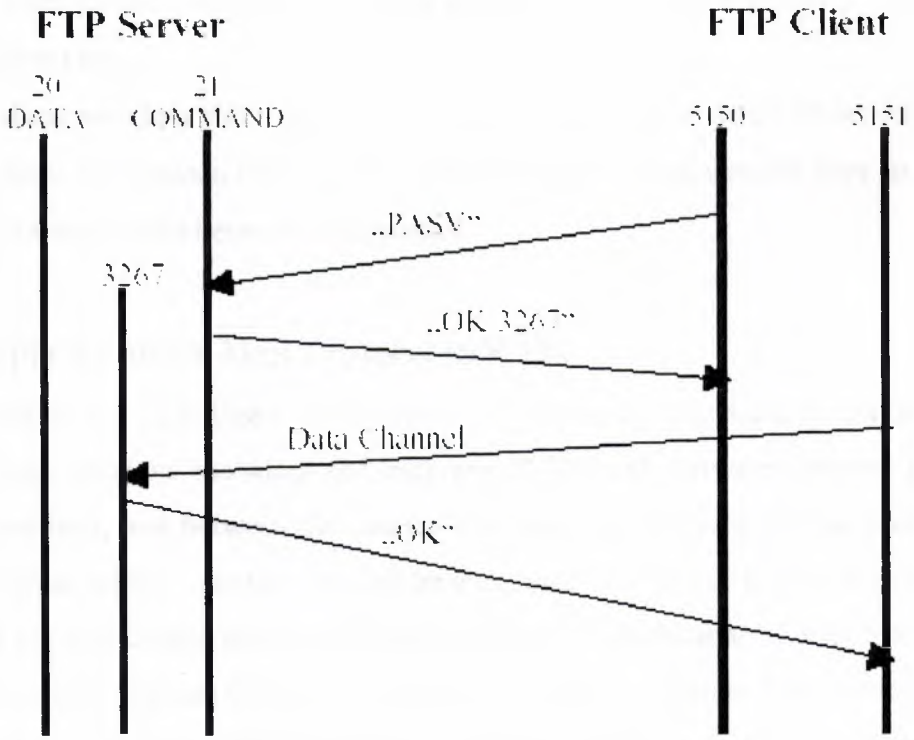

Figure 8.2 A Passive Mode FTP Connection

If you have FTP clients that properly support passive mode, then allow internal hosts to contact external FTP servers via packet filtering. This is only safe if you can

filter on the TCP ACK bit, so that you can allow only outgoing TCP connections from ports above 1023 to ports above 1023.

If you have FTP clients that do not support passive mode, then use an FTP Proxy servers as the one in the TIS Internet Firewall Toolkit. Because, if you want to allow FTP via packet filtering, you will have to put a special case exception in your packet filtering rules to allow the server to open the data channel back in to the client. If you do so, you will still be vulnerable to attackers launching a connection from port 20 on the attacker's end to a port above 1023 on your end. Therefore, you should restrict this special case exception as much as possible, i.e., by trying it to the address of the particular client a server that doesn't support passive mode. (Even an exception for a single server makes you vulnerable to forged connections from that server) Recommendations

- Consider providing FTP access via both packet filtering and proxies, supporting passive mode via packet filtering and normal mode via proxies.
- If you want to allow incoming FTP, use packet filters to allow incoming FTP only to your Bastion Host.
- If you allow incoming FTP (anonymous or user), use an up to date FTP server.
- If you allow anonymous FTP users to write files, protect the writable area so it cannot be used to transfer files between third parties.

8.2 Simple Transfer Mail Protocol (SMTP)

SMTP is a TCP based service (port 25). Normally, you want to configure your packet filters to allow incoming and outgoing SMTP only between external hosts and the bastion host, and between the bastion host and your internal mail servers. Do not allow external hosts to contact random internal hosts via SMTP. If you cannot filter on the ACK bit, you cannot safely allow outgoing SMTP connections directly from random internal hosts. If you can filter on the ACK bit, you allow internal hosts to send mail to external hosts, but there isn't advantage in doing so. Although it shouldn't increase your vulnerability, it increase the likelihood that you're going to send misformatted mail, because the mail (mis)configurations of all your machines would be visible to the external world, and the chances that all your internal machines do all the right things with mail headers (particularly in adding fully qualified domain names to addresses and "Message-ID:" lines) are low. Sending outgoing mail via the bastion host allows the

bastion host theopportunity to clean up the headers before the mail is loosed upon the world.

