Internet Firewalls and Security

A Technology Overview

Internet Firewalls and Security
A Technology Overview

Contents

Internet Firewalls and Security

A Technology Overview

By Chuck Semeria

Chuck Semeria has worked for 3Com for the past six years. In his position as a marketing engineer in the network systems division, he develops classroom and independent study courses for the education services department in the customer services organization.

Prior to joining 3Com, Chuck was the senior course developer and instructor for Adept, a robotics and vision systems company. Before that, he taught mathematics and computer science in California high schools and junior colleges. Chuck is a graduate of the University of California at Davis.

Security has become one of the primary concerns when an organization connects its private network to the Internet. Regardless of the business, an increasing number of users on private networks are demanding access to Internet services such as the World Wide Web (WWW), Internet mail, Telnet, and File Transfer Protocol (FTP). In addition, corporations want to offer WWW home pages and FTP servers for public access on the Internet.

Network administrators have increasing concerns about the security of their networks when they expose their organization's private data and networking infrastructure to Internet crackers. To provide the required level of protection, an organization needs a security policy to prevent unauthorized users from accessing resources on the private network and to protect against the unauthorized export of private information. Even if an organization is not connected to the Internet, it may still want to establish an internal security policy to manage user access to portions of the network and protect sensitive or secret information.

Internet Firewalls

An Internet firewall is a system or group of systems that enforces a security policy between an organization's network and the Internet. The firewall determines which inside services may be accessed from the outside, which outsiders are permitted access to the permitted inside services, and which outside services may be accessed by insiders. For a firewall to be effective, all traffic to and from the Internet must pass through the firewall, where it can be inspected (Figure 1). The firewall must permit only authorized traffic to pass, and the firewall itself must be immune to penetration. Unfortunately, a firewall system cannot offer any protection once an attacker has gotten through or around the firewall.

It is important to note that an Internet firewall is not just a router, a bastion host, or a combination of devices that provides security for a network. The firewall is part of an overall security policy that creates a perimeter defense designed to protect the information resources of the organization. This security policy must include published security guidelines to inform users of their responsibilities; corporate policies defining network access, service access, local and remote user authentication, dial-in and dial-out, disk and data encryption, and virus protection measures; and employee training. All potential points of network attack must be protected with the same level of network security. Setting up an Internet firewall without a comprehensive security policy is like placing a steel door on a tent.

Benefits of an Internet Firewall

Internet firewalls manage access between the Internet and an organization's private network (Figure 2). Without a firewall, each host system on the private network is exposed to attacks from other hosts on the Internet. This means that the security of the private network

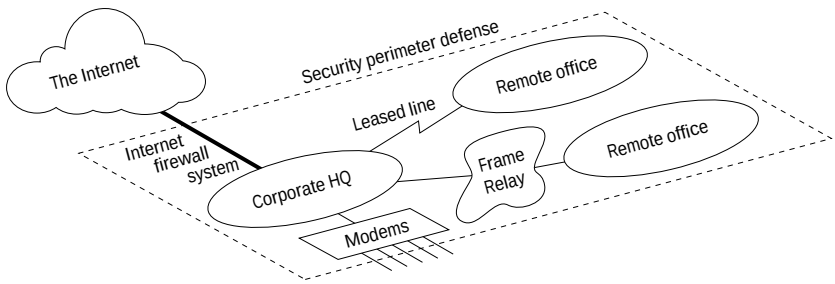

Figure 1. *Security Policy Creates a Perimeter Defense*

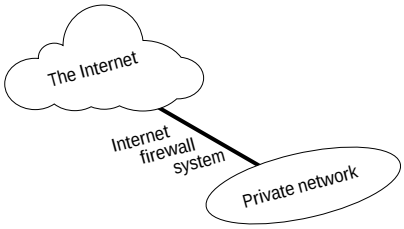

- Concentrates network security
- Serves as centralized access "choke point"
- Generates security alarms
- Monitors and logs Internet usage
- Good location for Network Address Translator (NAT)
- Good location for WWW and FTP servers

Figure 2. *Benefits of an Internet Firewall*

would depend on the "hardness" of each host's security features and would be only as secure as the weakest system.

Internet firewalls allow the network administrator to define a centralized "choke point" that keeps unauthorized users such as hackers, crackers, vandals, and spies out of the protected network; prohibits potentially vulnerable services from entering or leaving the protected network; and provides protection from various types of routing attacks. An Internet firewall simplifies security management, since network security is consolidated on the firewall systems rather than being distributed to every host in the entire private network.

Firewalls offer a convenient point where Internet security can be monitored and alarms generated. It should be noted that for organizations that have connections to the Internet, the question is not whether but when attacks will occur. Network administrators must audit and log all significant traffic through the firewall. If the network administrator doesn't take the time to respond to each alarm and examine logs on a regular basis, there is no need for the firewall, since the network administrator will never know if the firewall has been successfully attacked!

For the past few years, the Internet has been experiencing an address space crisis that has made registered IP addresses a less plentiful resource. This means that organizations wanting to connect to the Internet may not be able to obtain enough registered IP addresses to meet the demands of their user population. An Internet firewall is a logical place to deploy a Network Address Translator (NAT) that can help alleviate the address space shortage and eliminate the need to renumber when an organization changes Internet service providers (ISPs).

An Internet firewall is the perfect point to audit or log Internet usage. This permits the network administrator to justify the expense of the Internet connection to management, pinpoint potential bandwidth bottlenecks, and provide a method for departmental charge-backs if this fits the organization's financial model.

An Internet firewall can also offer a central point of contact for information delivery service to customers. The Internet firewall is the ideal location for deploying World Wide Web and FTP servers. The firewall can be configured to allow Internet access to these services, while prohibiting external access to other systems on the protected network.

Finally, some might argue that the deployment of an Internet firewall creates a single point of failure. It should be emphasized that if the connection to the Internet fails, the organization's private network will still continue to operate—only Internet access is lost. If there are multiple points of access, each one becomes a potential point of attack that the network administrator must firewall and monitor regularly.

