Understanding Network Security

LESSON 4

OBJECTIVE DOMAIN MATRIX

Skill/Concept	Exam Objective	Objective Number
Using Dedicated Firewalls to Protect a Network	Understand dedicated firewalls	3.1
Using Isolation to Protect the Network	Understand network isolation	3.2
Protecting Data with Protocol Security	Understand protocol security	3.3
Understanding Denial-of-Service Attacks	Understand protocol security	3.3
Securing the Wireless Network	Understand wireless security	1.4

KEY TERMS

application-level firewall
ARP spoofing
backdoor attack
buffer overflow attack
circuit-level firewall
cross-site scripting (XSS) attack
demilitarized zone (DMZ)
denial-of-service (DoS) attack
distributed denial-of-service (DDoS) attack
DNS poisoning
DNS Security Extensions (DNSSEC)
DNS spoofing
email denial-of-service attack
firewall
honey net
honeypot
host firewall
HTTP flood
ICMP (ping) flood
IEEE 802.1x
intrusion detection system (IDS)
intrusion prevention system (IPS)
IP address spoofing
MAC address
man-in-the-middle attack
network firewall
Open Systems Interconnect (OSI) reference model
personal firewall
ping of death
remote code execution attack
replay attack
social engineering
spoofing
SQL injection attack
stateful inspection
stateless inspection
SYN flood
tunneling
UDP flood
User Datagram Protocol (UDP)
Wired Equivalent Privacy (WEP)
Wi-Fi Protected Access (WPA)
Wi-Fi Protected Access version 2 (WPA2)
zero-day attack

Traditionally, when looking at building information security infrastructure, the first point of focus is the network. As soon as networks began interconnecting, it was obvious that the network offered the main vector of attack. It was the one way to get to an internal network from the outside.

The philosophy around network protection was originally reminiscent of the castles of old. The best way to secure a castle was to build strong walls, dig moats, and control access to the castle through the main gate. In network terms, this meant deploying multiple layers of firewall, and then controlling who could enter the network with firewall rules, access controls, and DMZs. This practice is known as securing the perimeter, or defense in depth.

This model worked quite well until the next round of technology evolution came about. In the late 1990's, the concept of virtual private networks (VPNs) was introduced. VPNs allowed companies to securely extend their network across untrusted networks like the internet, but also impacted the perimeter of the network.

Next came wireless network technologies, moving the perimeter that needed to be protected literally into the air, offering additional challenges to the network perimeter and the layered security model.

The good news is that as network technologies have evolved that have made securing the perimeter more challenging, the security technologies available for addressing those challenges have evolved as well. In this lesson, we will discuss these security solutions, and how they can be used to address the challenges you may encounter.

■ Using Dedicated Firewalls to Protect a Network

THE BOTTOM LINE

Firewalls remain the foundation of network security technologies. There are a number of options, types, and technologies associated with selecting, implementing, and maintaining the firewalls in a network. There are also a number of factors that help determine the proper solution to meet business requirements.

CERTIFICATION READY
Where would most companies place their dedicated firewall?
Objective 3.1

One of the first things that comes to mind when people talk about information security is the firewall. Firewalls have long been the foundation of a company's network security infrastructure. But what exactly is a firewall?

A ***firewall*** is a system that is designed to protect a computer or a computer network from network-based attacks. A firewall does this by filtering the data packets traversing the network. A typical perimeter firewall is implemented with two (or more) network connections (see Figure 4-1):

- A connection to the network being protected
- A connection to an external network

Figure 4-1
Example of a firewall implementation

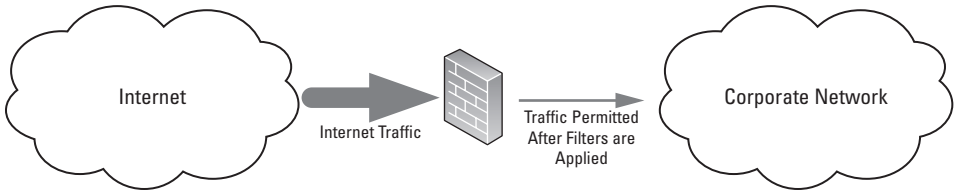

There are numerous variations on this model, but ultimately a firewall protects hosts on one network from hosts on another network.

These network connections may be referenced using different labels: internal and external, clean and dirty, secure and unsecure, local and remote, and so on. They all refer to the same model, but occasionally may need translation into familiar terminology.

In today's networks, firewalls are being used for a number of things beyond securing the perimeter. Many of today's corporate networks are being divided into security zones, secured by firewalls. Sometimes, these firewalls are not only securing internet and extranet connections, but also creating secure zones for financial systems, to secure research and development servers, or sometimes even to secure the production network from the development and test networks.

Given the widely varying uses for firewalls in today's networks, there are a variety of different firewall types. But before we get into discussing the different types of firewalls, we need to discuss the OSI model.

Understanding the OSI Model

Any discussion about network security requires a discussion and understanding of the ***Open Systems Interconnect (OSI) reference model***. The OSI model is a conceptual model, created by the International Organization for Standardization (ISO) in 1978 and revised in 1984, to describe a network architecture that allows data to be passed between computer systems. While never fully utilized as the model for a protocol, the OSI model is the standard for discussing how networking works.

As shown in Figure 4-2, the OSI model is built in the same way it is usually discussed, from the bottom to the top. The layers are: physical, data-link, network, transport, session, presentation, and application. When they are discussed, the physical layer is referred to as Layer 1, and the application layer is Layer 7. This is important to remember because in discussions, routers are often referred to as "Layer 3 devices" or a specific type of firewall might be called a "Layer 7 device." This nomenclature refers to where on the OSI model that device interacts. As a result, it is important to be familiar with the high level concept of the OSI model and what occurs at each layer.

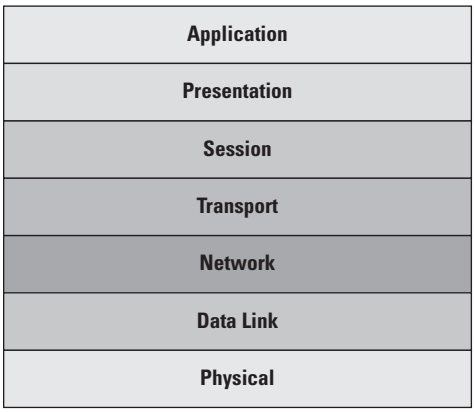

Figure 4-2
The seven-layer OSI model

Each layer of the OSI model has its own specific function. The following sections describe the function of each layer starting with the physical layer and working up the OSI model.

PHYSICAL LAYER (LAYER 1)

The physical layer of the OSI model is used to define the physical characteristics of the network, including the following specifications:

- **Media:** Cabling types, voltage, signal frequency, speed, bandwidth, and so on.
- **Hardware:** Type of connector, type of network interface card used, and so on.
- **Topology:** The topology to be used in the network, such as ring, mesh, star, and bus.

DATA-LINK LAYER (LAYER 2)

The data-link layer connects the data layer to the physical layer so that the data can be transmitted across the network. The data-link layer handles error detection, error correction, and hardware addressing (for example, the address of a network interface card).

The data-link layer is broken into two sublayers—the Media Access Control (MAC) sublayer and the Logical Link Control (LLC) sublayer.

- **MAC layer:** The MAC address is defined at this layer. The ***MAC address*** is the physical or hardware address burned into each NIC (for example, 96-4C-E5-48-78-C7). The MAC sublayer also controls access to the underlying network media.
- **LLC layer:** The LLC layer is the layer responsible for the error and flow control mechanisms of the data-link layer. The LLC layer is specified in the IEEE 802.2 standard.

TAKE NOTE*

The IEEE 802.x standards define a variety of networking technologies. For example, 802.1x defines a standard for wireless security. Ethernet is defined by the IEEE 802.3 standard.

NETWORK LAYER (LAYER 3)

The network layer is primarily responsible for routing. The network layer defines the mechanisms that allow data to be passed from one network to another. To be clear, it doesn't specify how the data is passed, but instead defines the mechanisms that permit it. How the data is passed is defined by the routing protocols (which we will discuss in more detail later in the lesson.) As a result, a router is typically known as a Layer 3 device.

TAKE NOTE*

It's important to remember that in addition to routing, that is, allowing traffic to select the best path, this layer of the OSI model specifies one other critical function. These protocols also set the addressing. In the case of TCP/IP, this is the layer where IP addresses are specified. While the data-link layer uses hard-coded MAC addresses to communicate on the physical layer, network protocols use software-configured addresses and routing protocols to communicate data across the network.

TRANSPORT LAYER (LAYER 4)

The transport layer does exactly what the name implies. It provides the mechanisms for carrying data across the network. It uses three main mechanisms to accomplish this:

- **Segmentation:** Downloading an MP3 file from a favorite music site involves dealing with a large block of data. In order to get from the music site to a PC, this file needs to be broken down into smaller, more manageable blocks, so the network can handle it. This process performed by the transport layer is called segmentation.
- **Service addressing:** Network protocols (TCP/IP, for example) provide several network services. These services are identified by ports. The transport layer ensures that when data traverses the network, it is passed to the correct service.
- **Error checking:** Transport layer protocols also perform error checking on the data and ensure that data is sent and received correctly.

The protocols operating at the transport layer come in two types:

- **Connection oriented:** A connection-oriented protocol, such as the Transmission Control Protocol (TCP), requires an end-to-end connection between hosts before data can be transmitted. Think of this like a telephone call—you don't start speaking to the person at the other end of a phone call until you are successfully connected to the person at the other end.
- **Connectionless:** A connectionless protocol, such as the User Datagram Protocol (UDP), allows for the transmission of data without requiring that a connection be established first. Connectionless protocols rely on the network to ensure the proper delivery of data from one host to another across the network. Think of a connectionless protocol like sending an email. You don't have to connect directly to the recipient before sending an email; instead you address the email, type it, and click Send. The network ensures that the email gets to the addressee.

The transport layer has an additional responsibility in the OSI model. It handles flow control of the data. Flow control determines how the receiving device accepts the data transmissions. There are two common methods of flow control—buffering and windowing:

- **Buffering:** Buffering flow control temporarily stores data in a buffer and waits for the destination device to become available. Buffering can be problematic if the sending device is able to transmit data much faster than the receiving device is able to receive. Too high a transmit rate can overload a buffer, which has a limited size, causing data loss.
- **Windowing:** In a windowing environment, data segments are grouped together, and when sent, require only one acknowledgment. The size of the window (that is, the number of segments that can be sent at one time) is agreed to by the two devices. In some cases, the window size is agreed to when the connection is first established; in others, the window size can vary based on network congestion and device resources. These types of windows are referred to as sliding windows. Windowing improves network performance by reducing the number of acknowledgements that need to be sent between devices.

TAKE NOTE*

If you are familiar with PC hardware, you may recognize these flow control methods. They are the same methods used for flow control in a PC when moving data into and out of the different types of data storage—like a hard drive, cache, and RAM.

SESSION LAYER (LAYER 5)

The session layer is responsible for data synchronization between the applications on the two devices. The session layer establishes, maintains, and breaks sessions between devices. The transport layer is responsible for connections between the two devices, and the session layer handles the same functions for the application transferring the data between the two devices.