For configuring your mail system to work with a firewall, here are the important steps to follow:

- Use DNS Mail Exchange (MX) records to specify that all your incoming mail should be directed to your bastion host(s).
- Configure the mailer on the bastion host to check the destination address on mail it receives.
- Configure your internal systems to send all outgoing mail to the bastion host.

Figure 8.3 (outbound SMTP) and Figure 8.4 (inbound SMTP) show how packet filtering works with SMTP:

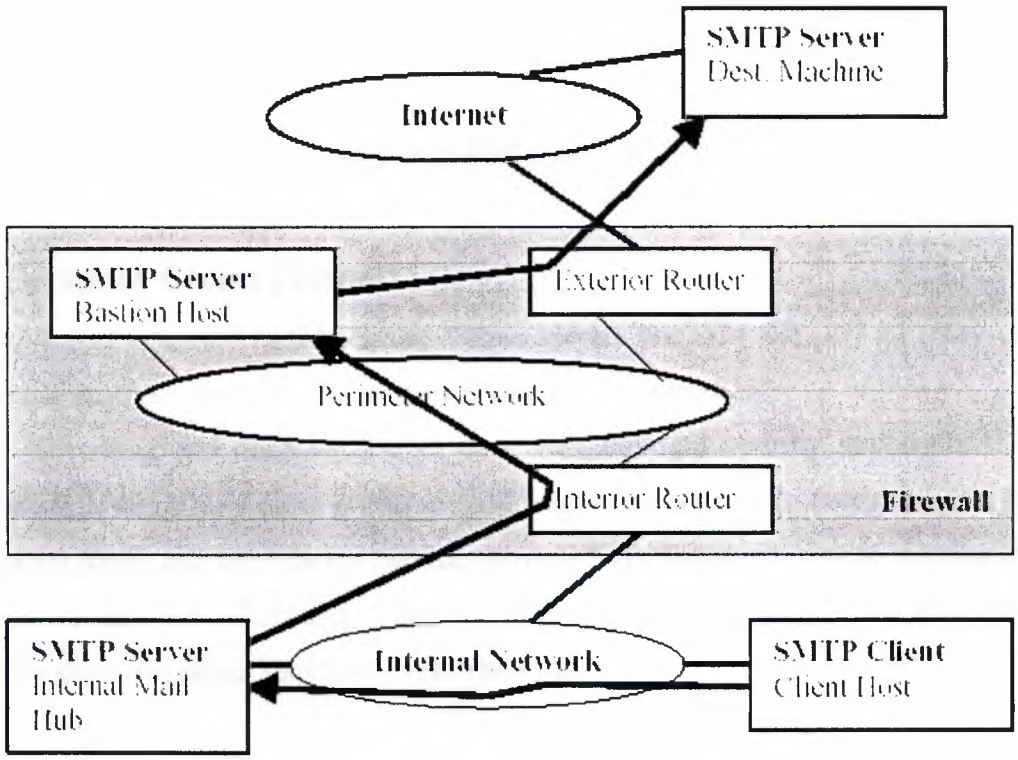

Figure 8.3 Outbound SMTP

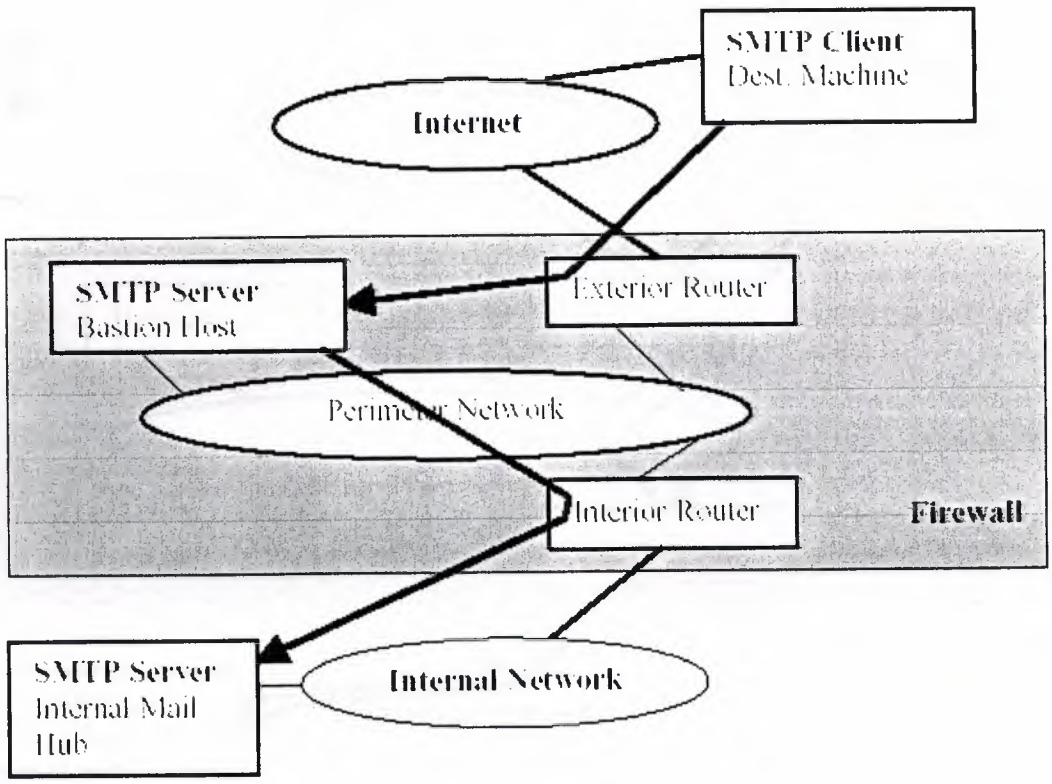

Figure 8.4 Inbound SMTP

8.3 Terminal Access (Telnet)

Telnet is a TCP based service. Telnet servers normally use port 23 (they can be set to use any port number). Telnet clients use ports above 1023.

Incoming and outgoing Telnet have very different security implications. Most sites want to allow their users access to outgoing Telnet service, so their users can get to command shells and information services provided via Telnet on remote systems on the Internet. (Figure 8.5 : illustrates outbound Telnet.) On the other hand, most sites do not want to allow incoming Telnet access to their site.