Limitations of an Internet Firewall

An Internet firewall cannot protect against attacks that do not go through the firewall. For example, if unrestricted dial-out is permitted from inside the protected network, internal users can make a direct SLIP or PPP connection to the Internet. Savvy users who become irritated with the additional authentication required by firewall proxy servers may be tempted to circumvent the security system

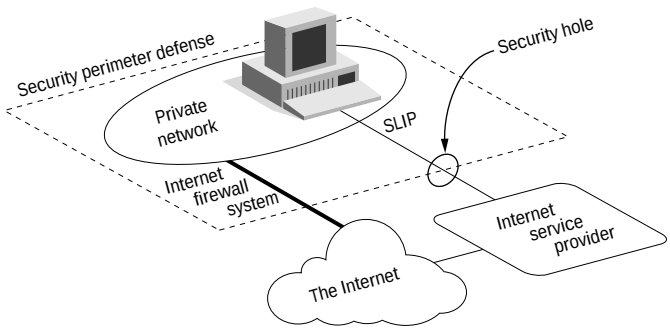

Figure 3. *A Connection Circumventing an Internet Firewall*

Glossary

Back door
A security hole in a compromised system that allows continued access to the system by an intruder even if the original attack is discovered.

Bastion host
A designated Internet firewall system specifically armored and protected against attacks.

Circuit-level gateway
A specialized function that relays TCP connections without performing any additional packet processing or filtering.

Internet firewall
A system or group of systems that enforces an access control policy between an organization's network and the Internet.

Packet filtering
A feature that allows a router to make a permit/deny decision for each packet based on the packet header information that is made available to the IP forwarding process.

Proxy service
Special-purpose, application-level code installed on an Internet firewall gateway. The proxy service allows the network administrator to permit or deny specific applications or specific features of an application.

Trojan horse
A packet sniffer that hides its sniffing activity. These packet sniffers can collect account names and passwords for Internet services, allowing a hacker to gain unauthorized access to other machines.

by purchasing a direct SLIP or PPP connection to an ISP. Since these types of connections bypass the security provided by the most carefully constructed firewall, they create a significant potential for back-door attacks (Figure 3). Users must be made aware that these types of connections are not permitted as part of the organization's overall security architecture.

Internet firewalls cannot protect against the types of threats posed by traitors or unwitting users. Firewalls do not prohibit traitors or corporate spies from copying sensitive data onto floppy disks or PCMCIA cards and removing them from a building. Firewalls do not protect against attacks where a hacker, pretending to be a supervisor or a befuddled new employee, persuades a less sophisticated user into revealing a password or granting them "temporary" network access. Employees must be educated about the various types of attacks and about the need to guard and periodically change their passwords.

Internet firewalls cannot protect against the transfer of virus-infected software or files. Since there are so many different viruses, operating systems, and ways of encoding and compressing binary files, an Internet firewall cannot be expected to accurately scan each and every file for potential viruses. Concerned organizations should deploy anti-viral software at each desktop to protect against their arrival from floppy disks or any other source.

Finally, Internet firewalls cannot protect against data-driven attacks. A data-driven attack occurs when seemingly harmless data is mailed or copied to an internal host and is executed to launch an attack. For example, a data-driven attack could cause a host to modify security-related files, making it easier for an intruder to gain access to the system. As we will see, the deployment of proxy servers on a bastion host is an excellent means of prohibiting direct connections from the outside and reducing the threat of data-driven attacks.

The Hacker's Toolbox

It is difficult to describe a typical hacker attack because intruders have different levels of technical expertise and many different motivations. Some hackers are intrigued by the challenge, others just want to make life more difficult for others, and still others are out to steal sensitive data for profit.

Information Gathering

Generally, the first step in a break-in is some form of information gathering. The goal is to construct a database of the target organization's network and gather information about the hosts residing on each of the networks. There are a number of tools that a hacker can use to collect this information:

- The SNMP protocol can be used to examine the routing table of an unsecured router to learn intimate details about the target organization's network topology.
- The TraceRoute program can reveal intermediate network numbers and routers in the path to a specific host.
- The Whois protocol is an information service that can provide data about all DNS domains and the system administrators responsible for each domain. However, this information is usually out of date.

- DNS servers can access a list of host IP addresses and their corresponding host names.
- The Finger protocol can reveal detailed information about the users (login names, phone numbers, time they last logged in, etc.) of a specified host.
- The Ping program can be employed to locate a particular host and determine its reachability. This simple tool can be used in a short scanning program that pings every possible host address on a network to construct a list of the hosts actually residing on the network.

Probing Systems for Security Weaknesses

After information about the targeted organization's network is gathered, the hacker attempts to probe each host for security weaknesses. There are a number of tools that a hacker can use to automatically scan the individual hosts residing on a network; for example:

- Since the list of known service vulnerabilities is rather short, a knowledgeable hacker can write a small program that attempts to connect to specific service ports on a targeted host. The output of the program is a list of hosts that support services that are exposed to attack.
- There are several publicly available tools, such as the Internet Security Scanner (ISS) or the Security Analysis Tool for Auditing Networks (SATAN), that scan an entire domain or subnetwork and look for security holes. These programs determine the weaknesses of each system with respect to several common system vulnerabilities. Intruders use the information collected from these scans to gain unauthorized access to the targeted organization's systems.

A clever network administrator can use these tools within their private network to discover potential security weaknesses and determine which hosts need to be updated with new software patches.

Accessing Protected Systems

The intruder uses the results of the host probes to target a specific system for attack. After gaining access to a protected system, the hacker has many options available:

- The intruder can attempt to destroy evidence of the assault and open new security holes or back doors in the compromised system in order to have continued access even if the original attack is discovered.
- The intruder can install packet sniffers that include Trojan horse binaries that hide the sniffing activity on the installed systems. The packet sniffers collect account names and passwords for Telnet and FTP services that allow the hacker to spread the attack to other machines.
- The intruder can find other hosts that trust the compromised system. This allows the hacker to exploit the vulnerabilities of a single host and spread the attack across the entire organization's network.
- If the hacker can obtain privileged access on a compromised system, he or she can read mail, search private files, steal private files, and destroy or corrupt important data.