PRESENTATION LAYER (LAYER 6)

The presentation layer converts application layer data into a format that permits the data to be transmitted across the network. Data formatted for transport across the network is not always natively readable by applications. Some common data formats that are converted by the presentation layer include the following:

- Graphics files
- Text and data files
- Music and video files

The presentation layer is also the layer where encryption and decryption of data is done.

APPLICATION LAYER (LAYER 7)

Finally, at the top of the OSI model is the application layer. The application layer takes data from the user and passes the data to the lower layers of the OSI model for transport. Responses are passed up through the layers and are displayed back to the user.

TAKE NOTE*

It's important to remember that the application layer of the OSI model is not the actual application displayed on the computer. The application layer is used to define how the applications running on a computer can take advantage of the network. For example, to print a document to a network printer, the word processing application would take the file information and pass it to the application layer, which would pass it down the layers so the data could be transmitted to the printer. Of course, there are applications that may use the network service or application that runs at application layer services, like web browsers.

While the OSI model gives us a framework to categorize technology, it is not fully implemented on today's networks. Instead, today's networks follow a simplified model usually consisting of the following four layers:

- **Link layer:** The link layer is the lowest layer of the TCP/IP model and is designed to be hardware independent. It is responsible for linking to the hardware network technology and transmits data. TCP/IP has been implemented on top of virtually any hardware networking technology in existence.
- **Internet layer:** The internet layer is responsible for connecting multiple networks together and for routing of packets between networks.
- **Transport layer:** The transport layer is responsible for end-to-end message transfer capabilities independent of the underlying network. It also handles error control, segmentation, flow control, congestion control, and application addressing (port numbers).
- **Application layer:** The application layer refers to the higher-level network protocols and services such as SMTP or FTP.

Now that you have an understanding of the OSI model, we can discuss the various networking technologies and their impact on your information security program.

Types of Hardware Firewalls and Their Characteristics

In today's network environment, the clear majority of production firewalls are hardware based. A hardware firewall is a firewall that runs on a dedicated platform, specifically designed, optimized, and hardened (the process of securing a system) to run the firewall application software.

While there are a variety of types of firewalls, with varying characteristics, firewalls share some basic functions. Firewalls filter traffic based on a set of configured rules. Generally, these rules are based on information contained in the data packets that are traveling across the network. The header information contained in those data packets provides the firewall the information it needs to properly apply these rules.

TAKE NOTE*

Don't get too hung up on the definitions of firewall types. Understand the functionality of the firewall types instead. What they are called is not as important as how the different flavors of firewalls function.

These rules are generally defined by a company's security policies and business requirements. While it is possible to configure a firewall to permit all traffic, and block specific traffic based on rules, virtually all firewalls will work based on the deny-all, permit-specific philosophy. This means that the firewall will by default, deny all traffic. Any traffic permitted to traverse the firewall will need to be explicitly configured in the firewall's rules.

There are a variety of different firewall types, and different people sometimes define firewall types in different ways. The key is to thoroughly understand the basics, because besides passing the certification test, you will generally not be called upon to identify firewall types in your day-to day-duties.

UNDERSTANDING PACKET FILTERING

The first type of firewall is known as the packet-filtering firewall. This type of firewall is considered the first-generation firewall, because the first firewalls functioned as packet filters. As we have discussed, the primary purpose of a firewall is to filter traffic. A packet-filtering firewall inspects the data packets as they attempt to traverse the firewall, and based on the rules that have been defined on the firewall, the firewall allows or denies each packet.

One of the very first versions of this firewall was the packet-filtering router. Routers can do some rudimentary packet filtering, such as permitting all outbound traffic while denying all inbound traffic, or blocking specific protocols from passing through the router, like telnet or ftp.

Firewalls significantly improve on the capabilities of a packet-filtering firewall, as they permit more granular rules. You might configure a packet-filtering firewall to block web browsing on the internet, except to your company's internet website, while permitting outbound web traffic from your internal network to the internet. Another option is to set up a rule that would drop any ping requests, unless they originate from someone on the network team's workstation.

When configuring a packet-filtering firewall rule, one (or more) of the following TCP/IP attributes should generally be used:

- Source IP addresses
- Destination IP addresses
- IP protocol (telnet, ftp, http, https, and so on)
- Source TCP and UDP ports (for example, the http protocol runs on TCP port 80)
- Destination TCP and UDP ports
- The inbound firewall network interface
- The outbound firewall network interface

Some of the more common protocols and ports that will be encountered in a production network include the following:

- FTP (file transfer) 20/tcp and 21/tcp
- Telnet (Terminal logon) 23/tcp
- DNS 53/udp and 53/tcp
- HTTP (web) 80/tcp
- HTTPS (web) 443/tcp
- SMTP (email) 25/tcp
- POP3 (email) 110/tcp
- IMAP3 (email) 220/tcp
- IMAP4 (email) 143/tcp
- LDAP (directory services) 389/tcp
- SQL Server 1433/tcp
- RDP (Terminal Services) 3389/tcp

This is not a comprehensive list, as there are thousands of different protocols and ports, but these are common protocols you will see when configuring rules on a packet-filtering firewall.

UNDERSTANDING CIRCUIT-LEVEL FIREWALLS

Circuit-level firewalls are typically considered a second-generation firewall technology. They work similarly to packet-filtering firewalls, but they operate at the transport and session layers of the OSI model.

Instead of analyzing each individual packet, a circuit-level firewall monitors TCP/IP sessions by monitoring the TCP handshaking between packets to validate the session. Traffic is filtered based on specified session rules and may be restricted to authorized computers only. When the session is established, the firewall maintains a table of valid connections and lets data pass through when session information matches an entry in the table. The table entry is removed, and the circuit is closed when the session is terminated. One unique feature of circuit-level firewalls is that sessions that cross this type of firewall appear to originate from that firewall. This allows the internal network to be hidden from the public network.

This type of firewall is also known as a transparent proxy, because all sessions appear to originate from the firewall. Circuit-level firewalls are almost always used in conjunction with other types of firewalls, because they are only able to permit sessions from authorized computers. Additional granularity is typically required in most production environments.

UNDERSTANDING APPLICATION-LEVEL FIREWALLS

Application-level firewalls (also known as proxy servers) work by performing a deep inspection of application data as it traverses the firewall. Rules are set based on analyzing client requests and application responses, then enforcing correct application behavior. Application-level firewalls can block malicious activity, log user activity, provide content filtering, and even protect against spam and viruses. Microsoft Internet Security and Acceleration Server is an example of an application-level firewall.

Now for the downside—deep inspection of application data is a resource-intensive activity, and can require significant processing power to reduce the chances of the firewall impacting network performance. The deeper the inspection, the higher the resource requirements, and the higher the possibility for network performance impacts. When deploying an application-level firewall, it is important to size it appropriately. Cutting corners on processor and RAM on your application-level firewall is an excellent formula for creating unhappy users, and it is always a better idea to go a little more powerful than your immediate needs. Remember to always plan for growth. Network utilization very seldom decreases over time. You usually don't want to go back to management in a year to fund an upgrade.

One capability available on some application-level firewalls that can help offset performance impacts of the deep inspection of application data is the addition of caching. Caching allows the firewall to store commonly downloaded data and provide it in response to requests from a user rather than having to retrieve the data from the internet. Most web browsers have this capability for local storage of commonly used pages; a caching firewall extends this capability to all users on the network. For example, if 50 employees all read the front page of the online Wall Street Journal when they come into the office, the firewall caches the first visit to the site, and then serves the stored page to the next 49 visitors.

Caching was a much more effective technology during the early days of the internet, when most of the content was static. With the advent of customizable views, mashups, and interactive content, the effectiveness of caching is becoming more and more limited.

UNDERSTANDING STATEFUL MULTI-LEVEL FIREWALLS

Stateful multi-level firewalls are designed to provide the best features of both packet-filtering and application-level firewalls. This type of firewall provides network-level packet filtering and is also capable of recognizing and processing application-level data. When configured correctly, these firewalls can provide the highest level of security of the firewall types discussed,

but are typically the most expensive firewalls. In addition, with all the available features, they can also be very complex to configure and maintain.

Understanding When to Use a Hardware Firewall Instead of a Software Firewall

Before we can look at when it's appropriate to utilize a hardware firewall instead of a software firewall, we need to look at what is meant by a software firewall. There are two basic types of software firewall:

- ***Host firewall*:** One type of software firewall is a firewall application installed on a host, used to protect the host from network-based attacks. An example of this type of software firewall would be the Windows firewall included with recent versions of Microsoft operating systems. Host firewalls are also known as ***personal firewalls***.
- ***Network firewall*:** The other type of software firewall is a firewall application installed on a server used to protect network segments from other network segments. These types of firewalls offer similar functionality to a hardware firewall. The most popular network firewalls are those produced by Cisco.

The one circumstance where it clearly doesn't make sense to use a hardware firewall is to protect a single host. To protect a single host, the best solution would be to install a software firewall on the host, with a specific set of rules based on what needs to be protected. If the host is part of a larger network, which they virtually always are, it will also be protected by any network firewalls deployed on the network.

Host firewalls aside, there are a variety of factors which will impact the decision on whether to use a software solution to protect a network. Many of these factors are related to some of the challenges associated with software firewalls. These factors include the following:

- **Host hardware:** Software firewalls run on the already busy server's general purpose hardware. This can lead to bottlenecks (such as processor, memory, or network), especially if the hardware hasn't been sized appropriately to address the traffic requirements associated with running a firewall application.
- **Host operating system:** While both hardware and software firewalls run operating systems, a hardware firewall runs a hardened operating system, providing a smaller attack surface than an unhardened operating system. In order to match the security level of the hardened OS provided by a hardware firewall, the software firewall server needs to be similarly hardened. This can require specialized expertise and additional investments in time and resources. As a result, most software firewalls have larger attack surfaces than their hardware counterparts.
- **Other applications:** Software firewalls must compete for resources with any other processes running on the host. A hardware firewall has dedicated hardware resources that are not shared with any other service. As a result, additional hardware may be needed to match the performance of the hardware firewall, due to the added resource requirements.
- **Availability/stability:** One of the potential issues associated with using a software firewall is that its reliability is tied to the reliability of the underlying operating system and associated hardware. While the hardware components in a host will generally be as reliable as the components found in a hardware firewall, they are not always available in a redundant configuration as hardware firewalls are. Operating systems have come a long way in terms of stability, but a general-purpose operating system that would be used with a software firewall is typically not as stable as the hardened operating system used on a hardware firewall.

With all the potential challenges associated with software firewalls, there are a couple of compelling reasons to use software firewalls. First, they are very cost effective. Second, they are generally less complex to install and support than their hardware counterparts.

So, in a medium to large network environment, where performance, availability, and reliability are critical, a hardware firewall is the best solution. Hardware firewalls exist in virtually every enterprise network.

For a small network, when trying to keep costs down or trying to secure a single host, using a software firewall may be the right answer.

Understanding Stateful Inspection and Stateless Inspection

As we have discussed, the most basic firewall system works by filtering packets. A packet-filtering firewall inspects the data packets as they attempt to traverse the firewall, and based on the rules that have been defined on the firewall, the firewall allows or denies each packet. The firewall doesn't consider any other information related to the packets when determining which packets are permitted to cross the firewall and which packets are blocked. This type of data packet inspection is known as stateless inspection.