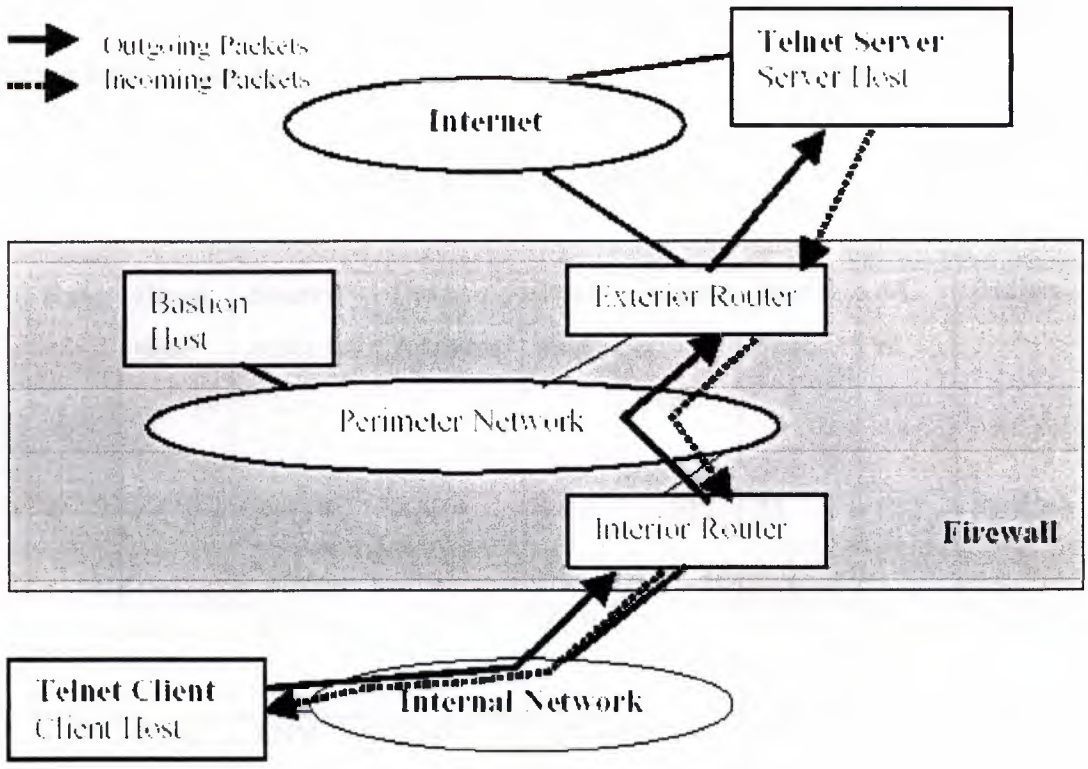

Figure 8.5 Outbound Telnet

Table 8.1 Types of packets involved in inbound and outbound Telnet services

Service Direction	Packet Direction	Source Address	Dest. Address	Packet Type	Source Port	Dest. Port	ACK Set
Outbound	Outgoing	Internal	External	TCP	Y	23	1
Outbound	Incoming	External	Internal	TCP	23	Y	YES
Inbound	Incoming	External	Internal	TCP	Z	23	1
Inbound	Outgoing	Internal	External	TCP	23	Z	YES

The table illustrates the various types of packets involved in inbound and outbound Telnet Services.

Note that Y and Z are both random (from the packet filtering system's point of view) port numbers above 1023.

If you want to allow outgoing Telnet, but nothing else, you would set up your packet filtering like this:

Table 8.2 Packet Filtering Rules

Rule	Direction	Source Address	Dest. Address	Protocol	Source Port	Dest. Port	ACK Set	Action
A	OUT	INTERNAL	ANY	TCP	>1023	23	EITHER	PERMIT
B	IN	ANY	INTERNAL	TCP	23	>1023	YES	PERMIT
C	EITHER	ANY	ANY	ANY	ANY	ANY	EITHER	DENY

- Rule A allows packets out to remote Telnet servers.
- Rule B allows the returning packets to come back in. Because it verifies that the ACK bit is set, rule B can't be abused by an attacker to allow incoming TCP connections from port 23 on the attacker's end to ports above 1023 on your end, i.e., an X11 server on port 6000.
- Rule C is the default rule. If none of the preceding rules apply, the packet is blocked. Remember from our discussion above that any blocked packet should be logged, and that it may or may not cause an ICMP message to be returned to the originator.

Recommendations:

- Restrict incoming Telnet as far possible; most sites have little or no need for it.
- Outgoing Telnet can safely be allowed via packet filtering or proxying.
- If you're concerned about the sensitivity of the data accessed over Telnet sessions, consider using an encrypting version of Telnet.

8.4 HTTP

