CHAPTER 10

Firewalls

Introduction

The Internet plays an important role in our daily life. Today, everyone is "connected" to everyone else almost at any given instant as we are connected to the Internet most of the time and interacting with others through e-mails or instant messengers like Skype or are using some applications on the web. With the innovation of high-speed computing devices, large-scale deployment of wireless networks, Web 3.0, Cloud computing, and social networks, "always connected" is a reality. The Internet continues to grow exponentially. Most of the businesses are connected on and through the Internet. E-commerce, e-business, and other Internet-related businesses are growing at a faster rate than ever before. According to an estimate by one of the leaders in network systems and services, the number of globally connected devices, which was around 8 billion in 2013, is expected to reach 25 billion by 2015, outnumbering the people by twice as much. And the number of devices that are going to be connected to the Internet is estimated to go as high as 50 billion by the year 2020.[1] According to the latest statistics, more than 75% of the world's population will be connected to the Internet by 2020. The Internet is bringing together people, processes, and data to make network connections more relevant to today's world. Demand for network-based applications and services is exponentially growing.

Though these applications have an immediate benefit to the end user, they can pose security risks to the individual user and the information resources of a company and government. Any information on the network is an asset and must be protected. Without adequate network security, many individuals, businesses, and governments are at risk of losing those assets. The goal of network security is to:

- Protect Confidentiality
- Maintain Integrity
- Ensure Availability

In the previous chapter, we discussed network security threats. Network attacks can be broadly classified into three categories:

- Network Access attacks: An intruder attempts to gain unauthorized access to a system to retrieve data. The attacker gains access by either cracking the network or system passwords or he already has access to the network but not to the resources. The attacker will use exploits to gain access. Improper configuration can often expose a service to substantial risk.
- Denial of Service (DoS) attacks: The purpose of the attacker is to bring the network and network resources down. The intent is to deny the authorized users access to the resources. He is attacking the availability of the network and resources. In Distributed DoS (DDoS) attack, the source consists of multiple systems that are spread across geographical boundaries.
- Reconnaissance attacks: The purpose is to identify the weakness of the system so that they can attack it at a later time. A typical reconnaissance attack consists of scanning the network devices for the open TCP port to see which application is running as well as to try to determine the operating system on which it is running.

The greatest threat arises from within the network, namely from insider attacks. Since an insider already has enough information about the network, network resources, and access to these resources, he/she can easily hack into the network. An outsider's skills combined with an insider's access could cause significant damage to an organization. Hence, the network has to be guarded from both insider and outsider attacks.

How Do You Protect a Network?

Information security has one purpose: to protect assets. How do you protect a network and be protected yourself? You build strong walls and fences to stop the enemy from entering your compound but provide a small, guarded door to friends. For a long time this was the strategy that was used for a "closed network," as illustrated in Figure 10-1.

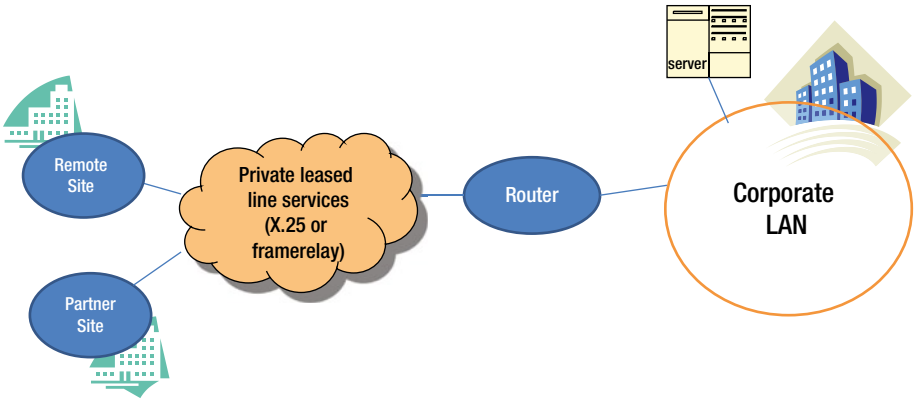

Figure 10-1. *Closed Network*

The original corporate network usually implemented a "closed network," which typically allowed secure access only to known parties and employees. In the early days, there was no outside connectivity and no World Wide Web. This worked well until the advent of World Wide Web and e-business.

As e-business, World Wide Web, and related applications continue to grow, a "closed network" was no longer closed and "private" networks started getting connected to the outside public Internet as well. Extranet connected internal and external business processes. Companies soon realized the value of supply chain management and Enterprise Resource Planning (ERP) systems to their business. Enterprises also realized the benefit of e-commerce applications to business partners and consumers, and connecting sales-force automation systems to mobile sales force. Today, an enterprise network demands an "Open Network" with the flexibility to connect to the Internet and web applications, and to support telecommuters and mobile sales force, accessing through mobile devices and much more as Figure 10-2 illustrates.

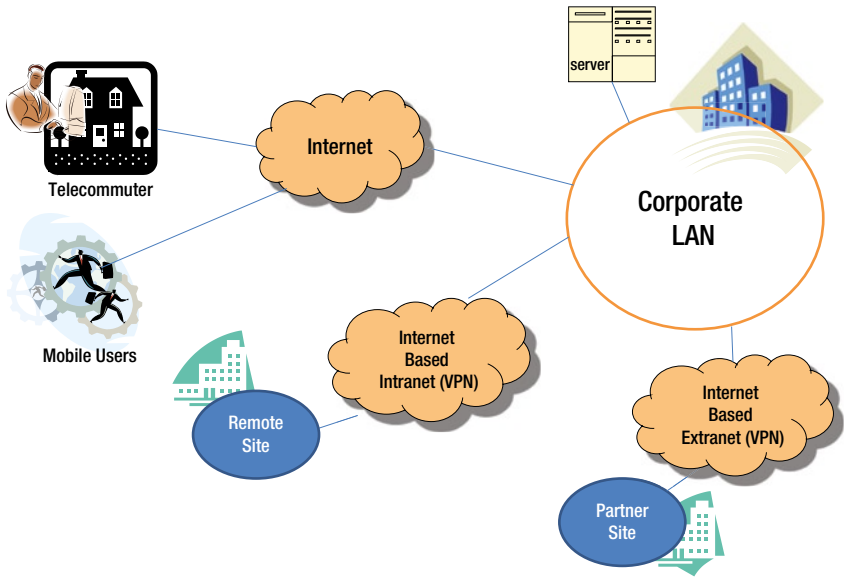

Figure 10-2. *Open Network*

Typical network architecture has three layers:

- Access Layer
- Distribution Layer
- Core Layer[2]

Network security can be translated to providing security at each of these layers – access or perimeter layer, distribution layer, and core layer. Different devices are deployed at each layer to protect the network and its assets.