In ***stateless inspection***, the data traversing the firewall are examined for information like:

- The IP address of the sending device
- The IP address of the receiving device
- The type of packet (TCP, UDP, and so on)
- The port number

Stateful inspection takes packet filtering to the next level. In addition to examining the header information of a packet traversing the firewall, a stateful inspection firewall also considers other factors when determining if traffic should be permitted across the firewall.

Stateful inspection also determines whether a packet is part of an existing session and that information can be used to determine whether to permit or deny a packet. The existing session is referred to as the state, and frequently occurs at Layer 4 of the OSI model, the transport layer. Many of today's stateful inspection firewalls can also track communications across Layers 5-7 as well.

It may sound simple, but it's a very complex process, which is why stateful inspection firewalls are typically more expensive and more challenging to configure. A stateful inspection firewall keeps track of all current sessions in a state table stored in memory. In other words, when a user initiates a connection to the MSN website to check today's headlines, the firewall stores the information regarding the session in a table. The same is done for every other connection occurring across the firewall.

As each packet is encountered by the firewall, it will be analyzed to determine whether it is part of an existing session (state) or not. If it is, and the session is permitted based on the current firewall rules, then it is passed. If it is not part of an existing session, and it is not a packet being used to initiate a permitted session, it will be dropped.

Another benefit of stateful inspection is that once a session is established, the firewall manages access based on the sessions rather than on the packets. This permits a simpler set of firewall rules when compared to traditional packet-filtering firewalls. A packet-filtering firewall requires a rule for each authorized packet. To permit a connection between Host A and Host B across a packet-filtering firewall, you need a rule that permits packets from Host A to Host B, and another rule that permits packets from Host B to Host A. Using a stateful inspection firewall, a rule can be defined that permits a connection from Host A to Host B, and then the firewall's state table management will automatically allow the return traffic.

Stateful inspection firewalls make excellent perimeter firewalls for protecting an internal network from the internet, for protecting DMZ-based hosts (discussed in more detail later in this lesson) from the internet, and for protecting extranets from connections to customers, vendors, or business partners.

■ Using Isolation to Protect the Network

THE BOTTOM LINE

In addition to protecting the perimeter of the network, there are other techniques that can be used to protect computing resources on an internal network. These technologies allow you to isolate portions of your network, provide a special use for your firewalls, or even supplement the security provided by your firewalls. Virtual local area networks (VLANs) and routing are network technologies that can help to segregate a network into security zones. Also, deploying technologies like honeypots helps to distract attackers from the important portions of the network. Firewalls can also play a part if it is necessary to create DMZs on the network. VPN, NAT, Server isolation, and Domain isolation are additional network concepts that can be used to secure the network.

CERTIFICATION READY
Which feature can be used to isolate a subnet, with all of its servers, from the rest of the network?
Objective 3.2

Understanding VLANs

Before we can discuss what a virtual LAN (VLAN) is, we need to quickly review what is meant by local area network (LAN). A LAN is a network of hosts covering a small physical area, like an office, a floor in a building, or a small group of buildings. LANs are used to connect multiple hosts. These LANs are then connected to other LANs using a router, which is a Layer 3 device.

One of the challenges with LANs as they grow larger is that each device on each LAN subnet broadcasts traffic onto that subnet. While these broadcasts will not cross a router, if there are enough hosts, the aggregate broadcast traffic can saturate a network. One solution is to deploy more routers as a way to divide the network into more manageable segments. But routers add latency to network traffic, and require a routing protocol (which we will discuss in the next section) for traffic to find its way from one part of the network to another.

Virtual LANs (VLANs) were developed as an alternate solution to deploying multiple routers. VLANs are logical network segments used to create separate broadcast domains, but still allow the devices on the VLANs to communicate at Layer 2, without requiring a router. VLANs are created by switches, and traffic between VLANs is switched, not routed, which creates a much faster network connection, as there is no need for a routing protocol to be involved. Even though the hosts are logically separated, the traffic between the hosts is switched directly as if the hosts were on the same LAN segment.

VLANs provide several benefits over a routed network, including the following:

- Higher performance on medium or large LANs due to reduced broadcast traffic
- Organized devices on the network for easier management
- Additional security because devices can be put on their own VLAN

There are several different ways to assign hosts to VLANs. These methods include the following:

- **VLAN membership by port:** The ports on the switch are defined as belonging to a specific VLAN, so any device plugged into a port would be assigned to the corresponding VLAN. For example, a 32-port switch might have ports 1-4 assigned to VLAN1, ports 5-16 assigned to VLAN2, and ports 17–32 assigned to VLAN3. While this seems like a straightforward method for organizing ports, it can be problematic if, for example, you work in an environment where users change office locations frequently. If the ports have been assigned in one section of cubes to Sales, and next week they decide to move Sales to the other side of the floor, the switch needs to be reconfigured to support them. In a relatively static environment, this model works very well.

- **VLAN membership by MAC address:** Under this model, membership in a VLAN is based on the MAC address of the host. When the VLAN is set up on the switch, the hosts are assigned based on their MAC address. When a workstation moves to another location, and connects to a different switch port, the switch automatically assigns the host to the appropriate VLAN based on the MAC address of the workstation. Because the MAC address is generally hard-coded into the host's NIC, this model is generally more usable in an environment where hosts move. One downside to this model is that it does require more initial work to set up, because it is necessary to obtain all the MAC addresses from the hosts and associate them with the appropriate VLAN.
- **Membership by IP subnet address:** In this type of VLAN association, membership is based on the Layer 3 header. The switch reads the Layer 3 IP address and associates the address range with the appropriate VLAN. Even though the switch accesses Layer 3 information in the header, the VLAN assignment is still done at Layer 2 of the OSI model and no routing takes place. This model is also conducive to an environment where there are frequent user moves. There can be an impact to performance, because the switch needs to read the Layer 3 header to determine which VLAN to assign the host to. This is generally not an issue with today's switch technologies, but it is good to be aware that there is additional overhead associated with this model.
- **Membership by protocol:** VLANs can also be organized based on protocol. This was a very useful solution when many LANs ran multiple network protocols, but with the dominance of TCP/IP in virtually every network, this model is almost never used on a modern network.

The next question to think about is, "How do VLANs help with security?" There are two basic ways to leverage VLANs in support of security.

First, because a VLAN provides logical separation, traffic on one VLAN is not directly accessible to hosts on another VLAN. However, this is of minimal use because there are techniques called VLAN hopping which can get access to traffic on other VLANs.

The second use of VLANs from a security perspective is to use VLANs to organize the hosts for assigning access permissions. This technique is used in conjunction with firewalls or access control lists. For example, create a VLAN for the section of the building in which the administrators work, and give it access through your firewalls so these employees can access all sections of the network. The Sales department might be on a VLAN with its access restricted to the Sales application servers, but is blocked from getting to the HR and Finance applications.

Understanding Routing

Routing takes one step up the OSI model from the VLAN—taking place at Layer 3. Routing is the process of forwarding a packet based on the packet's destination address. At each step in the route a packet takes across the network, a decision has to be made about where the packet is to be forwarded.

To make these decisions, the IP layer consults a routing table stored in the memory of the routing device. Routing table entries are created by default when TCP/IP initializes, then additional entries are added either manually by a system administrator or automatically through communication with routers.

But what exactly is a router? Above, we defined routing as the process of forwarding a packet based on the packet's destination address. In its simplest form, a router is any device that forwards packets from one interface to another. This is a very simple description for a very complex process.

Routers come in two basic types: software and hardware. A software router is a computer running an operating system and multiple services, including a routing service. Windows

Server 2016 supports routing, for example. Some benefits of a software router include the following:

- **Tight integration with the OS:** The routing service is frequently integrated with the operating system and other services.
- **Consistent/easier user interface:** No retraining is required on a new interface/operating system—the routing functions are configured through the standard user interface.
- **Low cost:** When adding routing to an existing server, it is not necessary to pay for dedicated hardware. This reduces the overall cost, although if you were to dedicate a software router for just routing, any cost savings would be negligible.
- **Flexibility:** Software routers allow multiple services to be configured and run on a single platform.

When is it a good idea to use a software router? Typically, software routers can be found in small offices that usually need an inexpensive, easy-to-manage solution. While there are a number of benefits to software-based routers, the drawbacks frequently outweigh them during the selection process. A hardware router will not typically be impacted by a virus, or be prone to performance problems due to a runaway process. Another circumstance in which you might use a software router is between two LAN segments, where traffic requirements are expected to be low. An example of this might be a lab segment, where you want to isolate the lab hosts, but do not want to invest in a dedicated hardware router.

While there are benefits to using a software router, there are also some significant drawbacks when compared to a hardware router. These drawbacks of a software router include the following:

- **Performance:** Due to the additional overhead associated with the operating system and any additional running services, software routers are typically slower than a hardware router.
- **Less reliable:** Any software router has the potential for issues with the operating system and other running services, as well as with the greater number of hardware components compared to a hardware router. As a result, software routers are typically less reliable than hardware routers.
- **Limited scalability:** Scaling a software router to multiple high-speed interfaces will be subject to the limitations of the computer hardware. Because most PC-based servers are not designed to route multiple high-speed network interface cards, software routers will generally not scale as easily or as much as a hardware router. Also, adding additional services like access control lists or firewall services impact a software router's performance to a greater degree than a comparable hardware router.
- **Limited protocol support:** Software routers typically do not support as many routing protocols as a hardware router. Windows Server 2016 is limited to the IP routing protocols RIP, OSPF, and BGP, and does not presently support any of the more advanced IP-based routing protocols, like BGP4.

A hardware router is a dedicated hardware device whose main function has been to route packets. This description is not as true as it was in years past. Many of today's hardware routers are multi-function devices, having additional functionality like VPN, DHCP, firewall, caching, or in some cases even intrusion detection services. The benefits of a hardware router include the following:

- **Higher performance:** Hardware routers run on custom-built, single-purpose hardware platforms with highly optimized hardware and operating systems.
- **Highly reliable:** Hardware routers are typically more reliable than their software counterparts, due in large part to the limited software capabilities, and dedicated hardware. A hardware router will typically have higher modularity than a software router.

Hardware routers also support more high availability capabilities, where they can be deployed in pairs—one will take over if the other fails. While this is theoretically possible with a software router, it is very seldom done.

- **Wide routing protocol support:** Hardware routers can typically be configured to support a larger range of routing protocols, as long as the appropriate functions are purchased. They also support a greater number of routing algorithms than a software router. In a larger network environment, this can be critical.

Using hardware routers is not always advantageous—sometimes there are drawbacks to using a hardware router:

- **Higher cost:** Hardware routers are typically dedicated platforms, which tends to make them more expensive than a software router that is also providing other services. This line is blurring as additional features continue to become available on hardware routers. However, a small hardware router can be relatively inexpensive.
- **Less user friendly:** Hardware routers are typically configured using a Secure Shell (SSH) connection, and are managed through a command-line interface. While there are graphical tools for managing routers, a lot of router configuration is still done through the command line, using an extremely complex list of commands. An experienced router support engineer can configure or troubleshoot a router without too much difficulty, but for someone new to routers, there is a steep learning curve.
- **More complex:** While an individual router may not actually be that much more complex than its software-based counterpart, when scaling to large networks, a hardware router environment can rapidly become very complex. This issue would also apply to a software router, but software routers are not common in the real world. In most network environments, hardware routers are used almost exclusively, with software routers being reserved for only the smallest networks or locations.