HTTP is a TCP based service. Most servers use port 80, but some don't. Your firewall will probably prevent people on your internal network from setting up their own servers at non-standard ports (you are not going to want to allow inbound connection to arbitrary ports above 1023). You could set up such servers on a bastion host, but wherever possible, it is kinder to other sites to leave your servers on the standard port.

The following table illustrates the various types of packets involved in inbound and outbound HTTP services.

Table 8.3 Types of packets involved in inbound and outbound HTTP services

Dire-ction	Source Addr.	Dest. Addr.	Pro-Tocol	Source Port	Dest. Port	ACK Set	Notes
In	Ext	Int	TCP	>1023	80₂	3	Incoming session, client to server
Out	Int	Ext	TCP	80₂	>1023	YES	Incoming session, server to client
Out	Int	Ext	TCP	>1023	80₂	3	Outgoing session, client to server
In	Ext	Int	TCP	80₂	>1023	YES	Outgoing session, server to client

Recommendations:

- If you are going to run an HTTP server, use a dedicated bastion host if possible.
- If you are going to run an HTTP server, carefully configure the HTTP server to control; in particular, watch out for ways that someone could upload a program to the system somehow (via mail or FTP, for example), and then execute it via the http server.

• Carefully control the external programs your HTTP server can access.

• You can not allow internal hosts to access all HTTP servers without allowing them to access all TCP ports, because some HTTP servers use non-standard port numbers. If you do not mind allowing your users access to all TCP ports, you can use packet filtering to examine the ACK bit to allow outgoing connections to those ports (but not incoming connections from those ports). If you do mind, then either restrict your users to servers on the standard port (80), or use proxying.

• Proxying HTTP is easy, and a caching proxy server offers network bandwidth benefits as well as security benefits.