In any fast-growing industry, changes are expected. As you keep adding more devices to the network, the risk also increases. Hackers are becoming smarter every day and many hacking tools are publically available on the World Wide Web. Attacking networks and associated resources has become relatively easy because of various tools and techniques available on the Internet. This is a cause for concern to organizations and enterprises. If the security of the network is compromised, there could be serious consequences, such as loss of privacy, breach of confidentiality, theft of identities, and legal liabilities. Figure 10-3 illustrates threats and their consequences.

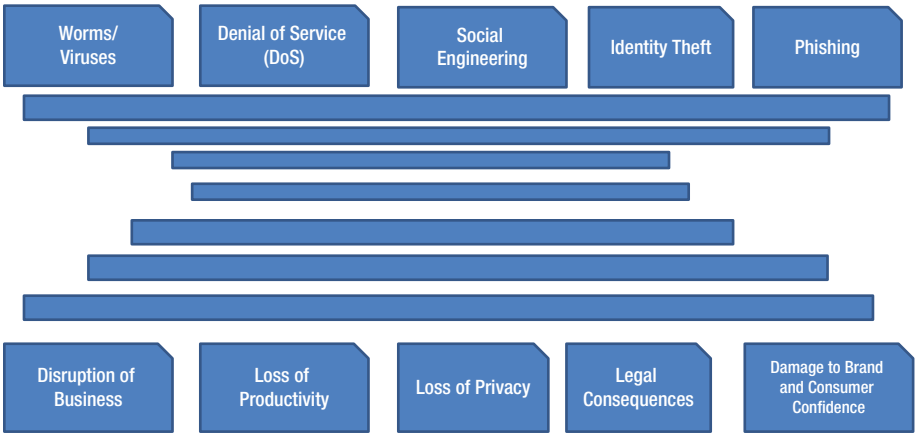

Figure 10-3. *Threats and consequences*

Various security devices such as Firewalls, Intrusion Detection Systems, Authentication and Access Control devices, and Virtual Private Networks (VPNs) are implemented at different layers. To understand these things better Table 10-1 provides an analogy.

Table 10-1. *Generic analogy of security devices*

Sl. No.	Description	Security Devices
1	Doorman, Lock and Key	Firewall
2	Passport and VISA	Access Cards, Biometric
3	Surveillance	Intrusion Detection System/Intrusion Prevention Systems
4	Escorting guests to your lab	Virtual Private Network (VPN)
5	Guards, Guard dogs	Intrusion Detection System/Intrusion Prevention Systems

In this chapter, we discuss firewalls. The chapters that follow cover Intrusion Detection System (IDS), Intrusion Prevention System (IPS), and Virtual Private Network (VPN), respectively.

Firewall

The term firewall, in the real world, means a wall built to protect from fire and intended to slow the spread of fire through a structure. The same concept is used in networks too. A network firewall is intended to stop unauthorized users from accessing the network and its services from other external networks. The most common deployment of firewalls is between a trusted network of an organization to an untrusted network, typically the Internet, as Figure 10-4 illustrates. Typically, the Internet Service Provider (ISP) connection terminates at a border router and then connects to a firewall.

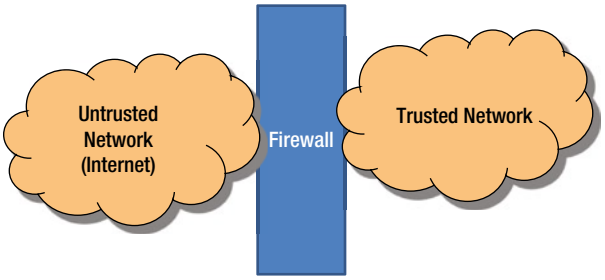

Figure 10-4. *Firewall Deployment*

Basic Functions of Firewall

A firewall in the networking world should examine the traffic that is entering into the network and pass the "Wall" based on some rules defined by the network and its resources. It acts as a security guard, who normally sits at the main gate, and checks your identity and access privileges and lets you in. Depending on the type of organization, the guards may screen people who are exiting the gate too. Many of the Internet and Information security concepts can be described using some of these practical examples.

If you talk to various vendors about a firewall and its function you will get several different answers or definitions. In its simplest form, a firewall is a combination of hardware and software devices, which bifurcates the internal network from the outside networks (Internet) and blocks certain traffic and allows some specific traffic. However, it has three basic functions (depending upon its type):

- Packet filtering: A firewall filters the IP packets. The IP headers of all the packets that enter or exit the network firewall are inspected. Firewall makes an explicit decision on each packet that enters as to whether to allow the packet or deny the packet.
- Stateful Packet Filtering: Here the packet filtering goes beyond basic packet filtering. This keeps track of state of connection flows for all the packets, in both directions. It also keeps track of all the IP addresses currently connected at any point of time.
- Application Level Gateways (Proxy): A firewall is also capable of inspecting application level protocols. This requires the firewall to understand certain specific application protocols.

