NIST Special Publication 800-10
U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology

Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

John P. Wack
Lisa J. Carnahan

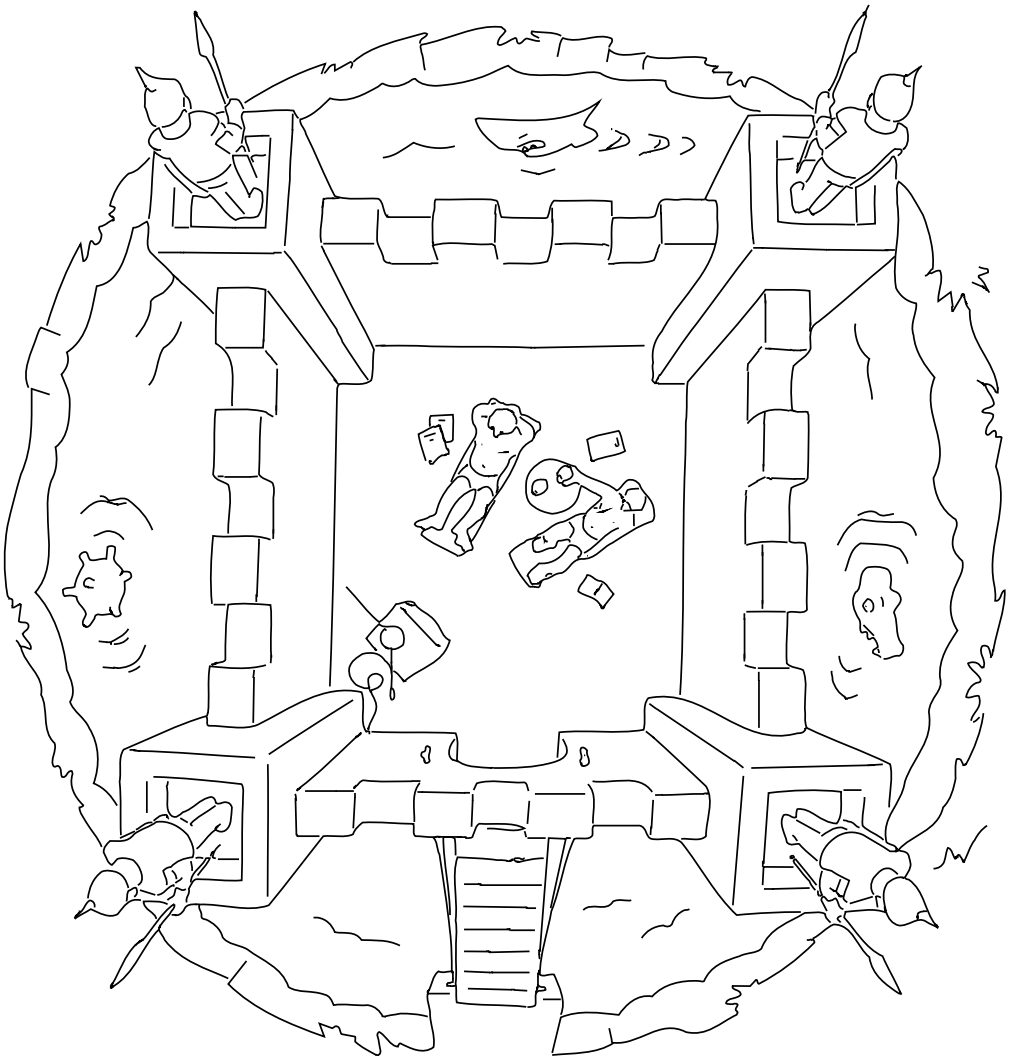

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Documents
Document Sharing Hub

Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Abstract

This document provides an overview of the Internet and security-related problems. It then provides an overview of firewall components and the general reasoning behind firewall usage. Several types of network access policies are described, as well as technical implementations of those policies. Lastly, the document contains pointers and references for more detailed information.

The document is designed to assist users in understanding the nature of Internet-related security problems and what types of firewalls will solve or alleviate specific problems. Users can then use this document to assist in purchasing or planning a firewall.

Documents
Document Sharing Hub

Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

This work is a contribution
of the National Institute of Standards and Technology,
and is not subject to copyright.

Because of the nature of this report, it is necessary to mention vendors and commercial products. The presence or absence of a particular trade name product does not imply criticism or endorsement by the National Institute of Standards and Technology, nor does it imply that the products identified are necessarily the best available.

Acknowledgments

The National Institute of Standards and Technology would like to thank the following individuals who reviewed drafts of this document and advised on document structure and content: David Curry of Purdue University, Uwe Ellermann of the DFN-CERT in Germany, and Stephen Weeber of the Department of Energy's Computer Incident Advisory Capability (CIAC).