• Configure your HTTP clients carefully and warn your users not to reconfigure them based on external advice.

8.5 IIOP

IIOP is a TCP based service. IIOP servers (Orbix Daemon) normally use port 1570. IIOP clients use ports defined above 1023. The following table illustrates the various types of packets involved in inbound and outbound IIOP services.

Table 8.4 Types of packets involved in inbound and outbound IIOP services

Dire-ction	**Source Addr.**	**Dest. Addr.**	**Pro-Tocol**	**Source Port**	**Dest. Port**	**ACK Set**	**Notes**
In	Ext	Int	TCP	>1023	1570	4	Incoming query, client to server
Out	Int	Ext	TCP	1570	>1023	YES	Outgoing response, server to client
Out	Int	Ext	TCP	>1023	1570	4	Outgoing query, client to server
In	Ext	Int	TCP	1570	>1023	YES	Incoming response, server to client

CONCLUSION

The only way to truly secure the computer is to isolate it from non-secured networks. Organisations are increasingly finding that they lose too much by adopting this approach and that an Internet connection is becoming a commercial requirement. Initially protection was provided by careful system management and clever router configuration. However as the number of, and demand for, Internet services grew proxy servers or packet filters have been the standard means of protecting an organisation from the Internet.

The role of the firewall is developing with the changing usage of the Internet as illustrated by the emergence of Intranets and the need to protect internal in addition to external boundaries. Firewalls are a stop gap measure that is needed because many services are developed that operate either with poor security or no security at all. Perhaps the most important lesson we can learn from firewalls is the need for strong session level authentication in applications and well designed application protocols.

Firewall administration requires a seasoned systems manager. While tools are fairly easy to install, it assumes an amount of expertise on the part of the administrator, since he must know how to interpret error conditions, configure the system, and disable potentially threatening services. While it is a temptation to make tools self-installing and selfconfiguring, doing so raises the possibility that someone might install it who lacks the basic skills necessary to know if they have in fact secured their network. Packaging tools as a set of components that can be used freely has proven effective, since it fills a need on the part of those experienced system managers who would have had to design, write, debug, and test their own implementations.

Firewall systems can increase the security in every network. But there can be noguarantee for total security. There is no universal firewall solution. Every solution is depending on the existing environment

Organisations are now looking beyond passively protecting themselves to see what they can do to use the Internet for competitive advantage. This will lead to increase authentication and encryption technology (possibly using kerberos protocolsfor making firewalls with a secure connecting between the different networks of virtual private networks.

ACRONYMS

AFS Andrew File System

ARP Address Resolution Protocol

ARPANET Advanced Research Projects Agency Network

ASCII American Standard Code for Information Interchange

ASIS American Society for Industrial Security

BSD Berkeley Software Distribution

CAGR Compound Annual Growth Rate

CCITT Consultive Committee on International Telegraphy and Telephony

CERT Computer Emergency Response Team

CORBA Common Object Request Broker Architecture

DARPA Defence Advanced Research Projects Agency

DISA Defence Information Systems Agency

DoD Department of Defence

DMZ Demilitarised Zone

DNS Domain Name System

FSP Sneaky File Transfer Protocol

FTP File Transfer Protocol

GIOP General Inter Object Request Broker Protocol

HTML HyperText Mark-up Language

HTTP HyperText Transfer Protocol

IANA Internet Assigned Numbers Authority

ICMP Internet Control Message Protocol

IESG Internet Engineering Steering Group

IETF Internet Engineering Task Force

IIOP Internet Inter-ORB Protocol

IOR Interoperable Object Reference

IP Internet Protocol

IRC Internet Relay Chat

ISN Initial Sequence Number

ISO/ International Standards Organisation

OSI Open System Interconnection Model

LAN Local Area Network

NAP Network Access Point

NAT Network Address Translation

NCSA National Computer Security Agency

NFS Network File System

NIS Network Information Service

NNTP Network News Transfer Protocol

NSFNET National Science Foundation Network

NTP Network Time Protocol

OMG Object Management Group

ONC Open Network Computing

ORB Object Request Broker

PCT Private Communications Technology Protocol

PEM Privacy Enhanced Mail

PIN Personal Identification Number

PKCS Public-Key Cryptography Standards

PPP Point-to-Point Protocol

RFC Request For Comment

RLOGIN Remote LOGIN

RPC Remote Procedure Call

S/MIME Secure Multipurpose Internet Mail Extensions

SLIP Serial Line Internet Protocol

SMTP Simple Mail Transfer Protocol

SSL Secure Sockets Layer

TCP Transmission Control Protocol

TELNET Standard Terminal Emulation Protocol

TFTP Trivial File Transfer Protocol

TOS Type of Service

UDP User Datagram Protocol

WWW World Wide Web

XNS Xerox Networking System

FIGURES

TABLES

GLOSSARY OF TERMS

· Abuse of privilege: when a user performs an action that they should not have, according to organisational policy or law.