Packet Filtering

As the name suggests, a packet filter filters the packets that are entering and leaving the network. The firewall inspects each IP packet and a decision is made. Each packet is compared with a set of filter rules and based on any match, the packet is either allowed, denied, or dropped. Packet filtering works on the network layer and transport layer of the OSI reference or TCP and IP layer of TCP/IP, as Figure 10-5 illustrates. It does not remember the state and hence it is called as ***stateless firewall***.

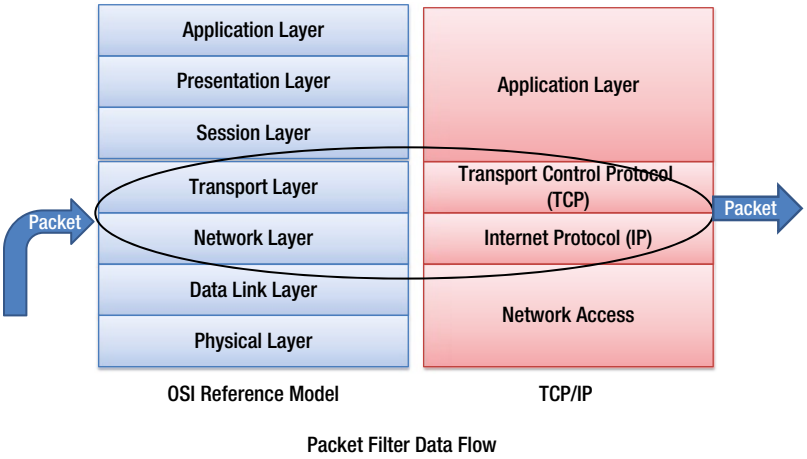

Figure 10-5. *Packet Filtering related Layers*

Packet filters usually permit or deny network traffic based on the following:

- Source and destination IP addresses
- Protocol such as TCP, UDP, or ICMP
- Source and destination TCP or UDP port addresses
- Flags in the TCP header – ACK, CLOSE, and SYNC
- IP fragmentation flag
- Direction of the packet – inbound or outbound
- Physical interface

How a packet filtering firewall works

Figure 10-6 illustrates how a packet filtering firewall works.

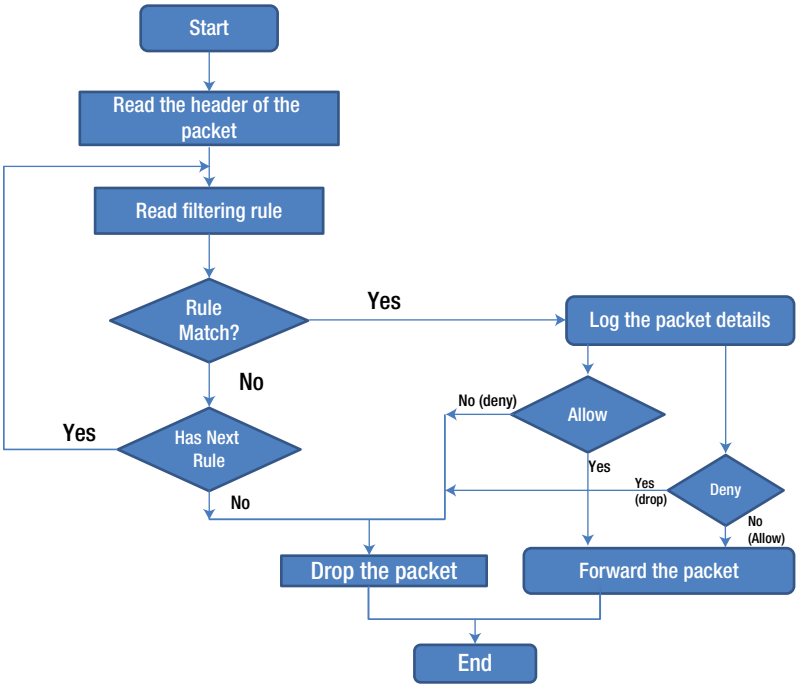

Figure 10-6. *Packet filtering flow diagram*

A packet filter firewall is configured with a set of rules that define when to accept a packet or deny. When the firewall receives a packet, the filter checks the rules defined against IP address, port number, protocol, and so on. If the rule matches "accept," then the packet is accepted in the network, otherwise it is dropped.

To understand configuring packet filtering rules, you need to first understand TCP/IP protocol, what an IP packet is and how they are handled at each layer. RFC 791[3] and RFC 793[4] provide the details of IP protocol and TCP protocol. From a packet filtering point of view, the IP header contains three important pieces of information:

- The IP source address – four bytes long, and typically written as 192.168.2.34
- The IP destination address – four bytes long, Just like source address
- The IP protocol – specifies whether it is a TCP packet or UDP packet, an Internet Control Message Protocol (ICMP) packet

If the network bandwidth is smaller than the source, IP may divide a packet into a series of smaller packets called *fragments.* Fragmenting does not change the structure of the IP packet, as shown in Figure 10-7, but it may set a flag inside the IP packet, stating that the body contains only a part of a packet.

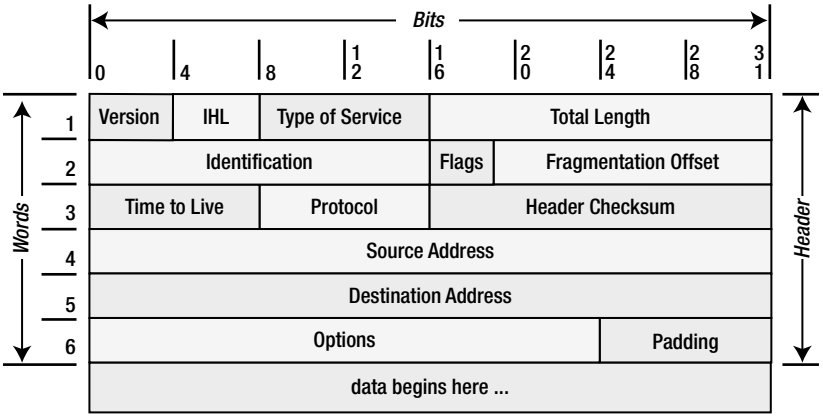

Figure 10-7. *IP Packet (source: RFC 791[3])*

TCP Layer

As Figure 10-8 shows, the TCP packet header contains the following information:

- The TCP source port – a two byte number, which specifies the application process that this packet belongs to
- The TCP destination port – a two byte number, which specifies the application process that this packet has to reach
- The TCP flags field – contains TCP protocol information such as connection establishment, closing connection, and size of packet

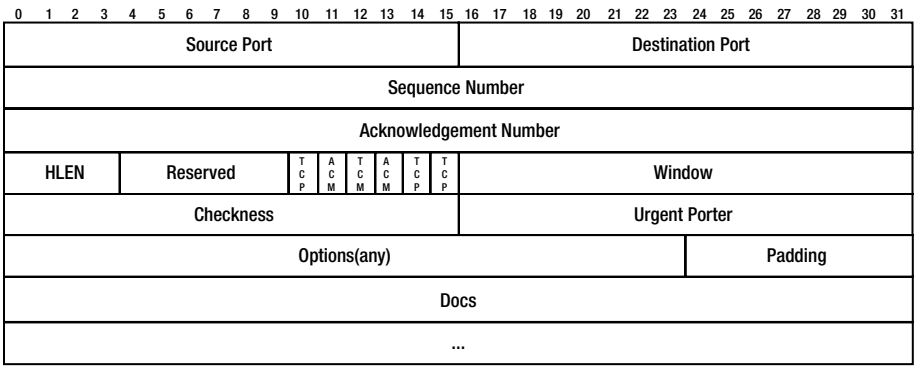

Figure 10-8. *TCP Header (source: RFP 793[4])*

An Example of Packet Filtering Rules

Let's say that you want to allow all the IP traffic between the external host (say 162.22.34.56) and the host on your internal network (class A 10.1.1.2). Table 10-2 lists the packet filtering rules.

Table 10-2. *Packet filtering rules*

Rule	Direction	Source Address	Destination Address	Application (TCP port)	Filter Set	Action
1	Inbound	Trusted external host (162.22.34.56)	Internal (10*.*)	Http	Any	Permit
2	Outbound	Internal	Trusted External host (162*.*)	SMTP	Any	Permit
3	Inbound or Outbound	Any	Any	TFTP	Any	Deny

Some examples of packet filtering rules are:

- Allow e-mail and HTTP (web) services, but block services such as TFTP and Telnet
- Block all incoming connections from outside systems except for SMTP connections (so that you can receive e-mails)
- Allow port 443 for all service destination addresses
- Allow port 80 for all service destination addresses

Advantages and Disadvantages of Packet filtering

The main advantage of the packet filter firewall is its simple rules: allow or deny:

- A strategically placed packet filtering firewall can protect the entire network. Most of the routers support packet filtering. If you have a border router placed just after Internet ISP, with the packet filtering enabled, you can protect an entire network regardless of the network size.
- Packet filtering is widely available in routers. Leading networking vendors like Cisco, Juniper, and HP, provide packet filtering on their core and edge routers known as Access Control Lists (ACL), which is configured in all the border routers.

There are several disadvantages:

- The packet filtering rules tend to be hard to configure. You need a lot of expertise and proper strategy to configure it right
- Once it is configured, it is difficult to comprehensively test and verify whether it is working correctly or not
- It is a stateless machine. It does not remember the state of the previous packet. Stateless packet filters are vulnerable to attacks. Hence, some of the attacks, such as spoofing attacks, can easily bypass firewall rules

Stateful Packet Filtering

The main disadvantage of basic packet filtering is that it is stateless. It does not remember the state of a telnet connection or an FTP connection flow already established or source port number of the client. In any application service, the TCP destination port is typically identified. For example, destination port for HTTP is 80 or FTP is 21. However, an FTP client can use any port and typically, this is dynamically chosen at run time. In basic packet filtering, since the firewall does not remember the previous state, and relies only on filtering rules, there is an amount of risk introduced, as some of the packets may bypass the firewall.

This type of risk is unavoidable for a basic packet filtering firewall. Therefore, all modern firewalls go beyond basic packet filtering, and are **stateful**. This means, the firewall keeps track of the **state** of connection flows for all the packets, in both directions – entering and exiting the firewall. The stateful firewall also keeps track of all the IP addresses currently being connected to the firewall.

A stateful firewall allows only those packets belonging to an allowed session. For example, instead of permitting any host to send data to the TCP port 8080, the firewall allows only those packets which already have the full TCP connection. Furthermore, it can check whether the packets are really of 8080 protocol traffic and it can enforce constraints at the application layer.

The main advantage of a stateful firewall is that the administrator no longer needs to write broad filtering rules, mentioning all the TCP services to allow or deny. The administrator needs to list the attributes of the flow's first packet in the rule base and the rest is taken care of by the firewall cache mechanism. An additional benefit is that the rules can be shorter. A single rule can describe the flow. Maintaining firewall rules becomes easy and prevents errors from creeping in. Finally, the stateful firewall provides better performance. With better structure of the table, a cache lookup can be made more efficient.

Network Address Translation (NAT)

An IP address is 32 bits long, and with the current schema, the maximum number of hosts you can have is about 4 billion different IP addresses. This puts a limit on the number of hosts that you can connect to the Internet. Since many companies have many hosts that need to communicate with the Internet, these 4 billion addresses are not enough and very quickly, these addresses get depleted. In 1994, RFC 1631[5] suggested a short term solution to this problem known as NAT (Network Address Translation). As it turned out, NAT not only solved the addressing problem, it became one of the ways to protect our internal network identity. RFC 791[3], defines a set of IP addresses in each class as private addresses, used only within the private network and the rest of the addresses can be used as public addresses. The reserved ranges for private network are:

- 10.0.0.0 – 10.255.255.255 (10/8 prefix)
- 172.16.0.0 – 172.31.255.255 (172.16/12 prefix)
- 192.168.0.0 – 192.168.255.255 (192.168/16 prefix)

The primary reason for this is to ensure that the IP addresses are allocated efficiently. A private range of IP address, for example, 10.0.0.0/8, which is not in use on the Internet, when tries to connect to the outside world, has to be replaced by a public address. This is done using NAT. For example, a host is listed as 10.62.1.3, which is the source IP address. After NATing, the source address is replaced by a public address of 23.2.32.3. The destination computer sees just this IP address; the internal network address is never known to the outside world. Therefore, NAT provides protection to the inside network resources. This process is shown in Figure 10-9.