Documents
Document Sharing Hub

Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Contents

Preface **ix**

1 Introduction to the Internet and Internet Security **1**
1.1 The Internet . . . 1
1.1.1 Common Services . . . 2
1.1.2 Internet Hosts . . . 3
1.2 Overview of TCP/IP Internals . . . 3
1.2.1 IP . . . 4
1.2.2 TCP . . . 5
1.2.3 UDP . . . 6
1.2.4 ICMP . . . 6
1.2.5 TCP and UDP Port Structure . . . 6
1.3 Security-Related Problems . . . 8
1.3.1 Security Incidents on the Internet . . . 8
1.3.2 Weak Authentication . . . 9
1.3.3 Ease of Spying/Monitoring . . . 10
1.3.4 Ease of Spoofing . . . 10
1.3.5 Flawed LAN Services and Mutually Trusting Hosts . . . 12
1.3.6 Complex Configuration and Controls . . . 12
1.3.7 Host-based Security Does Not Scale . . . 13
1.4 How Vulnerable Are Internet Sites? . . . 13

2 Introduction to Firewalls **15**
2.1 The Firewall Concept . . . 15
2.2 Why Firewalls . . . 16
2.2.1 Protection from Vulnerable Services . . . 16
2.2.2 Controlled Access to Site Systems . . . 17
2.2.3 Concentrated Security . . . 17
2.2.4 Enhanced Privacy . . . 17
2.2.5 Logging and Statistics on Network Use, Misuse . . . 18
2.2.6 Policy Enforcement . . . 18
2.3 Issues and Problems with Firewalls . . . 18
2.3.1 Restricted Access to Desirable Services . . . 18
2.3.2 Large Potential for Back Doors . . . 19
2.3.3 Little Protection from Insider Attacks . . . 19

Documents
Document Sharing Hub

Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

2.3.4 Other Issues . . . 19
2.4 Firewall Components . . . 20
2.4.1 Network Policy . . . 21
2.4.2 Advanced Authentication . . . 22
2.4.3 Packet Filtering . . . 24
2.4.4 Application Gateways . . . 29

3 Putting the Pieces Together: Firewall Examples 33
3.1 Packet Filtering Firewall . . . 33
3.2 Dual-homed Gateway Firewall . . . 34
3.3 Screened Host Firewall . . . 36
3.4 Screened Subnet Firewall . . . 38
3.5 Integrating Modem Pools with Firewalls . . . 40

4 Next Steps 43
4.1 Firewall Policy . . . 43
4.1.1 Steps in Creating a Service Access Policy . . . 43
4.1.2 Flexibility in Policy . . . 45
4.1.3 Remote User Advanced Authentication Policy . . . 45
4.1.4 Dial-in/out Policy . . . 46
4.1.5 Remote Network Connections . . . 46
4.1.6 Information Server Policy . . . 46
4.2 Procuring a Firewall . . . 47
4.2.1 What Should a Firewall Contain? . . . 47
4.2.2 To Buy or Build a Firewall . . . 49
4.3 Administration Issues with Firewalls . . . 50
4.3.1 System Management Expertise . . . 50
4.3.2 Site System Administration . . . 50
4.3.3 Incident Handling Contacts . . . 51

Bibliography 53

A On-Line Sources for More Information 55
A.1 Firewall-Specific Information . . . 55
A.2 NIST Computer Security Resource Clearinghouse . . . 55
A.3 Forum of Incident Response and Security Teams . . . 56

B Internet Firewalls Frequently Asked Questions 59

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

List of Figures

1.1 Conceptual View of Services and Layers in TCP/IP. 4
1.2 TELNET Port, IP Interaction. 7

2.1 Router and Application Gateway Firewall Example. 15
2.2 Use of Advanced Authentication on a Firewall to Preauthenticate TELNET, FTP Traffic. 23
2.3 Representation of Packet Filtering on TELNET and SMTP. 25
2.4 Virtual Connection Implemented by an Application Gateway and Proxy Services. 29

3.1 Packet Filtering Firewall. 34
3.2 Dual-homed Gateway Firewall with Router. 35
3.3 Screened Host Firewall. 37
3.4 Screened Subnet Firewall with Additional Systems. 39
3.5 Modem Pool Placement with Screened Host Firewall. 41
3.6 Modem Pool Placement with Screened Subnet and Dual-Homed Firewalls. 42

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Preface

The Internet is a world-wide collection of networks that all use a common protocol for communications. Many organizations are in the process of connecting to the Internet to take advantage of Internet services and resources. Businesses and agencies are now using the Internet or considering Internet access for a variety of purposes, including exchanging e-mail, distributing agency information to the public, and conducting research. Many organizations are connecting their existing internal local area networks to the Internet so that local area network workstations can have direct access to Internet services.

Internet connectivity can offer enormous advantages, however security needs to be a major consideration when planning an Internet connection. There are significant security risks associated with the Internet that often are not obvious to new (and existing) users. In particular, intruder activity as well as vulnerabilities that could assist intruder activity are widespread. Intruder activity is difficult to predict and at times can be difficult to discover and correct. Many organizations already have lost productive time and money in dealing with intruder activity; some organizations have had their reputations suffer as a result of intruder activity at their sites being publicized.

This publication focuses on security considerations for organizations considering Internet connections as well as for organizations already connected to the Internet. In particular, this document focuses on Internet *firewalls* as one of the mechanisms and methods used for protecting sites against Internet-borne threats. This document recommends that organizations use firewall technology and other related tools to filter connections and limit access. This document is an expansion of the issues and guidance contained in NIST CSL Bulletin, *Connecting to the Internet: Security Considerations* [NIST93].

Purpose

The purpose of this document is to provide a basis of understanding of how firewalls work and the steps necessary for implementing firewalls. Users can then use this document to assist in planning or purchasing a firewall. This document does not explain how to build a firewall; references are provided for more detailed information.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Audience

The intended audience of this publication is technical-level management, i.e., those individuals who may be responsible for implementing or maintaining Internet connections. This document would also be appropriate for other management who wish to learn more about Internet security issues.