HOW DOES ROUTING WORK?

When a router receives a packet that must be forwarded to a destination host, the router needs to make a decision. It needs to determine if it can deliver the packet directly to the destination host, or whether it needs to forward the packet to another router. To make this decision, the router examines the destination network address. If the router has an interface that is connected to the same network as the destination host, it can deliver the packet directly. Where it gets interesting is when the router is not connected to the same network as the destination host and it needs to determine the best route to the destination host so it can forward the packet correctly.

When one router needs to forward a packet to another router, it uses the information in its routing tables to choose the best path for forwarding the packet. The decision as to which router to forward the packet to is determined by a number of variables about each of the network paths to the destination host, including the number of hops, the cost of each hop, and so on.

TAKE NOTE*

Just because there is a route to a destination, there is no guarantee there is a route back. While not a common problem in networks with dynamic routing enabled, it can happen, particularly when working in a heavily firewalled network environment.

When a router needs to forward a packet to another router for delivery to a remote network, it consults a database of information known as the routing table. This database is stored in the router's memory to ensure this lookup process can be performed very quickly. As the packet travels across the network toward its destination, each router along the way decides about where to forward the packet by consulting its routing table. When a destination host sends a reply packet, it is possible that the same path may not be used to reach the original

sender. This depends on the metrics of each path along the route. In other words, the way to the destination host may not be the best path back.

Information in the routing table can be generated in one of two ways. The first method is to manually configure the routing table with the routes for each destination network. This is known as static routing. Static routing is more suited to small environments where the amount of information to configure is small, and the overhead of generating the routing information dynamically is unacceptable. Static routers do not scale well to large or frequently changing internetworks because of the requirement for manual administration.

The second method for generating routing table information is to make use of a dynamic routing protocol. Because dynamic routing protocols are quite a bit more complex than static routing, we need to take a more in-depth look at the subject.

A general definition of a protocol is an agreed-upon method for exchanging data between two devices. A routing protocol defines the method for exchanging routing information between two routing devices. A dynamic routing protocol is used to exchange routing information that is created and maintained in the routing table automatically. When using a dynamic routing protocol, routing information is exchanged between routers to update the information kept in their routing tables. This can be done either periodically (at scheduled intervals) or on demand. If set up correctly at the outset, dynamic routers will require little administration after they have been configured, outside of ensuring software updates are applied in a timely fashion. Because they learn routing information dynamically, and can route around failures when the network architecture will support it, dynamic routing is generally used in large network environments where it would not be practical to use static routing.

TAKE NOTE*

Remember—routers need to be patched too. Because routers run an operating system, there are security and functionality updates that need to be applied on a regular basis.

UNDERSTANDING ROUTING PROTOCOLS

Routing protocols are based either on a distance vector or link-state algorithm. The differences between the two algorithms include when routing information is exchanged, what information is sent when the routing information is exchanged, and how quickly the protocol can route around outages, when the network topology will support it.

Path selection involves applying a routing metric to multiple routes, in order to select the best route. Some of the metrics used are bandwidth, network delay, hop count, path cost, load, reliability, and communication costs. The hop count is the number of routers traversed by a packet between its source and destination.

Distance vector-based routing protocols require that each router inform its neighbors of its routing table. This is done by sending the entire routing table when the router boots, and then at scheduled intervals afterward. Each router receives changes from its neighboring routers and then updates their own routing tables based on the information they receive. Using the information from these updates, a router can build a network map in its routing table, and determine hop counts (the distance to any network) for each network entry in the routing table. RIP is an example of a distance vector-based routing protocol, and is supported by Windows Server 2016. Routing updates sent using a distance vector-based routing protocol are unacknowledged and unsynchronized, which is one of the drawbacks of these protocols. Some other drawbacks of this type of routing protocol include the following:

- **High overhead:** Because every router on the network sends its entire routing table when it sends an update, distance vector-based protocols produce very large routing tables. This adds overhead to the router memory needed to store these tables, and the router processing power needed to maintain these tables. Large routing tables can also hamper an administrator trying to determine the source of an issue when problems arise.
- **Not scalable:** Distance vector-based networks are limited to 15 hops (router traversals) for any given route. In a large network (like the internet) it is very easy to have network segments that are greater than 15 hops away, and these would be unreachable in a distance vector-based network.

- **Network bandwidth intensive:** Distance vector-based protocols require that routers exchange their entire routing table whenever they do an update. On a large network, with large routing tables, these routing updates can utilize significant amounts of bandwidth, especially across slower WAN connections or demand-dial links.
- **Long convergence time:** Convergence is the amount of time it takes a routing algorithm to detect and route around a network failure. Distance vector-based protocols typically have longer convergence times than link-state-based protocols.
- **Routing loop issues:** Distance vector-based protocols can also suffer from routing loop issues when there are multiple paths to a network.
- **Count-to-infinity issues (routing loops):** Count-to-infinity issues occur when there is a network outage, and the routing algorithm cannot calculate a new router. One router will broadcast a route and increment the hop count for the router, then a second router will broadcast the same route to the first router, also incrementing the hop count, and so on, until the route metric (hop count) reaches 16 and the route is discarded.

Distance vector-based routing protocols have additional mechanisms that allow them to avoid the count-to-infinity issues as well as improving convergence. These mechanisms include the following:

- **Split horizon:** The split horizon mechanism prevents routes from being broadcast out the interface from which they were received. Split horizon eliminates count-to-infinity issues and routing loops during convergence in single-path internetworks and reduces the chances of count-to-infinity issues in multipath internetworks.
- **Split horizon with poison reverse:** The split horizon with poison reverse mechanism allows routes to be broadcast back to the interface from which they were received, but they are announced with a hop count of 16, which indicates that the network is unreachable (in other words, the route has been poisoned and is unusable through that interface).
- **Triggered updates:** Triggered updates allow a router to announce changes in metric values almost immediately, rather than waiting for the next periodic announcement. The trigger is a change to a metric in an entry in the routing table. For example, networks that become unavailable can be announced with a hop count of 16 through a triggered update. If triggered updates were sent by all routers immediately, each triggered update could cause a cascade of broadcast traffic across the IP internetwork.

The advantages of distance vector-based routing are that it requires low maintenance and is easy to configure, making it popular in small network environments.

Link-state routing was designed to overcome the disadvantages of distance vector-based routing. Routers using link-state routing protocols learn about their network environment by "meeting" their neighboring routers. This is done through a "hello" packet, which tells the neighboring router what networks the first router can reach. Once the introduction is complete, the neighboring router will send the new network information to each of its neighboring routers using a link-state advertisement. Open Shortest Path First (OSPF) is an example of a link-state routing protocol. The neighboring routers copy the contents of the packet and forward the link-state advertisement to each attached network, except for the one on which the link-state advertisement was received. This is known as flooding.

Routers using a link-state routing protocol build a tree, or map, of shortest paths with itself as the root. The tree is based on all the link-state advertisements seen. The tree contains the route to each destination in the network. Once this tree has been built, routing information is only sent when changes to the network occur, instead of periodically as in the distance vector-based protocols.

There are a few advantages to this method, especially when compared to the distance vector-based routing protocols. These advantages include the following:

- **Smaller routing tables:** Because the router only maintains a table of link states, rather than a copy of every route on the network, it needs to maintain much smaller routing tables.
- **Highly scalable:** Link-state protocols do not suffer from the 15-hop issue that distance vector-based protocols do, so they are able to scale to much larger networks.
- **More efficient use of network bandwidth:** Because link-state information is not exchanged after the network has converged, routing updates do not consume precious bandwidth, unless there is an outage that forces the network to re-converge.
- **Faster convergence:** Link-state routing protocols will converge faster than distance vector-based protocols, because updates are sent as soon as a change to the network occurs, instead of having to wait for the periodic update used in the distance vector-based protocols.

The disadvantages of link-state-based protocols are that they are more complex to understand and configure than distance vector-based protocols. They also require additional processing power on the router, due to the need to calculate the routing tree.

Routing can be a key component of network security because it lets you determine which parts of the network can be accessed by other parts of the network. For example, if you have a business partner connection to a third-party network, the third-party network will need to have routing information in order to access any systems that have been advertised on your extranet DMZ for them to access. While a firewall is the best way to secure this connection, you can add an additional layer of security by restricting the routing available to the third party. In other words, if you only tell the third party's network the routes to the extranet, they will not be able to send packets to those parts of your network where they should not have access.

Understanding Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS)

Two other security technologies available to secure networks are ***intrusion detection systems (IDS)*** and ***intrusion prevention systems (IPS)***. An IDS is a solution designed to detect unauthorized user activities, attacks, and network compromises.

An intrusion prevention system (IPS) is very similar to an IDS, except that in addition to detecting and alerting, an IPS can also take action to prevent the breach from occurring.

There are two main types of IDS/IPS technologies:

- **Network-based:** A network-based IDS (NIDS) monitors network traffic using sensors that are located at key locations within the network, often in the demilitarized zone (DMZ) or at network borders. These sensors capture all network traffic and analyze the contents of individual packets for malicious traffic. An NIDS will gain access to network traffic by connecting to a hub, a network switch configured for port mirroring, or a network tap.
- **Host-based:** A host-based IDS (HIDS) generally has a software agent that acts as the sensor. This agent monitors all activity of the host on which it is installed, including monitoring the file system, the logs, and the kernel, to identify and alert upon suspicious behavior. An HIDS is typically deployed to safeguard the host on which it is installed.

There are two common deployment methodologies used when placing an IDS/IPS to protect a network from the internet. Each has its own advantages and disadvantages:

- **Unfiltered:** An unfiltered IDS/IPS installation examines the raw internet data stream before it crosses the firewall. This provides the highest amount of visibility to attacks, but also means that there is a significantly higher volume of data to be monitored, with a higher possibility of false positives. There is also a chance that during periods of high traffic, the IDS/IPS might not be able to process all the packets, and attacks can be missed.
- **Screened:** A screened IDS/IPS solution monitors the traffic that gets through the screening firewall. The advantage to this model is it dramatically reduces the amount of traffic that needs to be monitored, reducing the chances of false positives and lost packets during high traffic volumes. There is a loss of visibility with this model, because attacks cannot be seen on the screening firewall.

TAKE NOTE*

IDS and IPS are solutions that have historically been used to secure or provide alerts on internet connections, because those connections have typically presented the largest threat to the network. However, with the interconnectivity of networks beyond the internet and the threat of an insider attack, it may make sense to deploy IDS/IPS in strategic locations on your internal network. Give this some serious consideration if your internal network has connections to third-party networks like customers, vendors, or business partners.