· Application-level gateway firewall: a firewall system in which service is provided by processes that maintain complete TCP connection state and sequencing. Application level firewalls often re-address traffic so that outgoing traffic appears to have originated from the firewall, rather than the internal host.

· Authentication: the process of determining the identity of a user that is attempting to access a system.

· Authentication token: a portable device used for authenticating a user. Authentication tokens operate by challenge/response, time-based code sequences, or other techniques. This may include paper based lists of one time passwords.

· Authorisation: the process of determining what types of activities are permitted. Usually, authorisation is in the context of authentication: once you have authenticated a user, they may be authorised different types of access or activity.

· Bastion host: a system that has been hardened to resist attack, and which is installed on a network in such a way that it is expected to potentially come under attack. Bastion hosts are often components of firewalls, or may be "outside" Web servers or public access systems. Generally, a bastion host is running some form of general purpose operating system (e.g., UNIX, VMS, WNT, etc.) rather than a ROM-based or firmware operating system.

· Challenge/response: an authentication technique whereby a server sends an unpredictable challenge to the user, who computes a response using some form of authentication token.

· Chroot: a technique under UNIX whereby a process is permanently restricted to an isolated subset of the filesystem.

· Cryptographic checksum: a one way function applied to a file to produce a unique "fingerprint" of the file for later reference. Checksum systems are a primary means of detecting filesystem tampering on UNIX.

· Data driven attack: a form of attack in which the attack is encoded in innocuousseeming data which is executed by a user or other software to implement an attack. In the case of firewalls, a data driven attack is a concern since it may get through the firewall in data form and launch an attack against a system behind the firewall.

· Defence in depth: the security approach whereby each system on the network is secured to the greatest possible degree. May be used in conjunction with firewalls.

· DNS spoofing: assuming the DNS name of another system by either corrupting the name service cache of a victim system, or by compromising a domain name server for a valid domain.

· Dual homed gateway: a system that has two or more network interfaces, each of which is connected to a different network. In firewall configurations, a dual homed gateway usually acts to block or filter some or all of the traffic trying to pass between the networks.

· Encrypting router: see tunneling router and virtual private network.

· Firewall: a system or combination of systems that enforces a boundary between two or more networks, controlling access from one to the other.

· Gateway: a system that provides and controls access and from a network. Also, an application-level gateway, a machine or set of machines that relays services between the internal and external networks by means of proxy applications.

· GGP: Gateway to Gateway Protocol, the protocol used primarily between gateways to control routing and other gateway functions.

· Header: control information at the beginning of a message, segment, datagram, packet or block of data.

· Host-based security: the technique of securing an individual system from attack. Host based security is operating system and version dependent.

· ICMP: Internet Control Message Protocol, implemented in the internet module, the ICMP is used from gateways to hosts and between hosts to report errors and make routing suggestions.

· Insider attack: an attack originating from inside a protected network.

· Internet Address: a four octet (32 bit) source or destination address consisting of a Network field and a Local Address field.

· Internet datagram: the unit of data exchanged between a pair of internet modules (includes the internet header).

· Internet, the: global network of computers that is the basis for universal electronic mail, the World Wide Web, and numerous forms of electronic commerce. Typically, we reserve the term Internet for the TCP/IP-based descendant of ARPAnet's marriage to CSnet in 1982, now serving tens of millions of users via hundreds of thousands of host machines.

· Internet Protocol: one of two major protocols in the Internet Protocol Suite, otherwise known as TCP/IP, of which Internet Protocol is the IP. See IP.

· Internet Protocol Suite: official name of TCP/IP, as used in Internet standards documents, see TCP/IP.

· Internetwork: the process of connecting two networks together. The result is referred to as an internet without a capital `I.'