Some technical background in computer security and computer network communications is assumed. However, this document is intended to be a starting point; more detailed information about Internet security and firewalls can be found in the references section.

Document Structure

This document begins with an overview of the Internet and common services. It describes Internet-related security problems in detail by examining problems with various TCP/IP services and by examining other factors that have caused the Internet to grow less secure. Chapter 2 discusses firewalls, their benefits as well as their disadvantages, and then the various firewall components, including advanced authentication measures and network access policy. Chapter 3 describes several firewall configurations that illustrate how the firewall components fit together and can be used to implement various policies. Chapter 4 discusses procurement, administrative issues, and other actions sites should take to secure their Internet-connected systems. Appendix A provides pointers to other books and information about firewalls and Internet security. Appendix B contains a collection of frequently asked questions about firewalls that is available on-line (see Appendix B for more information).

Terminology

Internet firewalls are often referred to as *secure Internet gateways* in other literature. This document uses *firewall* to refer to a secure Internet gateway.

A firewall, as defined in this document, includes a number of items such as policy, network arrangement, and technical controls and procedures. This document uses *firewall system* when referring to the hosts or routers that implement the firewall.

This document, when referring to a network protected by a firewall, uses *protected subnet* or *protected LAN* (Local Area Network).

Some people dispute whether TCP/IP protocols should be referred to as protocols or services. It could be argued, for example, that TELNET is a protocol, a service, or a command. Where it makes obvious sense, this document uses *protocol*, otherwise it uses

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

service.

This document uses *application gateways* to refer to some firewall systems as opposed to *bastion hosts.*

As much as possible, this document avoids using terms such as *hacker* and *cracker*, and uses instead the less ambiguous *intruder* and *attacker.*

Background

The Internet is a vital and growing network that is changing the way many organizations and individuals communicate and do business. However, the Internet suffers from significant and widespread security problems. Many agencies and organizations have been attacked or *probed*[1] by intruders, with resultant high losses to productivity and reputation. In some cases, organizations have had to disconnect from the Internet temporarily, and have invested significant resources in correcting problems with system and network configuration. Sites that are unaware of or ignorant of these problems face a significant risk that they will be attacked by network intruders. Even sites that do observe good security practices face problems with new vulnerabilities in networking software and the persistence of some intruders.

A number of factors have contributed to this state of affairs. The fundamental problem may be that the Internet was not designed to be very secure, i.e., open access for the purposes of research was the prime consideration at the time the Internet was implemented. However, the phenomenal success of the Internet in combination with the introduction of different types of users, including unethical users, has aggravated existing security deficiencies to the extent that wide-open Internet sites risk inevitable break-ins and resultant damages. Other factors include the following:

- **vulnerable TCP/IP services** - a number of the TCP/IP services are not secure and can be compromised by knowledgeable intruders; services used in the local area networking environment for improving network management are especially vulnerable,
- **ease of spying and spoofing** - the majority of Internet traffic is unencrypted; e-mail, passwords, and file transfers can be monitored and captured using readily-available software, intruders can then reuse passwords to break into systems,
- **lack of policy** - many sites are configured unintentionally for wide-open Internet access without regard for the potential for abuse from the Internet; many sites

[1]Intruders have been observed to target specific sites for intrusions by methodically scanning host systems for vulnerabilities. Intruders often use automated probes, i.e., software that scans all host systems connected to a site's network. This is sometimes referred to as *probing a site.*

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

permit more TCP/IP services than they require for their operations and do not attempt to limit access to information about their computers that could prove valuable to intruders, and

- **complexity of configuration** - host security access controls are often complex to configure and monitor; controls that are accidentally misconfigured often result in unauthorized access.

Solutions

Fortunately, there are readily-available solutions that can be used to improve site security. A firewall system is one technique that has proven highly effective for improving the overall level of site security. A firewall system is a collection of systems, routers, and policy placed at a site's central connection to a network. A firewall forces all network connections to pass through the gateway where they can be examined and evaluated, and provides other services such as advanced authentication measures to replace simple passwords. The firewall may then restrict access to or from selected systems, or block certain TCP/IP services, or provide other security features. A well-configured firewall system can act also as an organization's "public-relations vehicle" and can help to present a favorable image of the organization to other Internet users.

A simple network usage policy that can be implemented by a firewall system is to provide access from internal to external systems, but little or no access from external to internal systems. However, a firewall does not negate the need for stronger system security. There are many tools available for system administrators to enhance system security and provide additional logging capability. Such tools can check for strong passwords, log connection information, detect changes in system files, and provide other features that will help administrators detect signs of intruders and break-ins.

Recommendations

NIST recommends that agencies and organizations, prior to connecting to the Internet, develop policy that clearly identifies the Internet services they will be using and how those services will be used. The policy should be clear, concise, and understandable, with a built-in mechanisms for changing the policy. Organizations should strongly consider using firewall systems as part of the implementation of that policy. NIST recommends also that agencies and organizations use advanced authentication measures, i.e., smartcards, or authentication tokens, or other one-time password mechanisms, as an integral part of firewalls for authenticating connections to site systems.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

1

Introduction to the Internet and Internet Security

While Internet connectivity offers enormous benefits in terms of increased access to information, Internet connectivity is not necessarily a good thing for sites with low levels of security. The Internet suffers from glaring security problems that, if ignored, could have disastrous results for unprepared sites. Inherent problems with TCP/IP services, the complexity of host configuration, vulnerabilities introduced in the software development process, and a variety of other factors have all contributed to making unprepared sites open to intruder activity and related problems.