Understanding Honeypots

Honeypots, honey nets, and padded cells are complementary technologies to IDS/IPS deployments. A ***honeypot*** is a trap for hackers. A honeypot is designed to distract hackers from real targets, detect new vulnerabilities and exploits, and learn about the identity of attackers. A ***honey net*** is just a collection of honeypots used to present an attacker with an even more realistic attack environment. A padded cell is a system that waits for an IDS to detect an attacker and then transfers the attacker to a special host where they cannot do any damage to the production environment. They are all related technologies, used to add an additional layer to your security infrastructure.

A honeypot is valuable as a surveillance and early-warning tool. It is also a generic term to describe anything that would attract an attacker. While it is usually a reference to a host running special software for detecting and analyzing an attack, the term honeypot can refer to other things such as files or data records, or even unused IP address space.

There are a variety of different types of honeypots, including the following:

- **Production:** A production honeypot is a relatively easy solution to deploy. They are used to distract attackers from potentially vulnerable production systems, and are relatively easy to use. A production honeypot typically captures limited information, and can generally be found in corporate networks. This type of honeypot is typically used as an additional form of early-warning system, as an enhancement to an IDS/IPS system.
- **Research:** A research honeypot is more complex than a production honeypot, and is more difficult to deploy and maintain. This type of honeypot is used to capture extensive information which is used to develop attack signatures, identify new attack techniques and vulnerabilities, and develop a better understanding of the attacker's mindset. Research honeypots are used primarily for research by universities, the military, or other government organizations.

When deploying a honeypot, ensure that there is no production information or purpose for the server. This not only ensures that production data is secure, but because there is no legitimate reason for traffic or activity on the system, the only thing that will touch the honeypot will be malicious activity.

Also, be aware that honeypots can create risks to the environment. Because a honeypot is essentially being used as bait for an attacker, attackers are being lured into the network environment. As a result, be absolutely certain that the honeypots are isolated from the production environment. If they are not, the attacker can jump from the honeypot to the production environment and compromise critical systems or infrastructure. It's somewhat like trying to lure a bear to an adjoining empty campsite, to keep them away from yours—there's always a chance the bear may find your campsite, anyway.

One area where honeypots are especially useful is in the battle against spam. One of the challenges associated with spam and spam filtering is that the spammers are constantly changing the techniques they use to bypass spam filters. They also have a variety of techniques for harvesting email addresses from websites for inclusion in their spam target lists.

As a result, the people who develop spam filters spend much of their time working to identify those techniques and develop new filters to combat the new spam methods. Honeypots are an essential component of this fight, and there are two types of honeypots that can be used to combat spam:

- **Email address honeypot:** Any email address which is dedicated to receiving spam for analysis can be considered a spam honeypot. An example of this technique is Project Honey Pot, a distributed, open-source project that uses honeypot pages installed on websites around the world in conjunction with uniquely tagged email addresses for analyzing not only spam delivery, but also the email address harvesting techniques.
- **Email open relay honeypot:** Email open relays are servers that relay mail from one mail server to another mail server. An example of a mail relay server is using POP3 or IMAP to send an email through your personal ISP. In some instances, these servers are set up so they do not need credentials to send email, which is a significant prize for spammers. It allows them to relay their millions of spam emails anonymously. Setting up a honeypot that appears to be an open relay can potentially reveal the spammer's IP address and provide bulk spam capture. This allows for in-depth analysis of the spammers techniques, response URLs, email addresses, and other valuable information.

While these are all extremely exciting technologies, they do not get deployed in too many corporate environments. Generally, these are deployed by educational institutions and security research firms. Corporate information security professionals are so busy securing their environment from attacks that they don't spend a lot of time researching attack patterns. As long as the attack didn't succeed, they are satisfied. In cases of high-security environments, where there is extensive internet-based activity and data requiring additional layers of security, honeypots may be used as part of the layered security defense.

Understanding DMZ

When most people hear the term ***DMZ*** (short for demilitarized zone), images of barbed wire and machine gun emplacements come to mind. While not entirely accurate in the scope of information security, the concept is not that far from reality. In computer networking, a DMZ is a firewall configuration used to secure hosts on a network segment. In most DMZs, the hosts on the DMZ are connected behind a firewall which is connected to a public network like the internet. Another common configuration is to have the firewall connected to an extranet, with connections to customers, vendors, or business partners. DMZs are designed to provide access to systems without jeopardizing the internal network.

There are two typical DMZ configurations encountered in production environments:

- **Sandwich DMZ:** In a sandwich DMZ model (see Figure 4-3), there is an outer firewall and an inner firewall. The outer firewall secures the DMZ network segment from the external (insecure) network. Servers that are meant to be accessed from the external network (like the internet) have the appropriate rules configured to permit secure access.

The inner firewall is used to add an additional layer of security between the servers on the DMZ and the internal (secure) network. The main benefit of this model is that in the event that the outer firewall and/or a server on the DMZ is compromised, there is an additional layer of firewall security protecting the internal network. Ideally, the outer and inner firewalls are from different vendors, in order to ensure that in the event an exploit is used to compromise the outer firewall, the same exploit cannot be used to compromise the inner firewall. The major drawbacks of this model are that it is a more complex architecture to implement and maintain, and it is more expensive, because additional training is needed for the different additional firewall.

Figure 4-3

An example of a sandwich DMZ

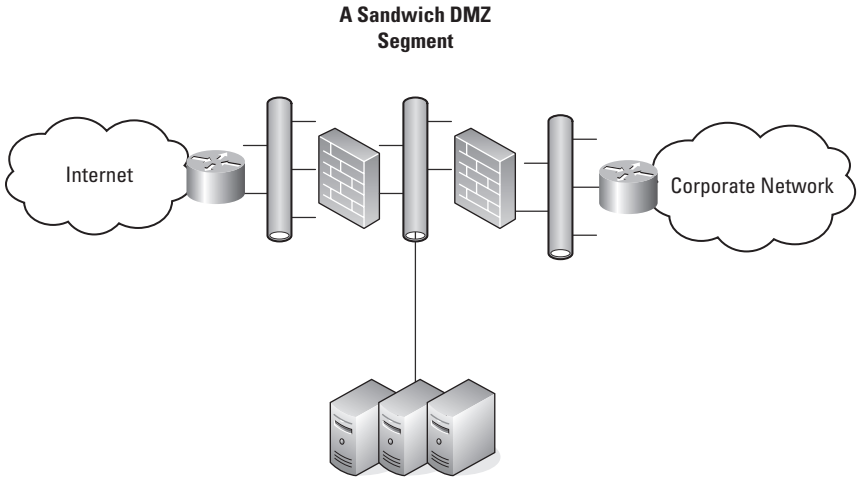

- **Single Firewall DMZ:** In a Single Firewall DMZ (see Figure 4-4), the DMZ is an additional network connection from the firewall. This provides an external network connection, an internal network connection, and a DMZ network connection, all connected to the same firewall. While this architecture still allows the firewall to control access to DMZ resources, if the firewall is compromised, access to the internal network may be breached. This model is less expensive than the sandwich model, but does not provide as high a level of security.

Figure 4-4

An example of a Single Firewall DMZ

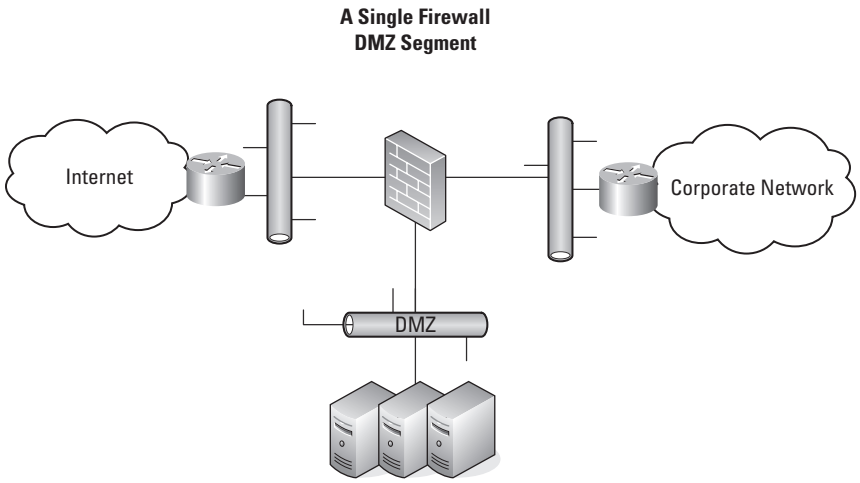

While it's easy to talk about the architecture of a DMZ, it's important to understand what types of servers and services might be placed on a DMZ. Some of the most common include the following:

- **Web servers:** Web servers are the most common servers found in DMZ networks. Accessed using HTTP over port 80 or HTTPS over port 443 for secure access, web servers are commonly internet-accessible. All web servers accessed on the internet are hosted on a DMZ somewhere. Web servers add an additional layer of complexity because many web applications need to communicate with an internal database or databases to provide some specialized services. A database would not be placed on the DMZ because it should not be accessed from the insecure network (the internet). An example of this might be an e-commerce application. When reaching the website, the catalog data, including product descriptions, prices, and availability are contained in the database (sometimes referred to as the backend database). If the database server contains critical information like Social Security numbers, financial information, credit card data, and so on, it's a good idea to add an application firewall between the web server and the database server. While this increases the cost and complexity of a solution, it adds an extra layer of security to protect the database.
- **Email relay servers:** Email servers are another type of server that needs to be accessed from the internet, to allow for sending and receiving internet email. In the early years of computer networking, it was not unusual for email to be restricted to the corporate network. Once companies and individuals got connected to the internet, the ability to send and receive email from other companies over the internet became critical to business success. By placing email relay servers, which communicate on port 25, on the DMZ, they can receive email from the internet, and then relay it to mail servers on the internal network securely. Spam filtering capabilities are frequently included on these relay servers.
- **Proxy servers:** Proxy servers are used to proxy, or act as an intermediary, for user requests from the internal network to the internet, and are typically used to retrieve website information. These are placed on the DMZ to provide additional security for web browsing. Some proxy servers will filter content, including inappropriate websites, add virus protection and anti-spyware security, and even improve performance by caching web requests.
- **Reverse proxy servers:** Reverse proxy servers are used to provide secure access to internal applications from an insecure network. While these have largely been replaced by VPN technologies, reverse proxy servers can be used to provide employees with access to web-based email servers on the internal network, provide access to internal web applications and, in some cases, even provide secure terminal services connections to the internal network.

Understanding NAT

Network Address Translation (NAT) is a technique used to modify the network address information of a host while traffic is traversing a router or firewall. This technique is used to hide the network information of a private network, while allowing traffic to be transferred across a public network like the internet.

NAT was originally created as a workaround for IP addressing issues. The internet relies on the TCP/IP protocol suite for communications between hosts. A critical component of this protocol suite is the IP addressing. The explosive growth of the internet threatened to exhaust the pool of IPv4 IP addresses, which would have crippled the expansion and use of the internet. Without unique addresses, the internet would be unable to successfully route TCP/IP traffic. NAT was the resultant workaround solution for preserving the number of IP addresses used on the internet.

In the early days of the internet, when the TCP/IP protocol and related addressing was being developed, the 32-bit addressing scheme (known as IPv4) was considered more than adequate for any potential network growth. Technically there were 4,294,967,296 unique addresses available using a 32-bit address, and even discounting the reserved ranges, there are still over 3 billion possible addresses. At the time, that was enough addresses to provide an address for every person on the planet, including children. Unfortunately, the designers of the addressing scheme dramatically underestimated the explosive growth of the internet, as well as the widespread adoption of TCP/IP in business and home networks, resulting in the depleting of IP addresses.