Basic Firewall Design Decisions

When designing an Internet firewall, there are a number of decisions that must be addressed by the network administrator:

- The stance of the firewall
- The overall security policy of the organization
- The financial cost of the firewall
- The components or building blocks of the firewall system

Stance of the Firewall

The stance of a firewall system describes the fundamental security philosophy of the organization. An Internet firewall may take one of two diametrically opposed stances:

- *Everything not specifically permitted is denied.* This stance assumes that a firewall should block all traffic, and that each desired service or application should be implemented on a case-by-case basis. This is the recommended approach. It creates a very secure environment, since only carefully selected services are supported. The disadvantage is that it places security ahead of ease of use, limiting the number of options available to the user community.

Acronyms

CERT
Computer Emergency Response Team

DNS
Domain Name Service

FAQ
Frequently Asked Questions

FTP
File Transfer Protocol

ICMP
Internet Control Message Protocol

ISP
Internet service provider

ISS
Internet Security Scanner

NAT
Network Address Translator

PCMCIA
Personal Computer Memory Card International Association

PPP
Point-to-Point Protocol

RFC
Request for Comment

SATAN
Security Analysis Tool for Auditing Networks

SLIP
Serial Line Internet Protocol

SMTP
Simple Mail Transfer Protocol

TCP
Transmission Control Protocol

UDP
User Datagram Protocol

Learning More About Internet Attacks

For the latest, up-to-date information concerning attacks on Internet sites, contact the Computer Emergency Response Team (CERT) Coordination Center. CERT periodically publishes warnings and summaries to draw attention to the various types of attacks that have been reported to their incident response staff. These reports also contain information and solutions for defeating each type of attack. New or updated files are available for anonymous FTP from ftp://info.cert.org, and past summaries are available from ftp://info.cert.org/pub/cert_summaries.

For more information concerning the techniques employed by hackers, track the following USENET newsgroups: comp.security.announce, comp.security.mis, comp.security.unix, alt.2600 , alt.wired, alt.hackers, and alt.security. *Finally, look for various hacker bulletin boards—they're everywhere!*

- *Everything not specifically denied is permitted.* This stance assumes that a firewall should forward all traffic, and that each potentially harmful service should be shut off on a case-by-case basis. This approach creates a more flexible environment, with more services available to the user community. The disadvantage is that it puts ease of use ahead of security, putting the network administrator in a reactive mode and making it increasingly difficult to provide security as the size of the protected network grows.

Security Policy of the Organization

As discussed earlier, an Internet firewall does not stand alone—it is part of the organization's overall security policy, which defines all aspects of its perimeter defense. To be successful, organizations must know what they are protecting. The security policy must be based on a carefully conducted security analysis, risk assessment, and business needs analysis. If an organization does not have a detailed security policy, the most carefully crafted firewall can be circumvented to expose the entire private network to attack.

Cost of the Firewall

How much security can the organization afford? A simple packet-filtering firewall can have a minimal cost since the organization needs a router to connect to the Internet, and packet filtering is included as part of the standard router feature set. A commercial firewall system provides increased security but may cost from U.S.$4,000 to $30,000, depending on its complexity and the number of systems protected. If an organization has the in-house expertise, a home-brewed firewall can be constructed from public domain software, but there are still costs in terms of the time to develop and deploy the firewall system. Finally, all firewalls require continuing support for administration, general maintenance, software updates, security patches, and incident handling.

Components of the Firewall System

After making decisions about firewall stance, security policy, and budget issues, the organization can determine the specific components of its firewall system. A typical firewall is composed of one or more of the following building blocks:

- Packet-filtering router
- Application-level gateway (or proxy server)
- Circuit-level gateway

The remainder of this paper discusses each of these building blocks and describes how they can work together to build an effective Internet firewall system.

Building Blocks: Packet-Filtering Routers

A packet-filtering router (Figure 4) makes a permit/deny decision for each packet that it receives. The router examines each datagram to determine whether it matches one of its packet-filtering rules. The filtering rules are based on the packet header information that is made available to the IP forwarding process. This information consists of the IP source address, the IP destination address, the encapsulated protocol (TCP, UDP, ICMP, or IP Tunnel), the TCP/UDP source port, the TCP/UDP destination port, the ICMP message type, the incoming interface of the packet, and the outgoing interface of the packet. If a match is found and the rule permits the packet, the packet is forwarded according to the information in the routing table. If a match is found and the rule denies the packet, the packet is

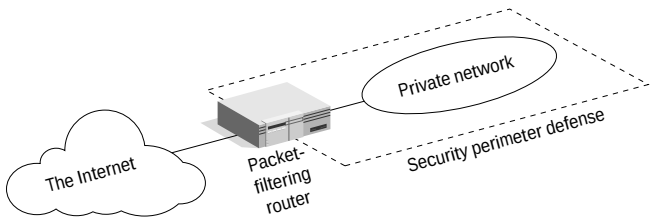

Figure 4. *Packet-Filtering Router*

discarded. If there is no matching rule, a user-configurable default parameter determines whether the packet is forwarded or discarded.

Service-Dependent Filtering

The packet-filtering rules allow a router to permit or deny traffic based on a specific service, since most service listeners reside on well-known TCP/UDP port numbers. For example, a Telnet server listens for remote connections on TCP port 23 and an SMTP server listens for incoming connections on TCP port 25. To block all incoming Telnet connections, the router simply discards all packets that contain a TCP destination port value equal to 23. To restrict incoming Telnet connections to a limited number of internal hosts, the router must deny all packets that contain a TCP destination port value equal to 23 *and* that do not contain the destination IP address of one of the permitted hosts.