The following sections present a brief overview of the Internet, TCP/IP, and then explain what some of the Internet security related problems are and what factors have contributed to their seriousness.

1.1 The Internet

The Internet is a world-wide "network of networks" that use the TCP/IP (Transmission Control Protocol/Internet Protocol) protocol suite for communications. The Internet was created initially to help foster communication among government-sponsored researchers. Throughout the 1980's, the Internet grew steadily to include educational institutions, government agencies, commercial organizations, and international organizations. In the 1990's, the Internet has undergone phenomenal growth, with connections increasing faster than any other network ever created (including the telephone network). Many millions of users are now connected to the Internet, with roughly half being business users [Cerf93]. The Internet is being used as the basis for the National Information Infrastructure (NII).

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

1.1.1 Common Services

There are a number of services associated with TCP/IP and the Internet. The most commonly used service is electronic mail (*e-mail*), implemented by the Simple Mail Transfer Protocol (SMTP). Also, TELNET (terminal emulation), for remote terminal access, and FTP (file transfer protocol) are used widely. Beyond that, there are a number of services and protocols used for remote printing, remote file and disk sharing, management of distributed databases, and for information services. Following is a brief list of the most common services:

- **SMTP** - Simple Mail Transfer Protocol, used for sending and receiving electronic mail,
- **TELNET** - used for connecting to remote systems connected via the network, uses basic terminal emulation features,
- **FTP** - File Transfer Protocol, used to retrieve or store files on networked systems,
- **DNS** - Domain Name Service, used by TELNET, FTP, and other services for translating host names to IP addresses,
- **information-based services, such as**
  - **gopher** - a menu-oriented information browser and server that can provide a user-friendly interface to other information-based services,
  - **WAIS** - Wide Area Information Service, used for indexing and searching with databases of files, and
  - **WWW/http** - World Wide Web, a superset of FTP, gopher, WAIS, other information services, using the hypertext transfer protocol (http), with *Mosaic* being a popular WWW client,
- **RPC-based services** - Remote Procedure Call services, such as
  - **NFS** - Network File System, allows systems to share directories and disks, causes a remote directory or disk to appear to be local, and
  - **NIS** - Network Information Services, allows multiple systems to share databases, e.g., the password file, to permit centralized management,
- **X Window System** - a graphical windowing system and set of application libraries for use on workstations, and
- **rlogin, rsh, and other "r" services** - employs a concept of mutually trusting hosts, for executing commands on other systems without requiring a password.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Although TCP/IP can be used equally well in a local area or wide area networking environment, a common use is for file and printer sharing at the local area networking level and for electronic mail and remote terminal access at both the local and the wide area networking levels. Gopher and Mosaic are increasingly popular; both present problems to firewall designers as will be discussed in later sections.

1.1.2 Internet Hosts

Many host systems connected to the Internet run a version of the UNIX operating system. TCP/IP was first implemented in the early 1980's for the version of UNIX written at the University of California at Berkeley known as the Berkeley Software Distribution (BSD). Many modern versions of UNIX derive their networking code directly from the BSD releases, thus UNIX provides a more-or-less standard set of TCP/IP services. This standard of sorts has resulted in many different versions of UNIX suffering from the same vulnerabilities, however it has also provided a common means for implementing firewall strategies such as IP packet filtering. It is important to note that BSD UNIX source code is fairly easy to obtain free from Internet sites, thus many good and bad people have been able to study the code for potential flaws and exploitable vulnerabilities.

Although UNIX is the predominant Internet host operating system, many other types of operating systems and computers are connected to the Internet, including systems running Digital Equipment Corporation's VMS, NeXT, mainframe operating systems, and personal computer operating systems such as for DOS, Microsoft Windows, and for Apple systems. Although personal computer systems often provide only client services, i.e., one can use TELNET to connect *from* but not *to* a personal computer, increasingly powerful personal computers are also beginning to provide, at low cost, the same services as larger hosts. Versions of UNIX for the personal computer, including Linux, FreeBSD, and BSDi, and other operating systems such as Microsoft Windows NT, can provide the same services and applications that were, until recently, found only on larger systems. The ramifications of this are that more people are able to utilize a wider array of TCP/IP services than ever before. While this is good in that the benefits of networking are more available, it has negative consequences in that there is more potential for harm from intruders (as well as uneducated but well-intentioned users who, to some sites, may appear to be intruders).

1.2 Overview of TCP/IP Internals

This section provides a simplified overview of TCP/IP for the purposes of later discussion on Internet-related security problems. [Com91a], [Com91b], [Ford94], [Hunt92],

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

and [Bel89] provide more complete descriptions; readers who wish to learn more should consult these references.