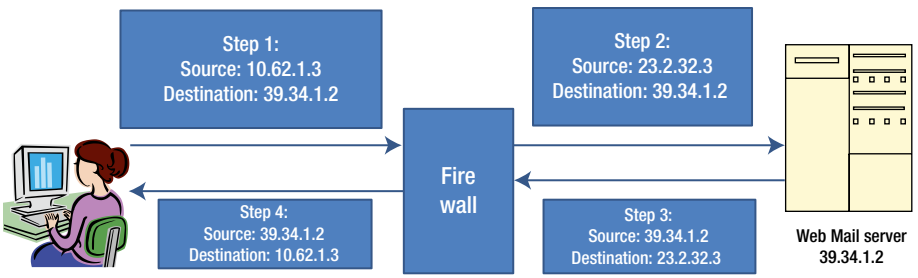

Figure 10-9. *NAT protecting internal network resources*

With NATing, companies can have only one public IP address. Any internal host that connects to the Internet or an outside network is NATed at the firewall. Internal hosts will never know that NATing has taken place. The destination host is also not aware of the NATing that has taken place. Thus NAT saves IP addresses. If more than one internal computer system is communicating with the Internet, as per the RFC 1631[5], all NAT firewall has to do is to change the source IP address, while it also needs to change the source port number with an unused port number above 1023 and keep track of this list temporarily during the connection.

NATing serves as a *basic security* measure that can make it a bit more difficult for an external attacker to map to the internal network IP addresses. When NATing is performed, the firewall rewrites the source IP address and stores both the altered source and destination IP addresses in the IP header. Internally, the firewall keeps track of the interface that is connected inside the network and interfaces connected to the outside network (global network). Global addresses are registered and assigned by an Internet Service Provider (ISP). The firewall internally identifies the packets as inbound or outbound, that is, in which direction the traffic is actually moving, and accordingly does the translations (NATing).

Static Translation

The NAT can use either a static or a dynamic mapping. In static NAT, configuration mapping is always fixed in a specific way. In a static NAT, a pool of inbound IP addresses are mapped to a pool of outbound IP addresses on a one-to-one basis. Once it is configured, it is fixed. This is particularly useful for a web server which has a consistent address that is accessible from the Internet. Figure 10-10 shows an example of a static NAT. Each inside address (172.16.1.1, 172.16.1.2, and 172.16.1.3) is mapped one-to-one with a global address (11.1.2.1, 11.1.2.2, and 11.1.3.3).

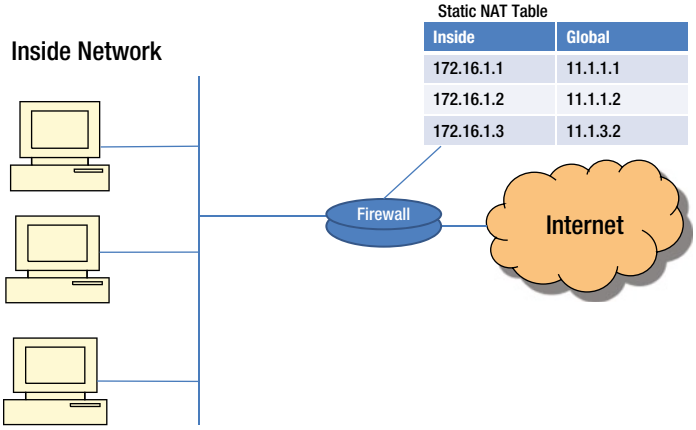

Figure 10-10. *Static NAT*

Dynamic Translation

In dynamic NATing, the mapping is not static. Mapping is based on the available IP address from a pool of public addresses. When a host inside the network requests access to the Internet, dynamic NAT picks up an IP address from a pool of addresses that has not been assigned and is not in use by any other host and assigns it to the host. Dynamic NATing is useful when fewer addresses are available and a larger number of hosts are to be connected to the Internet.

Port Address Translation (PAT)

When there is only one global IP address and multiple hosts inside the LAN trying to access the Internet, then we use what is called Port Address Translation (PAT). This situation is called overloading. The NAT/PAT box needs a way to keep track of the local addresses trying to connect to the Internet. This mapping is done using TCP/UDP ports. TCP/UDP uses 16 bits port numbers, which allows 65536 different services or source ports to be identified. When performing translation, PAT tries to use the original port number, if it is not used. If it is already in use, then the next available port number from the appropriate group is used.

The advantage of PAT is that multiple internal hosts can share a single global IP address. Global IP addresses are provided by ISP and they are expensive. Having one global IP address and with the help of PAT, organizations can save money. The second advantage is security. Internal networks are never exposed to the outside public network, making attacks from the outside more difficult and less frequent.

One disadvantage of PAT is the limitation on the number of hardware connections it supports. If too many internal hosts are trying to connect at the same time, then the hardware may run out of unused ports.

Application Level Gateways (Application Proxy)

As the name implies, an Application Level Gateway (ALG) inspects packets all the way up to the application layer and determines whether a packet is allowed or denied. It gives higher security than the packet filtering as the inspection is done all the way up to the application, as illustrated in Figure 10-11. However, this takes more CPU processing time and the necessity of having the knowledge of application protocol.