Some typical filtering rules include:

- Permit incoming Telnet sessions only to a specific list of internal hosts
- Permit incoming FTP sessions only to specific internal hosts
- Permit all outbound Telnet sessions
- Permit all outbound FTP sessions
- Deny all incoming traffic from specific external networks

Service-Independent Filtering

There are certain types of attacks that are difficult to identify using basic packet header information because the attacks are service independent. Routers can be configured to protect against these types of attacks, but they are more difficult to specify since the filtering rules require additional information that can be learned only by examining the routing table, inspecting for specific IP options, checking for a special fragment offset, and so on. Examples of these types of attacks include:

Source IP Address Spoofing Attacks. For this type of attack, the intruder transmits packets from the outside that pretend to originate from an internal host: the packets falsely contain the source IP address of an inside system. The attacker hopes that the use of a spoofed source IP address will allow penetration of systems that employ simple source address security where packets from specific trusted internal hosts are accepted and packets from other hosts are discarded. Source spoofing attacks can be defeated by discarding each packet with an inside source IP address if the packet arrives on one of the router's outside interfaces.

Source Routing Attacks. In a source routing attack, the source station specifies the route that a packet should take as it crosses the Internet. This type of attack is designed to bypass security measures and cause the packet to follow an unexpected path to its destination. A source routing attack can be defeated by simply discarding all packets that contain the source route option.

Tiny Fragment Attacks. For this type of attack, the intruder uses the IP fragmentation feature to create extremely small fragments and force the TCP header information into a separate packet fragment. Tiny fragment attacks are designed to circumvent user-defined filtering rules; the hacker hopes that a filtering router will examine only the first fragment and allows all other fragments to pass. A tiny fragment attack can be defeated by discarding all packets where the protocol type is TCP and the IP FragmentOffset is equal to 1.

Benefits of Packet-Filtering Routers

The majority of Internet firewall systems are deployed using only a packet-filtering router. Other than the time spent planning the filters and configuring the router, there is little or no cost for implementing packet filtering since the feature is included as part of standard router software releases. Since Internet access is generally provided over a WAN interface, there is little impact on router performance if traffic loads are moderate and few filters are defined. Finally, a packet-filtering router is generally transparent to users and applications, so it does not require specialized user training or that specific software be installed on each host.

Limitations of Packet-Filtering Routers

Defining packet filters can be a complex task because network administrators need to have a detailed understanding of the various Internet services, packet header formats, and the specific values they expect to find in each field. If complex filtering requirements must be supported, the filtering rule set can become very long and complicated, making it difficult to manage and comprehend. Finally, there are few testing facilities to verify the correctness of the filtering rules after they are configured on the router. This can potentially leave a site open to untested vulnerabilities.

Any packet that passes directly through a router could potentially be used launch a data-driven attack. Recall that a data-driven attack occurs when seemingly harmless data is forwarded by the router to an internal host. The data contains hidden instructions that cause the host to modify access control and security-related files, making it easier for the intruder to gain access to the system.

Generally, the packet throughput of a router decreases as the number of filters increases. Routers are optimized to extract the destination IP address from each packet, make a relatively simple routing table lookup, and then forward the packet to the proper interface for transmission. If filtering is enabled, the router must not only make a forwarding decision for each packet, but also apply all of the filter rules to each packet. This can consume CPU cycles and impact the performance of a system.

IP packet filters may not be able to provide enough control over traffic. A packet-filtering router can permit or deny a particular service, but it is not capable of understanding the context/data of a particular service. For example, a network administrator may need to filter traffic at the application layer in order to limit access to a subset of the available FTP or Telnet commands, or to block the import of mail or newsgroups concerning specific topics. This type of control is best performed at a higher layer by proxy services and application-level gateways.

Building Blocks: Application-Level Gateways

An application-level gateway allows the network administrator to implement a much stricter security policy than with a packet-filtering router. Rather than relying on a generic packet-filtering tool to manage the flow of Internet services through the firewall, special-purpose code (a proxy service) is installed on the gateway for each desired application. If the network administrator does not install the proxy code for a particular application, the service is not supported and cannot be forwarded across the firewall. Also, the proxy code can be configured to support only those specific features of an application that the network administrator considers acceptable while denying all other features.

This enhanced security comes with an increased cost in terms of purchasing the gateway hardware platform, the proxy service applications, the time and knowledge required to configure the gateway, a decrease in the level of service that may be provided to users, and a lack of transparency resulting in a less user-friendly system. As always, the network administrator is required to balance the organization's need for security with the user community's demand for ease of use.

It is important to note that users are permitted access to the proxy services, but they are *never* permitted to log in to the application-level gateway. If users are permitted to log in to the firewall system, the security of the firewall is threatened, since an intruder could potentially perform some activity that compromises the effectiveness of the firewall. For example, the intruder could gain root access, install Trojan horses to collect passwords, and modify the security configuration files of the firewall.

Bastion Host

Unlike packet-filtering routers, which allow the direct flow of packets between inside systems and outside systems, application-level gateways allow information to flow between systems but do not allow the direct exchange of packets. The chief risk of allowing packets to be exchanged between inside systems and

outside systems is that the host applications residing on the protected network's systems must be secured against any threat posed by the allowed services.