Part of the popularity of the TCP/IP protocol suite is due to its ability to be implemented on top of a variety of communications channels and lower-level protocols such as T1 and X.25, Ethernet, and RS-232-controlled serial lines. Most sites use Ethernet connections at local area networks to connect hosts and client systems, and then connect that network via a T1 line to a regional network (i.e., a regional TCP/IP backbone) that connects to other organizational networks and backbones. Sites customarily have one connection to the Internet, but large sites often have two or more connections. Modem speeds are increasing as new communications standards are being approved, thus versions of TCP/IP that operate over the switched telephone network are becoming more popular. Many sites and individuals use PPP (Point-to-Point Protocol) and SLIP (Serial Line IP), to connect networks and workstations to other networks using the switched telephone network.

TCP/IP is more correctly a *suite* of protocols including TCP and IP, UDP (User Datagram Protocol), ICMP (Internet Control Message Protocol), and several others. The TCP/IP protocol suite does not conform exactly to the Open Systems Interconnection's seven layer model, but rather could be pictured as shown in figure 1.1.

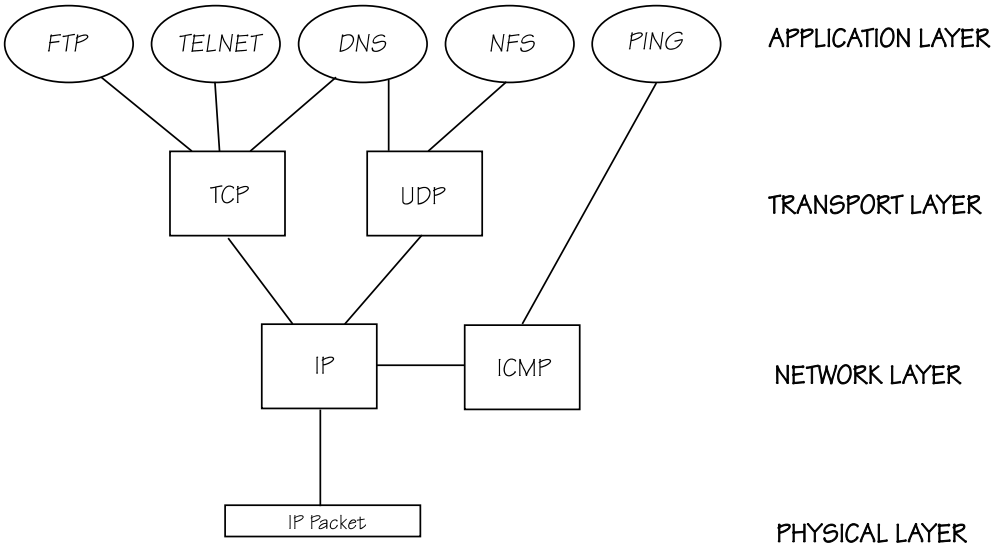

Figure 1.1: Conceptual View of Services and Layers in TCP/IP.

1.2.1 IP

The IP layer receives packets delivered by lower-level layers, e.g., an Ethernet device driver, and passes the packets "up" to the higher-layer TCP or UDP layers. Conversely,

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

IP transmits packets that have been received from the TCP or UDP layers to the lower-level layer.

IP packets are *unreliable datagrams* in that IP does nothing to ensure that IP packets are delivered in sequential order or are not damaged by errors. The IP packets contain the address of the host from which the packet was sent, referred to as the *source* address, and the address of the host that is to receive the packet, referred to as the *destination* address.

The higher-level TCP and UDP services generally assume that the source address in a packet is valid when accepting a packet. In other words, the IP address forms the basis of authentication for many services; the services trust that the packet has been sent from a valid host and that host is indeed who it says it is. IP does contain an option known as IP *Source Routing*, which can be used to specify a direct route to a destination and return path back to the origination. The route could involve the use of other routers or hosts that normally would not be used to forward packets to the destination. A source routed IP packet, to some TCP and UDP services, appears to come from the last system in the route as opposed to coming from the true origination. This option exists for testing purposes, however [Bel89] points out that source routing can be used to trick systems into permitting connections from systems that otherwise would not be permitted to connect. Thus, that a number of services trust and rely on the authenticity of the IP source address is problematic and can lead to breakins and intruder activity.

1.2.2 TCP

If the IP packets contain encapsulated TCP packets, the IP software will pass them "up" to the TCP software layer. TCP sequentially orders the packets and performs error correction, and implements *virtual circuits*, or connections between hosts. The TCP packets contain sequence numbers and acknowledgements of received packets so that packets received out of order can be reordered and damaged packets can be retransmitted.

TCP passes its information up to higher-layer applications, e.g., a TELNET client or server. The applications, in turn, pass information back to the TCP layer, which passes information down to the IP layer and device drivers and the physical medium, and back to the receiving host.

Connection oriented services, such as TELNET, FTP, rlogin, X Windows, and SMTP, require a high degree of reliability and therefore use TCP. DNS uses TCP in some cases (for transmitting and receiving domain name service databases), but uses UDP for transmitting information about individual hosts.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

1.2.3 UDP

As shown in figure 1.1, UDP interacts with application programs at the same relative layer as TCP. However, there is no error correction or retransmission of misordered or lost packets. UDP is therefore not used for connection-oriented services that need a virtual circuit. It is used for services that are query-response oriented, such as NFS, where the number of messages with regard to the exchange is small compared to TELNET or FTP sessions. Services that use UDP include RPC-based services such as NIS and NFS, NTP (Network Time Protocol), and DNS (DNS also uses TCP).