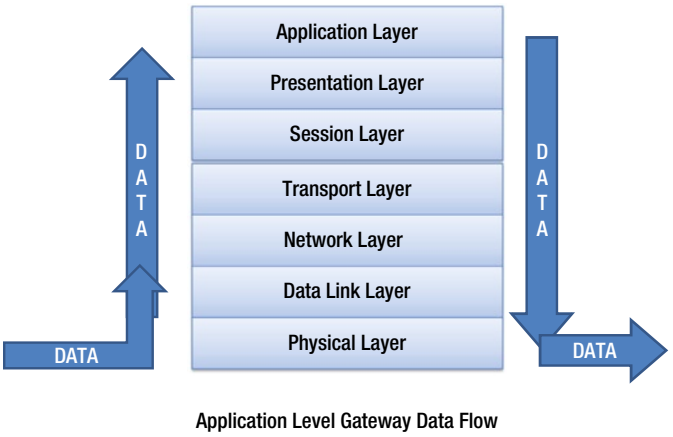

Figure 10-11. *How the Application Level Gateway Works*

An Application Level Gateway runs independently, copies and forwards information across the gateway and functions as a proxy server. It prevents a direct connection between a trusted server or client and an untrusted host. The proxies are application specific. Any new application that comes into the network needs to be informed to the application proxy, so that the rules may be set up and get executed for this application. It sits between a network firewall and a trusted host as shown in Figure 10-12. It can filter packets at the application layer.

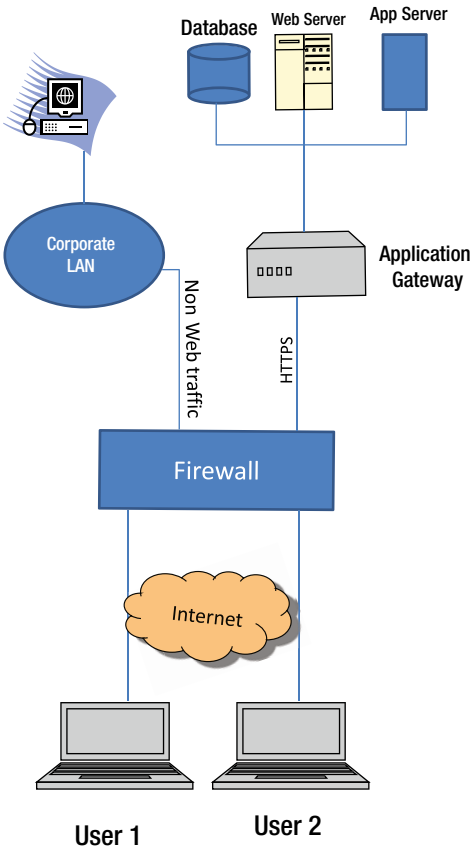

Figure 10-12. *Packet Filtering at the application layer using the Application Level Gateway*

An Application Layer Gateway maintains a complete TCP connection state and sequencing. ALG typically evaluates packets that are applied to the interface after security policies are applied. ALGs operate behind the NAT or a firewall.

Some of the advantages of ALG include:

- Direct connection between internal and external hosts are not allowed
- User-level authentication is supported
- Packet is inspected right up to application data payload

However, there are limitations. The disadvantages of ALG include:

- More processing power required
- Slower than packet filtering
- Not every application protocol is supported. Whenever there is a new application, corresponding proxies must also be implemented.

Firewall Deployment Architecture

Firewall is the first layer of protection to your internal network. Depending on the security strategy of the organization, firewalls can be deployed at different layers in the network. The following deployment scenarios are the most common.

Option 1: Bastion Host

This is the basic option where the firewall is placed in between the internal and external network as shown in Figure 10-13. This topology is well suited for simple networks. This has a single boundary, hence, once someone penetrates the firewall, they have gained unrestricted access to the protected network.

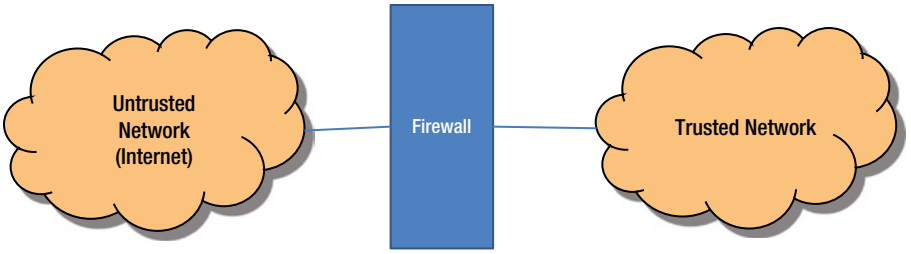

Figure 10-13. *Bastion Host*

Option 2: Staging Area or Demilitarized Zone (DMZ)

This topology allows organizations to host servers which face the internet directly, and separates the trusted network and the Internet (see Figure 10-14), thus allowing the users to access the internet securely. If a malicious user manages to compromise the firewall, he or she will not have access to the intranet services (provided the firewall is properly configured). This is the most commonly deployed architecture. The DMZ hosts all the servers offering public services, which face the Internet. The private zone contains all internal network resources such as the file server, the application server, the database servers, user workstations, and printers, which do not have any business connecting to the Internet. The DMZ zone hosts your public Web server, mail server, DNS servers, and other similar systems.

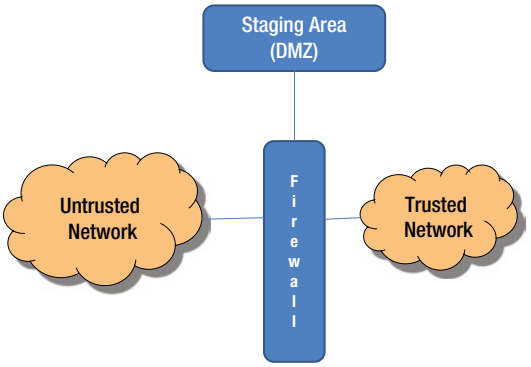

Figure 10-14. *Multiple Firewall Deployment - DMZ*

As more and more of the networks grow, the need to create a zone to protect internal assets has become imminent. Hence, many deployments now have a separate zone called Demilitarized Zone (DMZ) to separate the internal assets and the assets connecting to the Internet.