An application-level gateway is often referred to as a "bastion host" because it is a designated system that is specifically armored and protected against attacks. Several design features are used to provide security for a bastion host:

- The bastion host hardware platform executes a "secure" version of its operating system. For example, if the bastion host is a UNIX® platform, it executes a secure version of the UNIX operating system that is specifically designed to protect against operating system vulnerabilities and ensure firewall integrity.
- Only the services that the network administrator considers essential are installed on the bastion host. The reasoning is that if a service is not installed, it can't be attacked. Generally, a limited set of proxy applications such as Telnet, DNS, FTP, SMTP, and user authentication are installed on a bastion host.
- The bastion host may require additional authentication before a user is allowed access to the proxy services. For example, the bastion host is the ideal location for installing strong authentication using a one-time password technology where a smart card cryptographic authenticator generates a unique access code. In addition, each proxy service may require its own authentication before granting user access.
- Each proxy is configured to support only a subset of the standard application's command set. If a standard command is not supported by the proxy application, it is simply not available to the authenticated user.
- Each proxy is configured to allow access only to specific host systems. This means that the limited command/feature set may be applied only to a subset of systems on the protected network.
- Each proxy maintains detailed audit information by logging all traffic, each connection, and the duration of each connection. The audit log is an essential tool for discovering and terminating intruder attacks.
- Each proxy is a small and uncomplicated program specifically designed for network security. This allows the source code of the proxy application to be reviewed and checked for potential bugs and security holes. For example, a typical UNIX mail application may contain over 20,000 lines of code, while a mail proxy may contain fewer than 1000!
- Each proxy is independent of all other proxies on the bastion host. If there is a problem with the operation of any proxy, or if a future vulnerability is discovered, it can be uninstalled without affecting the operation of the other proxy applications. Also, if the user population requires support for a new service, the network administrator can easily install the required proxy on the bastion host.
- A proxy generally performs no disk access other than to read its initial configuration file. This makes it difficult for an intruder to install Trojan horse sniffers or other dangerous files on the bastion host.
- Each proxy runs as a nonprivileged user in a private and secured directory on the bastion host.

Example: Telnet Proxy

Figure 5 on page 10 illustrates the operation of a Telnet proxy on an bastion host. For this example, the outside client wants to Telnet to an inside server protected by the application-level gateway.

The Telnet proxy never allows the remote user to log in or have direct access to the internal server. The outside client Telnets to the bastion host, which authenticates the user employing one-time password technology. After authentication, the outside client gains access to the user interface of the Telnet proxy. The Telnet proxy permits only a subset of the Telnet command set and determines which inside hosts are available for Telnet access. The outside user specifies the destination host and the Telnet proxy makes its own connection to the inside server and forwards commands to the inside server on behalf of the outside client. The outside client believes that the Telnet proxy is the real inside server, while

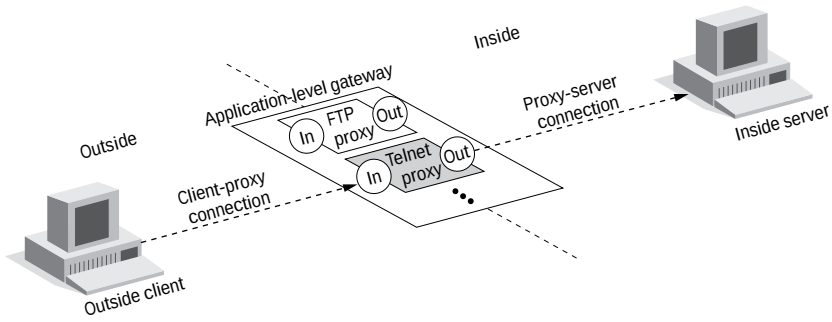

Figure 5. *Telnet Proxy*

the inside server believes that the Telnet proxy is the outside client.

Figure 6 shows the output to the outside client's terminal screen as the connection to the inside server is established. Note that the client is not performing a logon to the bastion host; the user is being authenticated by the bastion host and a challenge is issued before the user is permitted to communicate with the Telnet proxy. After passing the challenge, the proxy server limits the set of commands and destinations that are available to the outside client.

Authentication can be based on either something the user knows (like a password) or something the user physically possesses (like a smart card). Both techniques are subject to theft, but using a combination of both methods increases the likelihood of correct user authentication. In the Telnet example, the proxy transmits a challenge and the user, with the aid of a smart card, obtains a response to the challenge. Typically, a user unlocks the smart card by entering their PIN number and the card, based on a shared "secret" encryption key and its own internal clock, returns an encrypted value for the user to enter as a response to the challenge.

Benefits of Application-Level Gateways

There are many benefits to the deployment of application-level gateways. They give the network manager complete control over each

```
Outside-Client > telnet bastion_host
Username: John Smith
Challenge Number "237936"
Challenge Response: 723456
Trying 200.43.67.17 ...

HostOS UNIX (bastion_host)

bh-telnet-proxy> help
Valid commands are:

connect hostname
help/?
quit/exit

bh-telnet-proxy> connect inside_server

HostOS UNIX (inside_server)

login: John Smith
Password: ######
Last login: Wednesday April 15 11:17:15
```

Figure 6. *Telnet "Session" Terminal Display*

service, since the proxy application limits the command set and determines which internal hosts may be accessed by the service. Also, the network manager has complete control over which services are permitted, since the absence of a proxy for a particular service means that the service is completely blocked. Application-level gateways have the ability to support strong user authentication and provide detailed logging information. Finally, the filtering rules for an application-level gateway are much easier to configure and test than for a packet-filtering router.

Limitations of Application-Level Gateways

The greatest limitation of an application-level gateway is that it requires either that users modify their behavior, or that specialized software be installed on each system that accesses proxy services. For example, Telnet access via an application-level gateway requires two user steps to make the connection rather than a single step. However, specialized end-system software could make the application-level gateway transparent by allowing the user to specify the destination host rather than the application-level gateway in the Telnet command.

Building Blocks: Circuit-Level Gateways

A circuit-level gateway is a specialized function that can be performed by an application-level gateway. A circuit-level gateway simply relays TCP connections without performing any additional packet processing or filtering.

Figure 7 illustrates the operation of a typical Telnet connection through a circuit-level gateway. The circuit-level gateway simply relays the Telnet connection through the firewall but does no additional examination, filtering, or management of the Telnet protocol. The circuit-level gateway acts like a wire, copying bytes back and forth between the inside connection and the outside connection. However, because the connection appears to originate from the firewall system, it conceals information about the protected network.