It is easier to spoof UDP packets than TCP packets, since there is no initial connection setup (handshake) involved (since there is no virtual circuit between the two systems) [Ches94]. Thus, there is a higher risk associated with UDP-based services.

1.2.4 ICMP

ICMP (Internet Control Message Protocol) is at the same relative layer as IP; its purpose is to transmit information needed to control IP traffic. It is used mainly to provide information about routes to destination addresses. ICMP *redirect* messages inform hosts about more accurate routes to other systems, whereas ICMP *unreachable* messages indicate problems with a route. Additionally, ICMP can cause TCP connections to terminate "gracefully" if the route becomes unavailable. *ping* is a commonly-used ICMP-based service.

[Bel89] discusses two problems with ICMP: older versions of UNIX could drop all connections between two hosts even if only one connection was experiencing network problems. Also, ICMP redirect messages can be used to trick routers and hosts acting as routers into using "false" routes; these false routes would aid in directing traffic to an attacker's system instead of a legitimate trusted system. This could in turn lead to an attacker gaining access to systems that normally would not permit connections to the attacker's system or network.

1.2.5 TCP and UDP Port Structure

TCP and UDP services generally have a client-server relationship. For example, a TELNET server process initially sits idle at a system, waiting for an incoming connection. A user then interacts with a TELNET client process, which initiates a connection with the TELNET server. The client writes to the server, the server reads from the client and sends back its response. The client reads the response and reports back to the user. Thus, the connection is bidirectional and can be used for reading and writing.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

How are multiple TELNET connections between two systems identified and coordinated? A TCP or UDP connection is uniquely identified by the following four items present in each message:

- **source IP address** - the address of the system that sent the packet,
- **destination IP address** - the address of the system that receives the packet,
- **source port** - the connection's port at the source system, and
- **destination port** - the connection's port at the destination system.

The port is a software construct that is used by the client or server for sending or receiving messages; a port is identified by a 16-bit number. Server processes are usually associated with a fixed port, e.g., 25 for SMTP or 6000 for X Windows; the port number is "well-known" because it, along with the destination IP address, needs to be used when initiating a connection to a particular host and service. Client processes, on the other hand, request a port number from the operating system when they begin execution; the port number is random although in some cases it is the next available port number.

As an example of how ports are used for sending and receiving messages, consider the TELNET protocol. The TELNET server listens for incoming messages on port 23, and sends outgoing messages to port 23. A TELNET client, on the same or different system, would first request an unused port number from the operating system, and then use this port when sending and receiving messages. It would place this port number, say 3097, in packets destined for the TELNET server so that the server, when responding to the client, could place the client's port number in its TCP packets. The client's host, upon receiving a message, would examine the port and know which TELNET client should receive the message. This is shown conceptually in figure 1.2.

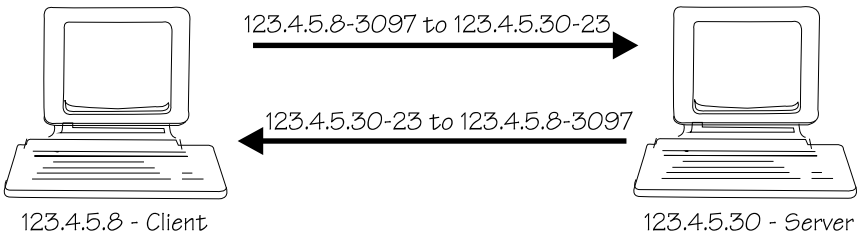

Figure 1.2: TELNET Port, IP Interaction.

There is a somewhat-uniform rule that only *privileged* server processes, i.e., those processes that operate with UNIX superuser privileges, can use port numbers less than 1024 (referred to as *privileged ports*). Servers mostly use ports numbered less than 1024,

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

whereas clients generally must request unprivileged port numbers from the operating system. Although this rule is not firm and is not required in the TCP/IP protocol specifications, BSD-based systems adhere to it. As an accidental but fortuitous result, firewalls can block or filter access to services by examining the port numbers in TCP or UDP packets and then routing or dropping the packet based on a policy that specifies which services are permitted or denied (this is covered in more detail in Chapter 2).

Not all TCP and UDP servers and clients use ports in as straightforward a fashion as TELNET, but in general the procedure described here is useful in the firewalls context. For example, many personal computer operating systems have no UNIX superuser concept, but still use ports as described (although there is no standard that requires this).

1.3 Security-Related Problems

As stated earlier, the Internet suffers from severe security-related problems. Sites that ignore these problems face some significant risk that they will be attacked by intruders and that they may provide intruders with a staging ground for attacks on other networks. Even sites that do observe good security practices face problems with new vulnerabilities in networking software and the persistence of some intruders.

Some of the problems with Internet security are a result of inherent vulnerabilities in the services (and the protocols that the services implement), while others are a result of host configuration and access controls that are poorly implemented or overly complex to administer. Additionally, the role and importance of system management is often short-changed in job descriptions, resulting in many administrators being, at best, part-time and poorly prepared. This is further aggravated by the tremendous growth of the Internet and how the Internet is used; businesses and agencies now depend on the Internet (often more than they realize) for communications and research and thus have much more to lose if their sites are attacked. The following sections describe problems on the Internet and factors that contribute to these problems.