Multiple Firewall

In this scenario, you will deploy two or more firewalls to create two or more zones, as shown in Figure 10-15. Since you have more zones, the network is more secured and you can plan your organization security policy better. One division is to place your sensitive resources in a separate zone, for example, all accounting and finance servers in one zone, public facing servers such as the Web server, the Mail server, and the DNS server in a more secured DMZ zone. Systems that provide services to the general public (web server) may be placed in a different zone than systems which offer authenticated users services such as intranet applications.

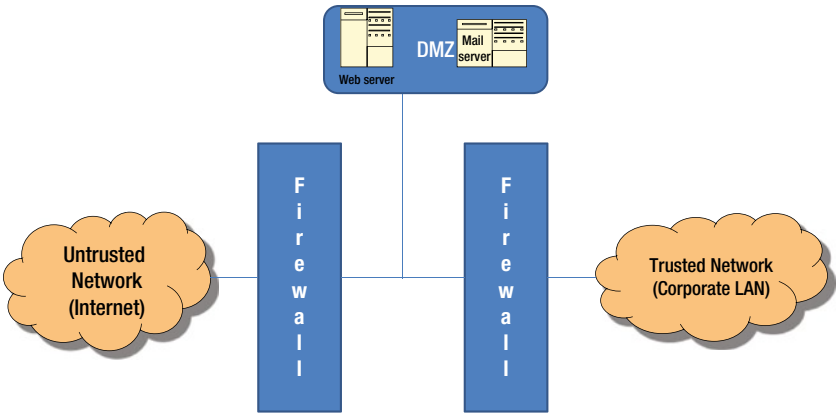

Figure 10-15. *Multiple Firewall Architecture*

Personal Firewall

A personal firewall, sometimes referred to as the desktop firewall, is a software-based firewall to protect a single system from intruders. Personal firewall protection is useful for those users who are always connected to the Internet using a Digital Subscriber Line (DSL) or a cable modem. Such connections use a static IP address, which makes the system more vulnerable. Personal firewall is better than anti-virus software as personal firewalls control the traffic on the Internet by filtering inbound and outbound traffic, and strengthening the user's preferences.

A personal firewall monitors the traffic that is going in and out of your system and grants access only to those who have passed the firewall policy set on the system. It identifies and blocks malicious software including viruses and worms. It also alerts when an unauthorized program attempts to hijack your system.

Antivirus software detects worms, viruses, and Trojans, but a personal firewall protects from intruders who attempt to hijack your system. Antivirus software along with a personal firewall can thwart many attacks on your system by blocking undesirable traffic.

Firewall Best Practices

Conceptually, firewall technology divides the enterprise network into a multiple segmentation of networks. Each segment represents a different level of trust and protection. Firewall rules allow for better control over the traffic on the network. If a hardware-based firewall is deployed, one should understand the security features supported by that particular hardware and configure the rules to provide maximum protection to your network. For the best protection of your network and firewall, the first thing is to change the default administration passwords and configure recommended firewall rules (never use the default configuration of the firewall).

The following types of traffic are always recommended to be blocked by the firewall rules (firewall configuration):

- Inbound traffic from a non-authenticated source system, for example, if you have branch offices in four different locations, your firewall rules should allow inbound traffic accessing your headquarters' network from these four locations and block others.
- Inbound traffic containing ICMP (Internet Control Message Protocol) traffic. For example, someone from outside can simply send continuous ping packets to the access router and overload the router thus denying genuine traffic.
- Inbound or outbound traffic from a host using source address which is the private address range as defined in RFC 1918[6] to avoid any spoofing attacks:
  - 10.0.0.0 to 10.255.255.255 (Class A)
  - 172.16.0.0 to 172.31.255.255 (class B)
  - 192.168.0.0. to 192.168.255.255 (class C)
- In addition, the following addresses also should be blocked:
  - 169.254.0.0
  - 192.0.2.0
  - 255.0.0.0
- Block any inbound with SNMP traffic
- Deny all traffic by default, and only enable those services as required
- Limit the number of applications that run on the firewall (VPN, DHCP, content filtering, etc.)
- Enable firewall logging and monitoring. It is also recommended to use Log monitoring solution to monitor security alerts
- Update firewall operating system patches regularly
- It is also recommended to run penetration tests on the firewall to check how secure your network is from attacks
- Schedule regular firewall audits
- Consider firewall in conjunction with other security devices such as:
  - Network based Intrusion Detection System (IDS)
  - Anti-virus software
  - Web content filtering software/hardware
  - URL filtering software/hardware
  - Separate authentication system such as LDAP or RADIUS

The following best practices can be applied to Firewalls:

- Understand firewall features while implementing them, including their capabilities and limitations. Firewall configuration requires understanding of information security and what you are trying to secure; understanding of TCP/IP protocol, applications that are running on your network, the type of users (contractors, staff, students etc.); knowledge of the type of firewall you are deploying and how you are deploying. Planning and architecture of network deployment is essential before configuring firewall rules. Also firewall features can vary from product to product. Available features need to be understood before configuring the firewall.