· Intranet: a closed network of computers that uses similar technology to the Internet, such as Web servers and browsers, to make information available to users to a controlled group of users. An intranet may have a connection to the Internet, or it may exist on the Internet, achieving controlled access through passwords or other means.

· Intrusion detection: detection of break-ins or break-in attempts either manually or via software expert systems that operate on logs or other information available on the network.

· IP: Internet Protocol, provides host to host communication. IP is referred to as an unreliable datagram service, meaning that upper-level protocols should not depend upon IP to deliver the packet every time. IP does its best to make the delivery to the Glossary requested destination host, but if it fails for any reason, it just drops the packet

· IP spoofing: an attack whereby a system attempts to illicitly impersonate another system by using its IP network address. For example, someone might determine the IP address of a legitimate user inside the firewall, then forge packets from outside the firewall which the firewall allows to pass because they are from a legitimate user.

· IP splicing /hijacking: an attack whereby an active, established, session is intercepted and co-opted by the attacker. IP Splicing attacks may occur after an authentication has been made, permitting the attacker to assume the role of an already authorised user. Primary protections against IP splicing rely on encryption at the session or network layer.

· Least privilege: designing operational aspects of a system to operate with a minimum amount of system privilege. This reduces the authorisation level at which various actions are performed and decreases the chance that a process or user with high privileges may be caused to perform unauthorised activity resulting in a security breach.

· Logging: the process of storing information about events that occurred on the firewall or network.

· Log retention: how long audit logs are retained and maintained.

· Log processing: how audit logs are processed, searched for key events, or summarised.

· Network-level firewall: a firewall in which traffic is examined at the network protocol packet level.

· Perimeter based security: the technique of securing a network by controlling access to all entry and exit points of the network.

· Policy: organisation-level rules governing acceptable use of computing resources, security practices, and operational procedures.

· Protocol: a formal description of messages to be exchanged and rules to be followed for two or more systems to exchange information.

· Proxy: a software agent that acts on behalf of a user. Typical proxies accept a connection from a user, make a decision as to whether or not the user or client IP address is permitted to use the proxy, perhaps does additional authentication, and then completes a connection on behalf of the user to a remote destination.

· Screened host: a host on a network behind a screening router. The degree to which a screened host may be accessed depends on the screening rules in the router.

· Screened subnet: a subnet behind a screening router. The degree to which the subnet may be accessed depends on the screening rules in the router.

· Screening router: a router configured to permit or deny traffic based on a set of permission rules installed by the administrator.

· Session stealing: See IP splicing.

· Social engineering: An attack based on deceiving users or administrators at the target site. For example, telephoning users or operators pretending to be an authorised user, to attempt to gain illicit access to systems.

· S/Key: freely available authentication system, developed at Bellcore (based on a paper by Leslie Lamport of DEC) that avoids many types of password attack [Amor96].

· S/WAN: emerging standard for secure firewall-to-firewall communication.

· TCP/IP: Transmission Control Protocol/Internet Protocol, otherwise known as the Internet Protocol Suite.

· Transmission Control Protocol: protocol that provides reliable transmission of packets over IP [Amor96].

· Trojan horse: a software entity that appears to do something normal but which, in fact, contains a trapdoor or attack program.

· Tunneling router: A router or system capable of routing traffic by encrypting it and encapsulating it for transmission across an distrusted network, for eventual deencapsulation and decryption.

· UDP: User Datagram Protocol: a user level protocol for transaction oriented applications.

· Virtual private network: a network that appears to be a single protected network behind firewalls, which actually encompasses encrypted virtual links over distrusted networks.

· Virus: a self-replicating code segment. Viruses may or may not contain payloads, attack programs or trapdoors.

REFERENCES

[Amor96] Edward Amoroso and Ronald Sharp, "Intranet and Internet Firewall Strategies", (Ziff Davis Press, 1996)

[Bern96] Berners-Lee T, Fielding R., Nielsen : "Hypertext Transfer Protocol – HTTP/1.0", draft-ietf-http-v10-spec-05.txt, 20 Feb 96

[Bild96] Bilodeau Anne : "Hacking the Web - A security guide", Web Developer, Vol. 1, No. 1, Winter 1996

[Book96] Booker Ellis : "Security auditing business is booming", Web Week Vol. 2, Issue 3, Mar 96

[Bray96] Bray Paul : "A quick easy and cheap solution", The Sunday Times, 28 Apr 96