1.3.1 Security Incidents on the Internet

As evidence of the above, three problems have occurred within months of each other. In the first, persistent vulnerabilities in the UNIX *sendmail*[2] program were discussed openly on Internet discussion lists. Sites that had not corrected their sendmail programs were

[2] *sendmail* is the mail transport software for most UNIX hosts. It is a very large, complex program that has been found repeatedly to contain vulnerabilities that have permitted intruder access to systems that run sendmail.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

forced to scramble to correct the programs before attackers used the vulnerabilities to attack the sites. However, due to the complexity of the sendmail program and networking software in general, three subsequent versions of sendmail were found to still contain significant vulnerabilities [CIAC94a]. The sendmail program is used widely, and sites without firewalls to limit access to sendmail are forced to react quickly whenever problems are found and vulnerabilities revealed.

In the second, a version of a popular and free FTP server was found to contain a Trojan Horse that permitted privileged access to the server. Sites using this FTP server, but not necessarily the contaminated version, were again forced to react very carefully and quickly to this situation [CIAC94c]. Many sites rely on the wealth of free software available on the Internet, especially security-related software that adds capability for logging, access control, and integrity checking that vendors often do not provide as part of the operating system. While the software is often high quality, sites may have little recourse other than to rely on the authors of the software if it is found to have vulnerabilities and other problems.[3]

The third problem has the strongest implications: [CERT94] and [CIAC94b] reported that intruders had broken into potentially *thousands* of systems throughout the Internet, including gateways between major networks, and installed *sniffer* programs to monitor network traffic for usernames and static passwords typed in by users to connect to networked systems. The intruders had used various known techniques for breaking into systems, as well as using passwords that had been "sniffed." One of the implications of this incident is that *static or reusable passwords are obsolete* for protecting access to user accounts. In fact, a user connecting to a remote system across the Internet may be unintentionally placing that system at risk of attack by intruders who could be monitoring the network traffic to the remote system.

The following sections go into more detail on problems with Internet security. [Garf92], [Cur92], [Bel89], [Ches94], and [Farm93] all provide more background and detail; readers are encouraged to consult these references.

1.3.2 Weak Authentication

Incident handling teams estimate that many incidents stem from use of weak, static passwords. Passwords on the Internet can be "cracked" a number of different ways, however the two most common methods are by cracking the encrypted form of the password and by monitoring communications channels for password packets. The UNIX operating system usually stores an encrypted form of passwords in a file that can be read by normal

[3]It should be pointed out that even vendor-supported software has such problems and may be even harder to get fixed in a timely fashion.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

users. The password file can be obtained by simply copying it or via a number of other intruder methods. Once the file is obtained, an intruder can run readily-available password cracking programs against the passwords. If the passwords are weak, e.g., less that 8 characters, English words, etc., they could be cracked and used to gain access into the system.

Another problem with authentication results from some TCP or UDP services being able to authenticate only to the granularity of host addresses and not to specific users. For example, an NFS (UDP) server cannot grant access to a specific user on a host, it must grant access to the entire host. The administrator of a server may trust a specific user on a host and wish to grant access to that user, but the administrator has no control over other users on that host and is thus forced to grant access to all users (or grant no access at all).

1.3.3 Ease of Spying/Monitoring

It is important to note that when a user connects to her account on a remote host using TELNET or FTP, the user's password travels across the Internet unencrypted, or in *plaintext*. Thus, another method for breaking into systems is to monitor connections for IP packets bearing a username and password, and then using them on the system to login normally. If the captured password is to an administrator account, then the job of obtaining privileged access is made much easier. As noted previously, hundreds and possibly thousands of systems across the Internet have been penetrated as a result of monitoring for usernames and passwords.

Electronic mail, as well as the contents of TELNET and FTP sessions, can be monitored and used to learn information about a site and its business transactions. Most users do not encrypt e-mail, yet many assume that e-mail is secure and thus safe for transmitting sensitive information.

The X Window System is an increasingly popular service that is also vulnerable to spying and monitoring. X permits multiple windows to be opened at a workstation, along with display of graphics and multi-media applications (for example, the WWW browser *Mosaic*). Intruders can sometimes open windows on other systems and read keystrokes that could contain passwords or sensitive information.

1.3.4 Ease of Spoofing

As noted in section 1.2.1, the IP address of a host is presumed to be valid and is therefore trusted by TCP and UDP services. A problem is that, using IP source routing, an attacker's host can masquerade as a trusted host or client. Briefly, IP source routing is

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

an option that can be used to specify a direct route to a destination and return path back to the origination. The route can involve the use of other routers or hosts that normally would not be used to forward packets to the destination. An example of how this can be used such that an attacker's system could masquerade as the trusted client of a particular server is as follows:

1. the attacker would change her host's IP address to match that of the trusted client,
2. the attacker would then construct a source route to the server that specifies the direct path the IP packets should take to the server and should take from the server back to the attacker's host, using the trusted client as the last *hop* in the route to the server,
3. the attacker sends a client request to the server using the source route,
4. the server accepts the client request as if it came directly from the trusted client and returns a reply to the trusted client,
5. the trusted client, using the source route, forwards the packet on to the attacker's host.

Many UNIX hosts accept source routed packets and will pass them on as the source route indicates. Many routers will accept source routed packets as well, whereas some routers can be configured to block source routed packets.