The practical use for NAT is that it allows the use of one set of IP addresses on the internal LAN, and a second set of IP addresses for the internet connection. There is a device (usually a router or firewall) in between the two networks that provides NAT services, managing the translation of internal addresses to external addresses. This allows companies to use large numbers of unregistered internal addresses while only needing a fraction of that number of addresses on the internet, thus conserving the addresses. This allows for the re-use of addresses within private networks while ensuring that the addresses used on the internet remain unique.

TAKE NOTE*

Network Address Translation (NAT) is supported under Windows Server 2016 by the Routing and Remote Access Service.

The long-term solution for this issue is IPv6 or Internet Protocol Version 6, the next generation protocol for the internet. It's designed to provide several advantages over IPv4, including support for addresses that are 128 bits long. This permits 2^{128} unique IPv6 addresses, or over 340 trillion addresses.

However, the adoption of IPv6 has been slow, in large part to the successful use of NAT and proxy servers to conserve the number of IPv4 addresses used on the internet today.

There are two main types of NAT:

- **Static NAT:** Static NAT maps an unregistered IP address on the private network to a registered IP address on the public network, using a one-to-one basis. This is used when the translated device needs to be accessible from the public network. For example, a web server on a DMZ network might have an unregistered address of 10.20.30.40 that is translated by a NAT-capable device to an internet-facing address of 12.4.4.234. A user trying to connect to that website can enter 12.4.4.234, and the router or firewall at the other end will translate that address to 10.20.30.40 when the packet reaches it. This version of NAT is typically used in conjunction with DMZ or extranet networks.
- **Dynamic NAT:** Dynamic NAT maps an unregistered IP address on the private network to a registered IP address that is selected by the routing device providing the NAT service from a pool of registered IP addresses. This is more commonly used when many hosts on the internal network need to access the internet and don't have a requirement for a static address. The workstation's address is translated to the next available registered address in the pool as soon as it initiates a connection to the public network.

There are two major security implications associated with the use of NAT. First, NAT can be used to hide private network addresses, which makes it more difficult for an attacker to successfully penetrate a private network. The addresses that are visible to an internet-based attacker are the NAT addresses typically stored on the firewall, which should be one of the more secure devices on a network.

NAT also presents a unique issue when working with the IPsec protocol, which we will be discussing in more detail later in the lesson. Early implementations of IPsec did not support NAT, so the IPsec protocol could not be used when NAT was enabled in the environment. NAT traversal capability was added in later versions of the IPsec protocol, but IPsec still requires that some special steps be taken in order to successfully work with NAT.

Understanding VPN

VPN (Virtual Private Network) is a technology that uses encrypted tunnels to create secure connections across public networks like the internet. There are a variety of uses for this technology—three of the most common uses appear in Figure 4-5.

Figure 4-5

Some common uses for VPN

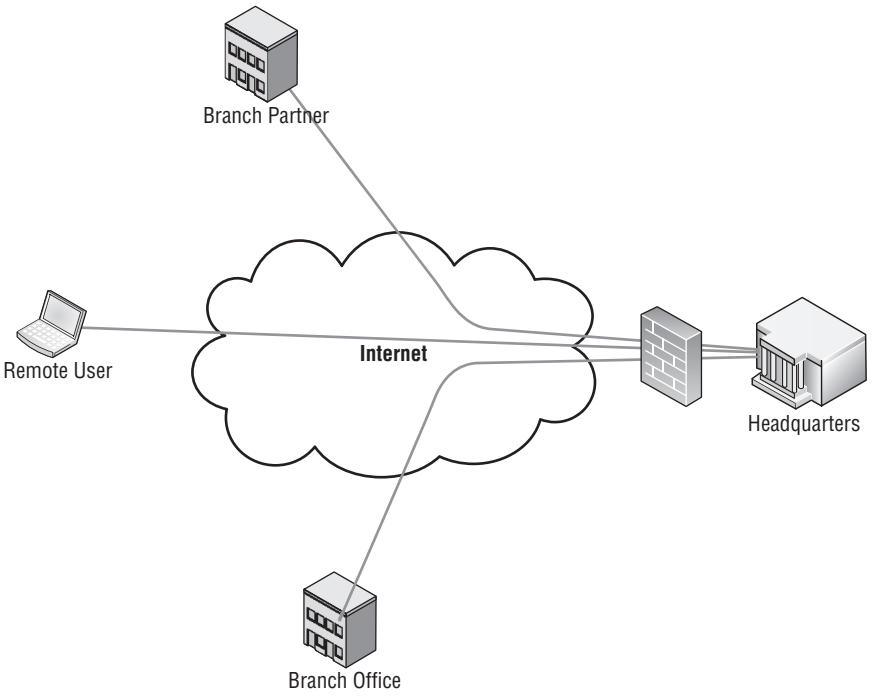

VPNs are commonly used by remote employees for access to the internal network, to create secure network-to-network connections for branch offices or business partner connections, or even to create secure host-to-host connections for additional security and isolation on an internal network. VPNs utilize encryption and authentication to provide confidentiality, integrity, and privacy protection for data.

Remote access VPNs were first introduced in the late 1990's, and were initially used in conjunction with modems to provide more secure, more flexible connectivity to a corporate network. All that was required was a dial-up internet connection and a VPN client, and a user could connect to the corporate network over an encrypted connection. No more modem banks in the data center, and no more toll-free modem lines to be managed. A user who could get to the internet could get remote access up and running.

With the advent of high speed internet connections, the use of VPN technologies exploded. It was now possible in some cases to get a faster connection via a high-speed home internet connection than typical dedicated network connections from branch offices. It also allows businesses to migrate from expensive dedicated network connections to less expensive internet-based VPN connections.

The first standards-based VPNs were based on the IPsec protocol. The IPsec-based VPNs quickly overtook some of the proprietary-based VPNs that were the first products to market.

Understanding Other VPN Protocols

While IPsec can be considered the predominant protocol associated with VPNs, there are other protocols that can also be used to build VPNs, or provide VPN-like connectivity.

SECURE SOCKETS LAYER (SSL)/TRANSPORT LAYER SECURITY (TLS)

One of the key VPN protocols used today is SSL/TLS, which is the main alternative to IPsec for implementing a VPN solution.

The SSL protocol standard was originally proposed as a standard by Netscape. While this protocol is widely used to secure websites, it has since been formalized in the IETF standard known as Transport Layer Security (TLS). The SSL/TLS protocol provides a method for secure client/server communications across a network and prevents eavesdropping and tampering with data in transit. SSL/TLS also provides endpoint authentication and communications confidentiality using encryption.

HTTPS, the secure version of HTTP web browsing, uses the SSL protocol. This protocol provides 128-bit encryption, and is currently the leading security mechanism for protecting web traffic including banking, e-commerce, secure email, and essentially any other secure website that might be encountered.

In typical end-user/browser usage, SSL/TLS authentication is one way. Only the server is authenticated—the client compares the information entered to access a server to information on the SSL certificate on the server. Thus, the client knows the server's identity. However, the server does not do this for the client—the client remains unauthenticated or anonymous.

SSL/TLS can also perform bi-directional authentication by using client-based certificates. This is particularly useful when using this protocol to access a protected network, as it adds an additional layer of authentication to the access.

As we discussed in the section on IPsec, a VPN creates a secure tunnel through a public network like the internet. While SSL VPNs still leverage the concept of tunneling, they create their tunnels differently than IPsec. An SSL VPN establishes connectivity using the SSL protocol. IPsec works at Layer 3 of the OSI model, while SSH functions at Layers 4-5. SSL VPNs can also encapsulate information at Layers 6-7, which makes SSL VPNs very flexible.

One additional function of an SSL VPN is that an SSL VPN usually connects using a web browser, whereas an IPsec VPN generally requires that client software be installed on the remote system.

SSL VPNs are predominantly used for remote access VPN connections, where a client is connecting to applications on an internal network, as opposed to a site-to-site connection, where two gateways are used to connect disparate private networks across the internet.

Some benefits of SSL/TLS VPNs over IPsec VPNs include:

- **Less expensive:** Because an SSL VPN is typically clientless, there aren't the costs for rolling out, supporting, and updating client software.
- **Platform independent:** Because the access to an SSL VPN is granted through the standard SSL interface, which is a component of virtually every web browser, virtually any OS that runs a browser is supported.
- **Client flexibility:** Generally, IPsec clients are usually installed only on corporate systems. Due to the additional configuration flexibility, SSL VPNs can be configured to allow access from a variety of clients, including corporate systems, home systems, customer or supplier systems, or even a kiosk machine in a library or an internet cafe. This wider access can greatly increase employee satisfaction.

- **NAT support:** Historically Network Address Translation (NAT) can cause issues with IPsec VPNs. Virtually all IPsec vendors have created workarounds for this issue. An SSL VPN doesn't have these issues, because SSL works at a higher layer than IPsec, and thus is not impacted by NAT.
- **Granular access control:** This could be considered a benefit or, depending on the environment, it could be a drawback. SSL VPNs require a greater granularity of access than a typical IPsec VPN, because instead of creating a tunnel from the host to the internal network, SSL VPNs require that each resource accessed be explicitly defined. The upside is unless it has been explicitly defined, an SSL VPN user cannot access it, which has significant security benefits, but in a complex environment this could add significant overhead to VPN support.
- **Fewer firewall rules required:** In order to access an IPsec gateway across a firewall, open several ports to support the individual protocols for authentication and the tunnel. An SSL VPN only needs port 443 opened, which is generally easy to do, due to the prevalence of the HTTPS protocol.

SECURE SHELL (SSH)

The Secure Shell (SSH) protocol is a protocol for secure remote logon and other secure network services over the network. SSH can be used for several applications across multiple platforms, including UNIX, Microsoft Windows, Apple Mac, and Linux.

Some of the applications supported with SSH include the following:

- Secure logon
- Secure remote command execution
- Secure file transfer
- Secure backup, copy, and mirroring of files
- Creation of VPN connections (when used in conjunction with the OpenSSH server and client)

The SSH protocol consists of three major components:

- **Transport layer protocol:** This provides server authentication, confidentiality, and integrity with perfect forward secrecy.
- **User authentication protocol:** This provides authentication of the client to the server.
- **Connection protocol:** Multiplexes the encrypted tunnel into several logical channels.

Now that we've looked at some of the protocols that can be used to secure traffic across a network, and usually across a public network like the internet, we can look at a technique for providing additional security on an internal network.

Understanding Server and Domain Isolation

Security professionals are constantly being asked by businesses to allow greater and greater access to resources in order to facilitate business requirements. While wider and easier access to resources can increase the production of a business, it also presents significant security challenges. The risk of virus attacks, rogue users and devices, and unauthorized access to sensitive information associated with unauthorized or unmanaged devices are enough to keep an information security professional awake at night.