[CA-95:01] CERT : "IP Spoofing Attacks and Hijacked Terminal Connections", CERT, Jan 95

[CA-96:01] CERT : "UDP Port Denial-of-Service Attack", CERT, Jan 96

[CCIT88] CCITT : "Recommendation X.509 : The Directory – Authentication Framework.", Consultative Committee on Internation Telegraphy and Telephony , 1988

[Cerf74] Cerf Vinton G. and Kahn Robert G. : "A Protocol For Packet Network Interconnection", IEEE Transactions of Communications, May 1974, May 74

[Chad94] Chadwick David : "Understanding X.500 The Directory", Chapman & Hall, 1994

[Chap92] Chapman D. Brent. : "Network (In)Security Through IP Packet Filtering.", In USENIX Security Symposium III Proceedings, pages 63- 76. USENIX Association., September 14-16 1992

[Chap95] Chapman D. Brent and Zwicky Elizabeth D. : "Building Internet Firewalls", O'Reilly and Associates, Inc., Sebastopol, CA, Sep 95

[Ches94] Cheswick William R. and Bellovin Steven M. : "Firewalls and Internet Security-Repelling the Wily Hacker.", Addison-Wesley, Reading, MA, 1994

[Clar96] Clark Tim : "Digital ID Center Opens", Inter@ctive Week, 29 Apr 96

[Cobb95] Cobb Stephen : "NCSA Firewall Policy Guide", NCSA Security White Paper Series, 1995

[CSI95a] Computer Security Institute : "CSI's 1995 Firewall Product Matrix", Computer Security Issues and Trends, Fall 1995

[CSI96] Computer Security Institute : "CSI's 1996 Firewall Product Matrix", Computer Security Issues and Trends, Spring 1996

[Dah94] Dahn David : "Security and the World Wide Web.", http://www.tis.com, Jun 94

[Diff76] Diffie W. and Hellman M.E. : "New directions in cryptography", IEEE Transactions on Information Theory. IT-22, pp 644-654, 1976

[DOD85] US Department of Defense : "Trusted Computer System Evaluation Criteria (The Orange Book)", DoD5200.28-STD, 1985

[Eich89] Eichin Mark W. and Rochlis Jon A. : "With Microscope and Tweezers: An Analysis of the Internet Virus of November 1988", Available from URL ftp://athena-dist.mit.edu/pub/virus/mit.ps, 9 Feb 89

[Farm94] Farmer Daniel and Spafford Eugene H. : "The Cops Security Checker System", Purdue University Technical Report CSD-TR-993, 28 Jul 94

[Field96] Fielding R., Frystyk H., Berners-Lee T. : "Hypertext Transfer Protocol – HTTP/1.1", draft-ietf-http-v11-spec-02.txt, 23 Apr 96

[Frei96] Freier A., Karlton P., Kocher P. : "The SSL Protocol Version 3.0", draft-freier-ssl-version3-01.txt, 13 Mar 96

[Gruber96] Gruber, Peter : "Firewall-Systeme kaufen Hackern den Schneid ab", Computerwoche 46/96

[Hews96] Hewson David : "Health Warning : Be safe or be sorry", The Sunday Times, 28 Apr 96

[Holt96] Holtman K. : "Proposed Content Negotiation Definitions for HTTP/1.", draft-holtman-http-negotiation-00.txt, 35118

[Hopm96] Hopmann A. : "Persistent HTTP Connections", draft-ietf-http-ses-ext- 01.txt, 21 Feb 96

[Host96] Hostetler J., Franks J., Hallam-Baker P. : "A Proposed Extension To HTTP : Digest Access Authentication", draft-ietf-http-digest-aa-03.text, 22 Mar 96

[Hove96] Hovey R. and Bradner S. : "The Organizations Involved in the IETF Standards Process", <draft-ietf-poised95-ietf-orgs-00.txt>, 22 Feb 96

[Hugh95] Hughes Larry J. Jnr : "Actually Useful Internet Security Techniques", New Riders Publishing, Indianapolis, Indiana, 1995

[Hunt92] Hunt Craig : "TCP/IP Network Administration.", O'Reilly and Associates, Inc., Sebastopol, CA, 1992

[IDC96] Julian Ted : "Internet Commerce - The Worldwide Firewall Market : 1995-2000", International Data Corporation, Feb 96