An even simpler method for spoofing a client is to wait until the client system is turned off and then impersonate the client's system. In many organizations, staff members use personal computers and TCP/IP network software to connect to and utilize UNIX hosts as a local area network server. The personal computers often use NFS to obtain access to server directories and files (NFS uses IP addresses only to authenticate clients). An attacker could, after hours, configure a personal computer with the same name and IP address as another's, and then initiate connections to the UNIX host as if it were the "real" client. This is very simple to accomplish and likely would be an insider attack.

Electronic mail on the Internet is particularly easy to spoof and, without enhancements such as digital signatures[NIST94a], generally cannot be trusted. As a brief example, consider the exchange that takes place when Internet hosts exchange mail. The exchange takes place using a simple protocol consisting of ASCII-character commands. An intruder easily could enter these commands by hand by using TELNET to connect directly to a system's Simple Mail Transfer Protocol (SMTP) port. The receiving host trusts that the sending host is who it says it is, thus the origin of the mail can be spoofed easily by entering a sender address that is different from the true address. As a result, any user, without privileges, can falsify or spoof e-mail.

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

Other services, such as Domain Name Service, can be spoofed, but with more difficulty than electronic mail. These services still represent a threat that needs to be considered when using them.

1.3.5 Flawed LAN Services and Mutually Trusting Hosts

Host systems are difficult and time consuming to manage securely. To ease management demands and to enhance local area networking, some sites use services such as Network Information Services (NIS) and Network File System (NFS). These services can greatly reduce the amount of redundant management by permitting certain databases such as the password files to be managed in a distributed manner and by permitting systems to share files and data. Ironically, these services are inherently insecure and can be exploited to gain access by knowledgeable intruders. If a central server system is compromised, then the other systems trusting the central system could be compromised rather easily.

Some services such as *rlogin* allow for hosts to "trust" each other for the purposes of user convenience and enhanced sharing of systems and devices. If a system is penetrated or spoofed, and that system is trusted by other systems, it is simple for the intruder to then gain access to the other systems. As an example, a user with an account on more than one system can eliminate the need to enter a password at every system by configuring the accounts to trust connections from the user's primary system. When the user uses the *rlogin* command to connect to a host, the destination system will not ask for a password or account name, and the user's connection will be accepted. While this has a positive aspect in that the user's password does not get transmitted and could not be monitored and captured, it has a negative aspect in that if the user's primary account were to be penetrated, the intruder could simply use *rlogin* to penetrate the accounts on the other systems. For this reason, use of mutually-trusting hosts is discouraged [Bel89], [Ches94].

1.3.6 Complex Configuration and Controls

Host system access controls are often complex to configure and test for correctness. As a result, controls that are accidentally misconfigured can result in intruders gaining access. Some major UNIX vendors still ship host systems with access controls configured for maximum (i.e., least secure) access, which can result in unauthorized access if left as is.

A number of security incidents have occurred on the Internet due in part to vulnerabilities discovered by intruders (and subsequently, users, incident response teams, and vendors). Since most modern variants of UNIX derive their networking code from the BSD releases, and since the source code to the BSD releases is widely available, intruders have been able to study the code for bugs and conditions that can be exploited to gain access

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls

to systems. The bugs exist in part because of the complexity of the software and the inability to test it in all the environments in which it must operate. Sometimes the bugs are easily discovered and corrected, other times little can be done except to rewrite the application, which is usually the option of last resort (the sendmail program may be an example of the latter).

1.3.7 Host-based Security Does Not Scale

Host-based security does not scale well: as the number of hosts at a site increases, the ability to ensure that security is at a high level for each host decreases. Given that secure management of just one system can be demanding, managing many such systems could easily result in mistakes and omissions. A contributing factor is that the role of system management is often short-changed and performed in haste. As a result, some systems will be less secure than other systems, and these systems could be the weak links that ultimately will "break" the overall security chain.

If a vulnerability is discovered in networking software, a site that is not protected by a firewall needs to correct the vulnerability on all exposed systems as quickly as possible. As discussed in section 1.3.2, some vulnerabilities have permitted easy access to the UNIX root account; a site with many UNIX hosts would be particularly at risk to intruders in such a situation. Patching vulnerabilities on many systems in a short amount of time may not be practical and, if different versions of the operating system are in use, may not be possible. Such a site would be a "sitting duck" to intruder activity.

1.4 How Vulnerable Are Internet Sites?

As noted in the preceding sections, a number of the TCP and UDP services provide poor levels of security in today's Internet environment. With millions of users connected to the Internet, and governments and industry placing more reliance on Internet availability, the flaws in these services, as well as the availability of source code and tools to automate breaking into systems, can be devastating to sites that suffer break-ins. However, it is difficult to know or assess the true risks of using the Internet and, following, how vulnerable a site is to some form of attack from intruders and related activity. There are no firm statistics.

The Computer Emergency Response Team Coordination Center (CERT/CC) has maintained some base statistics on the number of incidents they have handled since their inception in 1988. The numbers have climbed quite steeply as each year has progressed, however at the same time, the Internet has also grown dramatically. In some cases, CERT counts multiple break-ins of the same pattern as all part of a single incident, thus a sin-

Documents
Document Sharing Hub
Full document details: Keeping Your Site Comfortably Secure: An Introduction to Internet Firewalls