Circuit-level gateways are often used for outgoing connections where the system administrator trusts the internal users. Their chief advantage is that a bastion host can be configured as a hybrid gateway supporting application-level or proxy services for inbound connections and circuit-level functions for outbound connections. This makes the firewall system easier to use for internal users who want direct access to Internet services, while still providing the firewall functions needed to protect the organization from external attack.

Firewall Example #1: Packet-Filtering Router

The most common Internet firewall system consists of nothing more than a packet-filtering router deployed between the private network and the Internet (Figure 8 on page 12). A packet-filtering router performs the typical routing functions of forwarding traffic between networks as well as using packet-filtering rules to permit or deny traffic. Typically, the filter rules are defined so that hosts on the private network have direct access to the Internet, while hosts on the Internet have limited access to systems on the private network. The external stance of this type of firewall system is usually that everything not specifically permitted is denied.

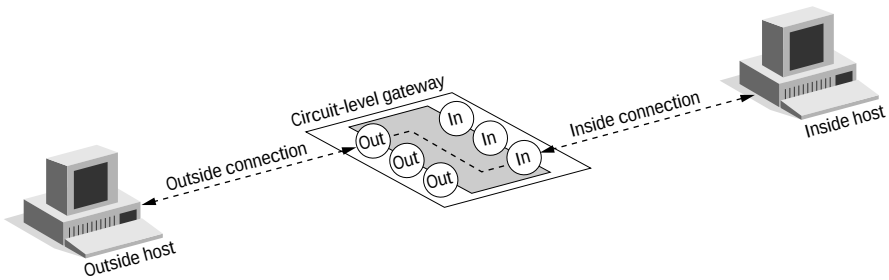

Figure 7. *Circuit-Level Gateway*

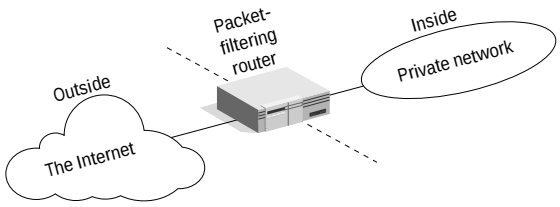

Figure 8. *Packet-Filtering Router Firewall*

Although this firewall system has the benefit of being inexpensive and transparent to users, it possesses all of the limitations of a packet-filtering router such as exposure to attacks from improperly configured filters and attacks that are tunneled over permitted services. Since the direct exchange of packets is permitted between outside systems and inside systems, the potential extent of an attack is determined by the total number of hosts and services to which the packet-filtering router permits traffic. This means that each host directly accessible from the Internet needs to support sophisticated user authentication and needs to be regularly examined by the network administrator for signs of an attack. Also, if the single packet-filtering router is penetrated, every system on the private network may be compromised.

Firewall Example #2: Screened Host Firewall

The second firewall example employs both a packet-filtering router and a bastion host (Figure 9). This firewall system provides a higher level of security than the previous example because it implements both network-layer security (packet-filtering) and application-layer security (proxy services). Also, an intruder has to penetrate two separate systems before the security of the private network can be compromised.

For this firewall system, the bastion host is configured on the private network with a packet-filtering router between the Internet and the bastion host. The filtering rules on the exposed router are configured so that outside systems can access only the bastion host; traffic addressed to all other internal systems is blocked. Since the inside hosts reside on the same network as the bastion host, the security policy of the organization determines whether inside systems are permitted direct access to the Internet, or whether they are required to use the proxy services on the bastion host. Inside users can be forced to use the proxy services by configuring the router's filter rules to accept only internal traffic originating from the bastion host.

One of the benefits of this firewall system is that a public information server providing Web and FTP services can be placed on the segment shared by the packet-filtering router and the bastion host. If the strongest security is required, the bastion host can run proxy services that require both internal and external users to access the bastion host before communicating with the information server. If a lower level of security is adequate, the router may be configured to allow outside users direct access to the public information server.

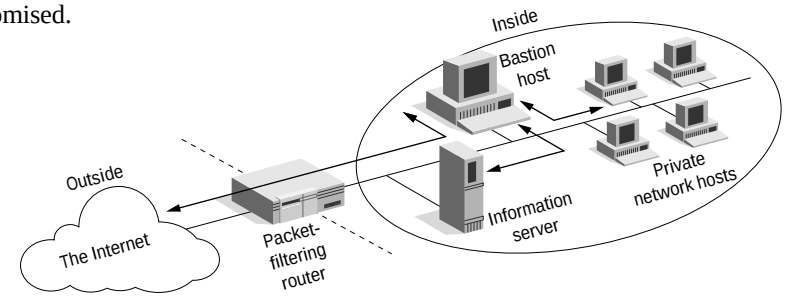

Figure 9. *Screened Host Firewall System (Single-Homed Bastion Host)*

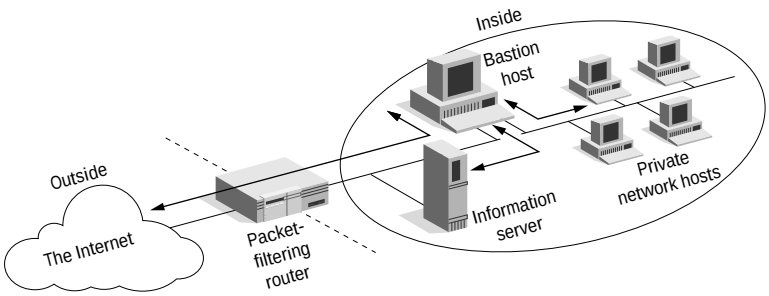

Figure 10. *Screened Host Firewall System (Dual-Homed Bastion Host)*

An even more secure firewall system can be constructed using a dual-homed bastion host system (Figure 10). A dual-homed bastion host has two network interfaces, but the host's ability to directly forward traffic between the two interfaces bypassing the proxy services is disabled. The physical topology forces all traffic destined for the private network through the bastion host and provides additional security if outside users are granted direct access to the information server.