An example of this might be a developer's workstation. Many developers feel they have unique requirements to do their job and as a result, they may run custom configurations, unsupported operating systems, open source applications, and not participate in the corporate patch and configuration management programs. In a typical environment, once this system is

TAKE NOTE*

To leverage isolation in an environment, be sure to take the time to do the appropriate planning. This can be a complex implementation and needs to be understood before protocols are enabled.

on the network, it would have access to any internal resources. Server and Domain isolation provides some additional security options.

Server and Domain isolation is a solution based on IPsec and the Microsoft Active Directory that enables administrators to dynamically segment their Windows environment into more secure and isolated logical networks. These logical networks are segmented based on policy and can be accomplished without needing to deploy firewalls, implement VLANs, or make other changes on the network. Through the use of authentication and encryption, internal servers and domains can be secured. This creates an additional layer of policy-driven protection, and provides another alternative to the security controls we have discussed thus far in this lesson.

Figure 4-6 provides an example of Server and Domain isolation. The isolated network can only be accessed by computers with the appropriate IPsec and Active Directory configuration.

Figure 4-6

An example of Server and Domain isolation

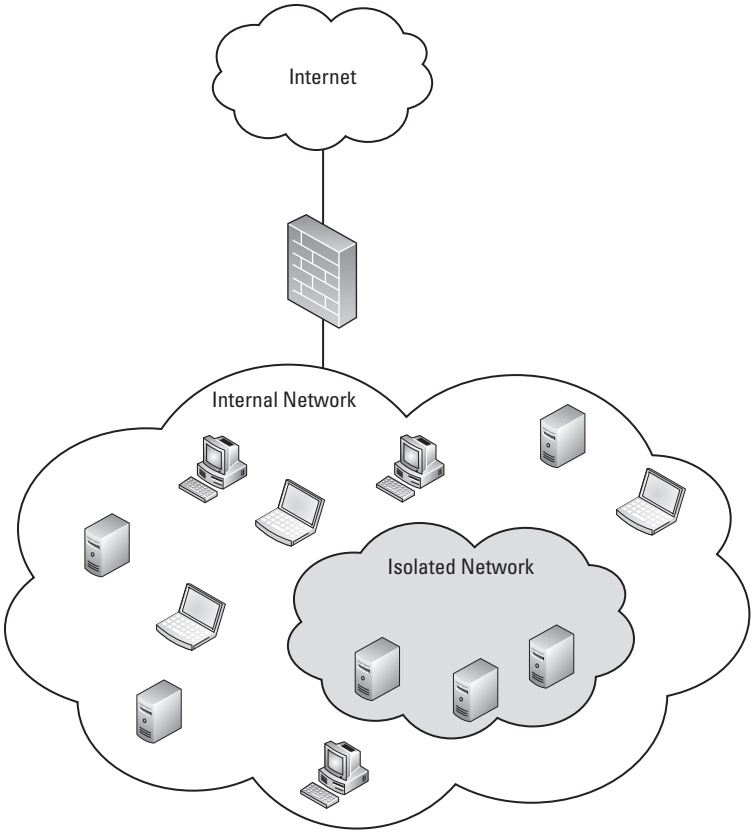

How does it work? Authentication to the isolated environment is based on the computer's machine credentials. The machine credentials can be an Active Directory issued Kerberos ticket or it can be an X.509 certificate automatically distributed to the computer by a Group Policy. Once the machine has authenticated, the associated isolation policies are enforced by the built-in IPsec functionality in Windows.

Recall that IPsec supports two modes. Tunnel mode is the most frequently used mode, because it supports the widely used remote access and site-to-site VPN solutions that are becoming ubiquitous in the corporate world. Transport mode is used for Server and Domain isolation, as it is the mode that supports secure host-to-host communications.

■ Protecting Data with Protocol Security

THE BOTTOM LINE

In this lesson, we have discussed several security protocols such as IPsec, SSL/TLS, and SSH. In this section, we are going to look at a couple more protocols that can be used to secure data. This includes looking at protocol spoofing, network sniffing, and some of the common attacks that might be encountered when working on securing a corporate computing environment.

CERTIFICATION READY
Which protocol can be used to protect confidential data from being sent between servers?
Objective 3.3

One of the more challenging topics for any information security professional to tackle is the idea of protocol security. This is an area which has long been the area of networking professionals, and while there is an obvious overlap between networking and information security, understanding protocol security can be a real challenge for information security professionals both new and old. In order to get an understanding of how network protocols can impact security, we need to start the discussion by looking at tunneling.

Understanding Tunneling

Tunneling is defined as the encapsulation of one network protocol within another. Tunneling can be used to route an unsupported protocol across a network, or to securely route traffic across an insecure network. VPNs use a form of tunneling when data is encapsulated in the IPsec protocol.

An example of tunneling that is used to move unsupported traffic across a network is the Generic Routing Encapsulation (GRE) protocol. GRE is an IP-based protocol frequently used to carry packets from unrouteable IP addresses across an IP network.

In order to understand why the GRE protocol is used, we need to discuss a little about IPv4 addressing. A component of the IPv4 addressing scheme is a set of addresses known as either private or reserved address ranges. These ranges include 10.0.0.0-10.255.255.255, 172.16.0.0-176.31.255.255, and 192.168.0.0-192.168.255.255. These ranges were assigned to help delay the exhaustion of the available IPv4 IP addresses, and are typically used for both home and office networks, where there is not a requirement for the addresses to be routed across a public network like the internet. These networks generally use NAT to permit internet access.

Another area where these addresses are used is for lab/development networks in an enterprise environment. Sometimes there is a requirement to route traffic from the lab/development network to another, but because these networks are addressed with private addresses, they may not be routable across the enterprise network. This is when GRE becomes useful. Traffic between the labs can be encapsulated in a GRE tunnel which can be routed over the enterprise network without requiring readdressing.

TAKE NOTE*

What is PPP? PPP, the Point-to-Point Protocol, was a protocol defined in the late 90's that provided a standard transport mechanism for point-to-point data connections. This was largely used in conjunction with modem connections, and has largely been phased out as modem connections have been replaced by high-speed internet connections.

PPTP (Point-to-Point Tunneling Protocol) is a proprietary VPN protocol originally developed by the PPTP Forum, a group of vendors that included Ascend Communications, Microsoft Corporation, 3Com, ECI Telematics, and U.S. Robotics. PPTP was designed as an extension of the Point-to-Point Protocol (PPP) to allow PPP to be tunneled through an IP network. At one time PPTP was the most widely used VPN protocol, but the release of IPsec had a significant impact on PPTP's use.

Another tunneling protocol which was widely used was L2TP (Layer 2 Tunneling Protocol), which combined the best features of PPTP and the L2F (Layer Two Forwarding) protocol, an early competing protocol for PPTP that was developed by Cisco Systems. Like PPTP, L2TP was designed as an extension of PPP to allow PPP to be tunneled through an IP network. L2TP support was first included in a Microsoft server product with the release of Windows 2000 Server. Prior to Windows 2000, PPTP was the only supported protocol. A number of hardware VPN vendors, including Cisco, also supported PPTP.

Understanding DNS Security Extensions (DNSSEC)

Anyone who has ever connected to a website by name has used the Domain Name Service (DNS). The DNS is a service used on the internet for resolving fully qualified domain names (FQDN) to their actual Internet Protocol (IP) addresses, using a distributed network of name servers. When entering a server name, like www.bing.com, DNS ensures the connection is directed to the appropriate servers. Although this service is largely invisible to end users, DNS is a critical element of how the internet functions.

Let's say you want to check the scores from your favorite sport on the Bing website. Before DNS, when asking "What's the address of the Bing website?" the answer might be 204.79.197.200. Most people might remember that number for less than 30 seconds and then would probably never find those sports scores. With DNS, type the server name to go to a site such as www.bing.com, and the DNS infrastructure of the internet will translate the name to the correct address. It's like a big phone book—put in a name to find the correct number.

However, DNS was developed during the early years of the internet, when functionality was the goal, not security. As a result, DNS was built without security. In recent years, this lack of security has been exploited with forged DNS data, which, among other things, redirects connections to malicious websites. After typing the address of your bank, it appears you have reached your destination. You enter your user ID and password to access your accounts, but can't log on. Next month you find out that your account has been cleaned out. What happened was that the initial connection was the result of a bad DNS entry. Instead of connecting to the bank's website, you connected to a clever duplicate, which captured your logon information and let the bad guys steal your life savings.

DNS Security Extensions (DNSSEC) adds security provisions to DNS so that computers can verify that they have been directed to proper servers. This new standard was published in March 2005, and is slowly being adopted by the internet domains. DNSSEC provides authentication and integrity checking on DNS lookups, ensuring that outgoing internet traffic is always sent to the correct server. This removes the issues of forged DNS data, because there is no way to forge the appropriate authentication. This not only addresses the issue of website redirection, but also addresses some challenges associated with spam and the use of faked mail domains.

DNSSEC provides authentication and integrity checking through the use of public key encryption. The domain name structure provides a hierarchy of authenticated keys, creating a chain of trust from the root of the DNS hierarchy to the domain being queried. DNSSEC will address many of the most problematic security issues associated with the internet's core infrastructure, but it comes at a significant cost. As with any large-scaled public key implementation, rolling this out to the entire internet will be an enormously complex, resource-intensive project. There are also challenges associated with maintaining the web of trust created by using public keys on such a large scale.

Understanding Protocol Spoofing

Another security area of concern with respect to protocols is the concept of protocol ***spoofing***. The word spoof can be defined as a hoax. Protocol spoofing is the misuse of a network protocol to perpetrate a hoax on a host or a network device. Some common forms of protocol spoofing include:

- ***ARP spoofing***: ARP (Address Resolution Protocol) spoofing (or ARP poisoning) is an attack on the protocol used to determine a device's hardware address (MAC address) on the network when the IP address is known. This is critical for the proper delivery of network data once the data has reached the proper LAN segment. An ARP spoofing attack occurs when an attacker modifies the network's ARP caches and takes over the IP address of the victim host. This permits the attacker to receive any data intended for the original host.
- ***DNS spoofing***: DNS spoofing occurs when an attacker is able to intercept a DNS request and respond to the request before the DNS server is able to. As a result, the victim host is directed to the wrong website, where additional malicious activities can take place. This attack is frequently used in conjunction with network sniffing, which will be discussed in the next section.
- ***IP address spoofing***: In an IP address spoofing attack, the attacker creates IP packets with a forged source IP address to either conceal the identity of the attacking host, or to impersonate the identity of a victim host. This attack was very popular in the early days of packet analysis firewalls—an attacker would spoof an internal IP address from the outside of a firewall, and if not configured correctly, the firewall would permit access to the internal network.

It is important to be aware that the term protocol spoofing has another definition within the computing arena. It's a term used to represent a technique associated with data compression and is used to improve network throughput and improve network performance. While a valuable tool in the appropriate circumstances, this form of protocol spoofing does not have information security implications.

Understanding Network Sniffing

Network sniffing is a type of network analysis that is a very useful tool for network administrators responsible for maintaining networks and identifying network issues. It involves connecting a device to the network with the appropriate software to allow access to the details of the packets traversing the network. Figure 4-7 shows an example of an Open Source network sniffing tool.