- Configure the firewall appropriately and carefully. Before configuring, you should answer the following questions: Do you need to configure DMZ (Demilitarized Zone)? How many network segments? Is it the entry or exit point firewall?
- Document the firewall build and configurations. The latest version of an operating system and firewall rules have to be documented. It is also recommended to regularly upgrade any patches available by the firewall vendor. Before making any changes to the firewall, rules need to be evaluated.
- Determine what you want to filter through the firewall and why. Once a decision is made then set the rules clearly on the firewall. Firewall rules are defined based on the security requirements and the type of applications you are running on your network. For example, do you need FTP? Do you want to allow your employees to access Facebook or Twitter? What are your security and IT policies related to pornographic sites? Do you want to allow remote access? Based on your applications and security policy your firewall rules are set.
- If there are multiple firewalls, then ensure that the rules set on each of these are complementary to each other and not conflicting with each other. For example, finance is on a separate segment and traffic control is more granular in nature. Only outbound traffic is allowed and limited inbound traffic is allowed.
- Ensure that the multiple entry points to the organization are appropriately protected with suitable equipment like firewalls
- All changes on the firewall have to go through a detailed impact analysis and have to be implemented only after approval from competent authorities or Change Control Board (as applicable to the organization). Before making the changes, business impact should be analyzed. How it is impacting the business should be assessed by all the business units and consensus should be taken for the changes.
- Change Control Board should approve the changes on the firewall. Changing firewall rules should have approval from all the stakeholders. Information security policy is defined based on the organization and business needs. Hence, any changes to firewall policy are changes to business policy and need proper process and approval. Without the proper approval process, firewall policies should not be altered. For example, a new partner network is allowed to access the internal network on a temporary basis, but it still needs to go through proper approval and business impact process.
- Relevant incident analysis should check on the working of the relevant firewall rules to discover any potential firewall weaknesses.
- Logging of firewall activities is essential. The logs have to be copied to a storage area and retained. Firewall logs have to be analyzed for any unwanted activities that would otherwise go unnoticed.

Auditing of Firewall

Periodic auditing of the firewall is a good practice at any organization. Some of the aspects to be considered during the audit are:

- Reasons for the firewall implementation – If rules were supposed to have been set, confirm that they are set appropriately
- Confirm that firewall configuration is appropriate and documented correctly

- Where multiple firewalls are deployed, are the rules defined complementary to each other and not conflicting with each other?
- Where there are multiple entry points to the network of the organization, are they all well protected appropriately?
- Confirm whether all changes to the firewall are carried out only after proper impact analysis and approval from competent authorities.
- Are the relaxations made to any of the rules of the firewall, including opening of specific ports for business purposes done only when justified and are carried out only for a limited period? Relaxations made are not left beyond the period for which such relaxations were permitted initially.
- Are firewall limitations, if any, considered during risk assessment and appropriate risk mitigations identified and applied?
- Are firewall logs held securely and analyzed regularly? Are the analysis results looked into and appropriate containment and corrective actions carried out?
- Are any issues identified through log analysis analyzed further and appropriate corrective actions taken?
- Are the firewall rules tested periodically? Are the test records maintained? Were there any issues observed during the testing? If so, are these issues resolved immediately by appropriate containment and corrective actions?

Chapter Summary

- We discussed how the Internet has transformed the way we carry out our business. While the Internet has substantial benefits we explored the fact that all these applications also expose the individuals and organizations to the information security risk. We then looked into the types of network attacks, including network access attacks, denial of service attacks, and reconnaissance attacks. We also looked at how insiders and outsiders potentially pose huge risks to organizations.
- We explored how we can protect the network and in turn get ourselves protected. We started the discussion of security of the network from the discussion on the "closed networks" implemented when there was no outside connectivity and no World Wide Web. We then looked into the fact as to how the development of World Wide Web and the applications thereon led to the necessity to allow connectivity from internal networks to the external network and vice versa. Hence, today, an enterprise network demands an "Open Network" with the flexibility to connect to the Internet and web applications, and to support telecommuters and mobile sales force, accessing mobile devices and much more. Then we looked into how this can lead to attacks on the networks by hackers using various tools available easily to them. We also mentioned that if the security of the network is compromised, there could be serious consequences, such as loss of privacy, breach of confidentiality, theft of identities, and legal liabilities. Then we mentioned that to protect the network, various devices like Firewalls, Intrusion Detection Systems, Authentication and Access Control devices, and Virtual Private Networks (VPN) are implemented at different layers.

- We discussed what a firewall is. A network firewall is intended to stop unauthorized users from accessing the network and its services from other external networks. The most common deployment of firewalls is between a trusted network of an organization to an untrusted network, typically the Internet. We then discussed the basic functions of firewall, including basic packet filtering, stateful packet filtering, and application level gateways (proxies). We then went on to discuss packet filtering and how packet filters allow or deny network traffic. Further, we elaborated as to how packet filtering firewalls work. We mentioned the filtering rules and how matching leads to the deny or allow rule to be triggered which in turn leads to dropping of the packet or forwarding of the packet. We also gave an example of packet filtering rules. We then went on to discuss the advantages and disadvantages of the packet filter firewalls.
- We also discussed the stateful packet filtering firewall, which keeps track of the state of connection flows for all the packets, in both directions and also keeps track of all the IP addresses currently connected to the firewall. A stateful firewall allows only those packets belonging to an allowed session. The main advantage of a stateful firewall is that the administrator no longer needs to write broad filtering rules, mentioning all the TCP services to allow or deny. The administrator needs to list the attributes of the flow's first packet in the rule base and the rest is taken care of by the firewall cache mechanism.
- We also discussed the Network Address Translation (NAT), both the static and dynamic translations pertaining to NAT and Port Address Translation (PAT) and how these features help in protecting the internal network details from the external parties.
- We examined an Application Level Gateway (ALG), how it inspects packets all the way up to the application layer, and determines whether a packet is allowed or denied. It gives higher security than the packet filtering as the inspection is done all the way up to the application. In this context we discussed that an Application Level Gateway runs independently, copies and forwards information across gateway, and functions as a proxy server. It prevents a direct connection between a trusted server or client and an untrusted host and that proxies are application specific. We then discussed some of the advantages and limitations of the ALG.
- We discussed firewall deployment architecture. In this context, we discussed bastion host, DMZ, and multiple firewalls. Then we went on to look at the personal firewall. A personal firewall, sometimes referred to as the desktop firewall, is a software-based firewall to protect a single system from intruders. We also looked at how this helps us.
- We explored some of the firewall-related best practices and then how firewall auditing can be carried out effectively.