Since the bastion host is the only internal system that can be directly accessed from the Internet, the potential set of systems open to attack is limited to the bastion host. However, if users are allowed to log on to the bastion host, the potential set of threatened systems expands to include the entire private network, since it is much easier for an intruder to compromise the bastion host if they are allowed to log on. It is critical that the bastion host be hardened and protected from penetration and that users never be allowed to log on to the bastion host.

Firewall Example #3: "Demilitarized Zone" or Screened-Subnet Firewall

The final firewall example employs two packet-filtering routers and a bastion host (Figure 11). This firewall system creates the most secure firewall system, since it supports both network- and application-layer security while defining a "demilitarized zone" (DMZ) network. The network administrator places the bastion host, information servers, modem pools, and other public servers on the DMZ network. The DMZ network functions as a small, isolated network positioned between the Internet and the private network. Typically, the DMZ is configured so that systems on the Internet and systems on the private network can access only a limited number of systems on the DMZ network, but the direct transmission of traffic across the DMZ network is prohibited.

For incoming traffic, the outside router protects against the standard external attacks (source IP address spoofing, source routing attacks, etc.) and manages Internet access to the DMZ network. It permits external systems

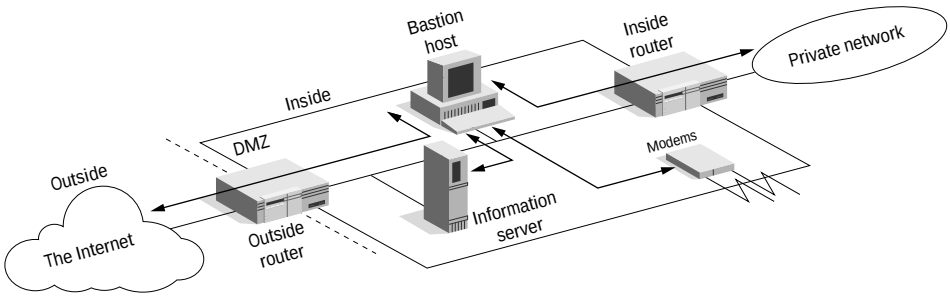

Figure 11. *Screened-Subnet Firewall System*

to access only the bastion host (and possibly the information server). The inside router provides a second line of defense, managing DMZ access to the private network by accepting only traffic originating from the bastion host.

For Internet-bound traffic, the inside router manages private network access to the DMZ network. It permits internal systems to access only the bastion host (and possibly the information server). The filtering rules on the outside router require use of the proxy services by accepting only Internet-bound traffic from the bastion host.

There are several key benefits to the deployment of a screened subnet firewall system:

- An intruder must crack three separate devices (without detection) to infiltrate the private network: the outside router, the bastion host, and the inside router.
- Since the outside router advertises the DMZ network only to the Internet, systems on the Internet do not have routes to the protected private network. This allows the network manager to ensure that the private network is "invisible," and that only selected systems on the DMZ are known to the Internet via routing table and DNS information exchanges.
- Since the inside router advertises the DMZ network only to the private network, systems on the private network do not have direct routes to the Internet. This guarantees that inside users must access the Internet via the proxy services residing on the bastion host.
- Packet-filtering routers direct traffic to specific systems on the DMZ network, eliminating the need for the bastion host to be dual-homed.
- The inside router supports greater packet throughput than a dual-homed bastion host when it functions as the final firewall system between the private network and the Internet.
- Since the DMZ network is a different network than the private network, a Network Address Translator (NAT) can be installed on the bastion host to eliminate the need to renumber or resubnet the private network.

Summary

There is no single correct answer for the design and deployment of Internet firewalls. Each organization's decision will be influenced by many different factors such as their corporate security policy, the technical background of their staff, cost, and the perceived threat of attack. This paper focused on many of the issues relating to the construction of Internet firewalls, including their benefits, limitations, building blocks, and examples of firewall system topologies. Since the benefits of connecting to the global Internet probably exceed its costs, network managers should proceed with an awareness of the dangers and an understanding that, with the proper precautions, their networks can be as safe as they need them to be. ❑

References

Textbooks

Building Internet Firewalls. D. Brent Chapman and Elizabeth Zwicky. O'Reilly & Associates, 1995.

Firewalls and Internet Security: Repelling the Wily Hacker. Bill Cheswick and Steve Bellovin. Addison-Wesley, 1994.

Practical UNIX Security. Simson Garfinkel and Gene Spafford. O'Reilly & Associates, 1991.

Requests for Comment

RFC 1244: *Site Security Handbook*. P. Holbrook and J. Reynolds, July 1991.

RFC 1636: *Report of IAB Workshop on Security in the Internet Architecture* (February 8–10, 1994). R. Braden, D. Clark, S. Crocker, and C. Huitema, June 1994.

RFC 1704: *On Internet Authentication*. N. Haller and R. Atkinson, October 1994.

RFC 1858: *Security Considerations for IP Fragment Filtering*. G. Ziemba, D. Reed, and P. Traina, October 1995.

Firewall and Security Papers

"Almost Everything You Ever Wanted to Know About Security (but were afraid to ask)." Maintained by Alec Muffett (http://www.cis.ohio-state.edu/hypertext/faq/usenet/ security-faq/faq.html).

"How to Set Up a Secure Anonymous FTP Site." Christopher Klaus, Internet Security Systems, Inc. (http://www.cis.ohio-state.edu/hypertext/faq/usenet/computer-security/anonymous-ftp-faq/faq.html).

"Internet Firewalls Frequently Asked Questions." Maintained by Marcus J. Ranum, Trusted Information Systems, Inc. (http://www.v-one.com/pubs/fw-faq/faq.htm).

"Thinking About Firewalls." Marcus J. Ranum, Trusted Information Systems, Inc.(http://www.telstra.com.au/pub/docs/security/ThinkingFirewalls/ThinkingFirewalls. html).

"A Toolkit and Methods for Internet Firewalls." Marcus J. Ramus and Frederick M. Avolio, Trusted Information Systems, Inc. (http://web1.cohesive.com/original/centri/usenix.htm).

"What If Your Machines Are Compromised by an Intruder." Christopher Klaus, Internet Security Systems, Inc. (http://www.cis.ohio-state.edu/hypertext/faq/usenet/ computer-security/compromise-faq/faq.html).

"The World Wide Web Security FAQ." Lincoln D. Stein (http://www-genome.wi. mit.edu/WWW/faqs/www-security-faq.html).

World Wide Web Index Pages to Security-Related Documents

http://lcweb.loc.gov/global/internet/security.html
Library of Congress page containing links to documents on computer security.

http://www.telstra.com.au/pub/docs/security/
Telstra page containing links to documents on computer security.

References (Continued)

http://mls.saic.com/docs.html
Science Applications International Corporation's (SAIC) page containing links to documents on computer security.

http://csrc.ncsl.nist.gov/first/resources/from-cd95/pap.htm
Forum of Incident Response and Security Teams (FIRST) page containing links to documents on network security.

http://web1.cohesive.com/original/centri/info.htm#applevel
Cohesive Systems' page containing links to documents on network security.

http://www.netsurf.com/nsf/v01/01/resource/firewall.html
General index page containing links to documents on firewalls.

http://burgau.inesc.pt/docs/security/firewall/index.html
General index page containing links to documents on firewalls.

http://burgau.inesc.pt/docs/security/IP-security/index.html
General index page containing links to documents on IP security.

ftp://ftp.uni-paderborn.de/doc/FAQ/comp.security.misc/
General index page containing links to security-related Frequently Asked Questions (FAQs).

3Com Corporation
P.O. Box 58145
5400 Bayfront Plaza
Santa Clara, CA
95052-8145
Phone: 800-NET-3Com
or 408-764-5000
Fax: 408-764-5001
World Wide Web:
http://www.3com.com

3Com ANZA
ANZA East
Phone: 61 2 9937 5000
Fax: 61 2 9956 6247
ANZA West
Phone: 61 3 9653 9515
Fax: 61 3 9653 9505

3Com Asia Limited
Beijing, China
Phone: 8610 8492568
Fax: 8610 8492789
Shanghai, China
Phone: 86 21 3740220
Ext. 6115
Fax: 86 21 3552079
Hong Kong
Phone: 852 2501 1111
Fax: 852 2537 1149
Indonesia
Phone: 6221 523 9181
Fax: 6221 523 9156
Korea
Phone: 82 2 319 4711
Fax: 82 2 319 4710
Malaysia
Phone: 60 3 732 7910
Fax: 60 3 732 7912
Singapore
Phone: 86 21 6374 0220
Ext. 6155
Fax: 86 21 6355 2079
Taiwan
Phone: 886 2 377 5850
Fax: 886 2 377 5860
Thailand
Phone: 622 231 8151 2
Fax: 622 231 8121

3Com Benelux B.V.
Belgium, Luxembourg
Phone: 32 2 716 4880
Fax: 32 2 716 4780
Netherlands
Phone: 31 030 6029700
Fax: 31 030 6029777

3Com Canada
Calgary
Phone: 403 265 3266
Fax: 403 265 3268
Montreal
Phone: 514 874 8008
Fax: 514 393 1249
Toronto
Phone: 416 498 3266
Fax: 416 498 1262
Vancouver
Phone: 604 434 3266
Fax: 604 434 3264

3Com European HQ
Phone: 44 1628 897000
Fax: 44 1628 897041

3Com France
Phone: 33 1 69 86 68 00
Fax: 33 1 69 07 11 54

3Com GmbH
Germany
Phone: 49 89 627320
Fax: 49 89 62732233
Berlin
Phone: 49 30 3498790
Fax: 49 30 34987999
Poland
Phone: 48 22 6451351
Fax: 48 22 6451352
Switzerland
Phone: 41 31 9984555
Fax: 41 31 9984550

3Com Ireland
Phone: 353 1 820 7077
Fax: 353 1 820 7107

3Com Japan
Phone: 81 3 3345 7251
Fax: 81 3 3345 7261

3Com Latin America
U.S. Headquarters
Phone: 408-764-6075
Fax: 408-764-5730
Argentina
Phone: 541 815 7164
Fax: 541 815 7165
Brazil
Phone: 55 11 546 0869
Fax: 55 11 246 6813
Chile
Phone: 562 633 9242
Fax: 562 633 8935
Colombia
Phone: 571 618 4584
Fax: 571 618 4534
Mexico
Phone: 525 520 7841
Fax: 525 520 7837

3Com Northern Latin America
Miami, Florida
Phone: 305-261-3266
Fax: 305-261-4901
Venezuela
Phone: 582 261 0710
Fax: 582 261 5257

3Com Mediterraneo
Milano, Italy
Phone: 39 2 253011
Fax: 39 2 27304244
Rome, Italy
Phone: 39 6 5917756
Fax: 39 6 5918969
Spain
Phone: 34 1 3831700
Fax: 34 1 3831703

3Com Middle East
Phone: 971 4 349049
Fax: 971 4 349803

3Com Nordic AB
Sweden
Phone: 46 8 632 91 00
Fax: 46 8 632 09 05
Norway
Phone: 47 22 18 40 03
Fax: 47 22 18 23 85
Denmark
Phone: 45 33 37 71 17
Fax: 45 33 32 43 70
Finland
Phone: 358 0 435 420 67
Fax: 358 0 435 422 00

3Com South Africa
Phone: 27 11 803 7404/5
Fax: 27 11 803 7411

3Com UK Ltd.
Buckinghamshire
Phone: 44 1628 897000
Fax: 44 1628 897003
Manchester
Phone: 44 161 873 7717
Fax: 44 161 873 8053
Scotland
Phone: 44 131 220 8228
Fax: 44 131 226 1410

Printed in U.S.A.

500619-001 7/96