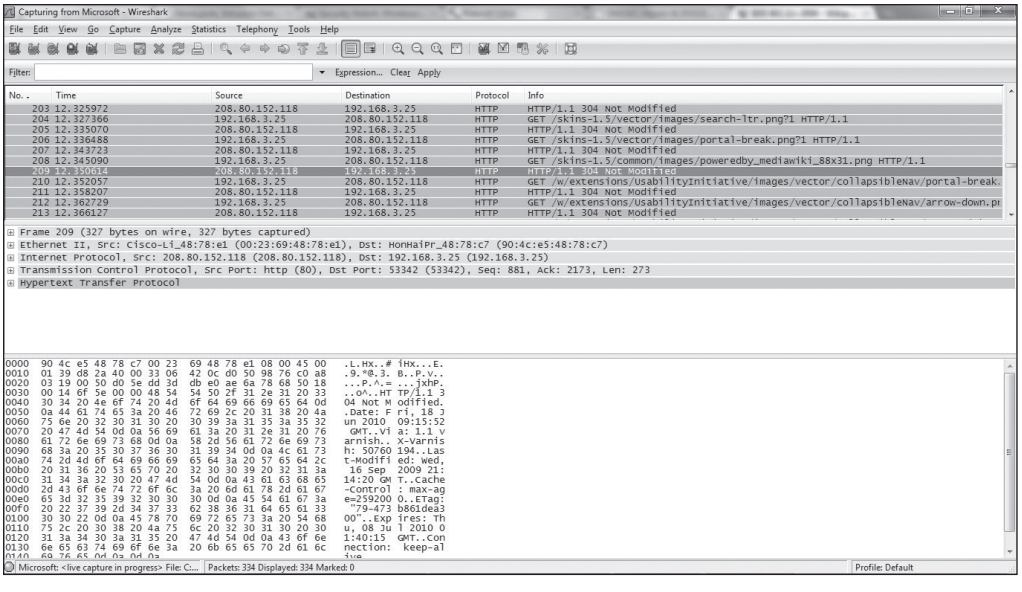

Figure 4-7

Wireshark—a commonly used open source network sniffing tool

As shown in the figure, this tool reveals a significant amount of information about the packet being analyzed. To a network administrator with an in-depth understanding of networking, this information can be used to identify application issues, network latency, and a variety of other network errors.

Unfortunately, to an attacker with similar skills, the information offered by network sniffing provides equally valuable information that can be used for attacking a network. For example, any data sent in clear text, that is, not encrypted, can generally be read directly from the network. In the early days of the internet, this was a significant amount of the traffic. Reading passwords from data packets was a trivial exercise.

Today, with the widespread use of encryption through secure websites and the use of VPNs for remote access, the risks presented by network sniffing are slightly mitigated, as the attacker can no longer read the data contents of a packet, but they can still get important information about the data packet that can be used in attacks.

It is important to be aware that a network sniffer can only see traffic that crosses the port to which it is connected. So, a sniffer placed on the LAN in a branch office cannot capture traffic from the headquarters network. In a switched environment, leveraging VLANs, the amount of traffic passing any one port can be limited. The ports that offer the most information are the ingress/egress points to the network, where all the traffic from the subnet is concentrated.

This means an attacker cannot directly capture traffic from your network, but that doesn't mean you're safe. A system on your internal network that is infected by a virus can end up running a network sniffer and providing the captured traffic to a remote host.

Another security challenge associated with a network sniffer is that they are passive devices. Unless the attacker has made modifications to the network in order to access more information, it is almost impossible to detect a network sniffer. In fact, there could be a network sniffer on a network node beyond your internal network that could be capturing packets about your internet access. In that case, there is no access to the network infrastructure to look for changes.

Be aware that wireless networks are particularly susceptible to network sniffing attacks, due to the lack of a port requirement. Once connected to a wireless network, an attacker has access to all the traffic on the network. It's an excellent idea to only use encrypted connections for anything done on a wireless network, beyond general web browsing.

Understanding Common Attack Methods

We have covered the information security challenges associated with computer networking throughout this lesson. The final piece of the network security puzzle relates to the types of attacks that can be expected when working to protect computer networks. While no list of attacks can be complete, if only because attackers are constantly coming up with new attacks, this list covers the major categories of attacks. Common attacks include the following:

- **Denial-of-service/*distributed denial-of-service (DoS/DDoS) attacks*:** The goal of a denial-of-service attack is to flood the network being attacked with overwhelming amounts of traffic, shutting down the network infrastructure like a router or firewall. Because the attacker isn't interested in receiving responses to their attack packets, DoS attacks are ideal opportunities for using spoofed addresses. Spoofed addresses are more difficult to filter, because each spoofed packet appears to come from a different address, and they hide the true source of the attack. This makes backtracking the attack extremely difficult. The new wrinkle to the DoS is the distributed DoS, which leverages botnets to generate DoS attacks from multiple sources. Not only does this make the attack more difficult to defend against, as multiple computers can generate significantly more traffic than a single computer, but it also makes it much more difficult to track down the source of the attack.

TAKE NOTE*

A botnet is a distributed network of computers that have been compromised by malicious software and are under the control of an attacker.

- **IP spoofing to bypass network security:** As we've discussed, IP spoofing is the modification of data packets so the data packets from the attacking computer appear to be from a trusted computer. By appearing as a trusted computer, the attacker is able to bypass network security measures, like a packet filter, or other solutions that rely on IP addresses for authentication. This method of attack on a remote system can be extremely difficult, because the attacker must modify thousands of packets in order to successfully complete the attack. This type of attack generally works best when there are trust relationships between machines. For example, it is not uncommon in some environment to have UNIX hosts on a corporate network that trust each other. Once a user successfully authenticates to one host, they are automatically trusted on the other hosts, and do not need a user ID or password to get into the system. If an attacker can successfully spoof a connection from a trusted machine, he may be able to access the target machine without an authentication. Identifying the trusted machine is frequently done using network sniffing.
- ***Man-in-the-middle attacks*:** A man-in-the-middle attack is a type of attack where the attacker breaks into the communication between the endpoints of a network connection. Once the attacker has broken into the communication stream, they can intercept data being transferred, or even inject false information into the data stream. These types of attacks are frequently used to intercept both HTTP and HTTPS connections. Systems connected to a wireless network are very susceptible to this form of attack.
- ***Backdoor attack*:** Backdoor attacks are attacks against an opening left in a functional piece of software that allows access into a system or software application without the owner's knowledge. Many times, these backdoors were left by the application developers, but current code testing has dramatically reduced the number of these found in commercial software. A more common version of this attack occurs when system administrators create system accounts that they can use in the event they are asked to leave the company. As an information security professional, one of your goals should be to validate system accounts belonging to employees at least once a year.
- ***DNS poisoning*:** A DNS poisoning attack is an attack against the cached information on a DNS server. When a DNS request is made, the result of the request is cached on the DNS server so subsequent DNS requests made for the same server can be returned more quickly, without requiring a lookup by an external DNS server. Unfortunately, these cache files are not particularly secure, and attackers target these files to insert a bogus IP address for a specific server entry into a cache. When this occurs, any host making a request for that site from the poisoned DNS server will be directed to the wrong site. The bogus entry in the cache will remain until the cache expires and is refreshed.
- ***Replay attack*:** A replay attack occurs when an attacker is able to capture an intact data stream from the network using a network sniffer, modify certain components of the data stream, and then replay the traffic back to the network to complete their attack. For example, an attacker could capture a session where a purchase is being made, modify the delivery address, and replay the traffic to place an order that would be delivered to their address.
- **Weak encryption keys:** An attack against weak encryption keys successfully occurs when the keys have a value that permits the breaking of the encryption. Once the encryption is broken, the attacker is able to access the data that is supposed to be encrypted. Probably the highest profile example of this attack was the weakness exploited in the Wired Equivalent Privacy (WEP) security standard used in conjunction with wireless networks. Intended to be used to secure the wireless network, instead WEP keys were found to be weak, and could be broken if 5 to 10 MB of wireless traffic could be captured. This traffic could then be run through one of the many tools published by the hacker community, and the result would be the WEP key, which permits the attacker to read the information protected with WEP. This is another example of an attack that relies on a network sniffer to successfully carry out the attack.

- ***Social engineering*:** Social engineering attacks occur when an attacker will contact an employee of the company and try to extract useful information from them. This information may later be used to help pull off a different attack. Social engineering attacks typically have an attacker trying to appear as harmless or respectful as possible. Generally, the attacker will ask a number of questions in an attempt to identify possible avenues to exploit during an attack. If they do not receive sufficient information from one employee, they may reach out to several others until they have sufficient information for the next phase of an attack.
- **Password cracking:** Password cracking is an attack that attempts to decrypt stored passwords. A successful password cracking attack requires access to the encrypted password database, and a tool designed to decrypt the database.
- **Dictionary attack:** A dictionary attack is similar to a password cracking attack, except instead of using a tool to try to decrypt the password, a dictionary attack uses a dictionary of common passwords and repeated logon attempts with those passwords to try to find a logon and password combination that work. A variation on this attack is the password guessing attack, where an attacker will gather information about the victim in an attempt to guess their password. This is why password policies typically prohibit using the names of relatives, pets, and so on for a password.
- **Brute force attack:** Brute force attacks are very similar to dictionary attacks, except instead of using a dictionary of common passwords, a brute force attack tries every single key combination known in order to break the password. The longer and more complex the password, the tougher it is for these password type attacks to be successful.
- **Software vulnerability attack:** An attack against a software vulnerability exploits a known or unknown vulnerability in an operating system or application to perform malicious activities. This is probably one of the most common avenues for attack, and is used frequently by viruses and worms. A solid patch management practice is the best defense against this type of attack, especially if coupled with a vulnerability management program.
- ***Buffer overflow attack*:** A buffer overflow attack exploits poorly written code by injecting data into variable fields and leveraging the response to access information in the application. This attack is made possible when the application developer doesn't limit or check the size of the data being entered in an application field. When data that is too long for the field is entered, it creates an error that can be exploited by the attacker to perform malicious actions against the application.
- ***Remote code execution attack*:** Remote code execution attacks are commonly run against web applications. When an application is improperly coded, an attacker is able to run arbitrary, system level code through the application and use the results to access data or perform other unintended actions against the application or application server.
- ***SQL injection attack*:** SQL injection attacks are one of the oldest attacks against web applications using the SQL Server database application. In this attack, control characters are entered into the web application and depending on the configuration of the database server, the attack can range from retrieval of information from the web server's database to allowing the execution of code or even full access to the server. This attack relies on database weaknesses as well as coding weaknesses.
- ***Cross-site scripting (XSS) attack*:** Cross-site scripting attacks are by far the most common and potentially the most dangerous current attack against web users. These attacks allow attackers to bypass the security mechanisms provided by the web browser. By injecting malicious scripts into web pages, and getting users to execute them, an attacker can gain elevated access privileges to sensitive page content, session cookies, and a variety of other information maintained by the browser.

Understanding Denial-of-Service (DoS) Attacks

THE BOTTOM LINE

A ***denial-of-service (DoS) attack*** is an attack whereby the attacker renders a machine or network resource unavailable. It is usually done by flooding the targeted machine or resources with superfluous requests in an attempt to overload the system. However, a DoS attack can also be caused by disconnecting a power or network cable. Today, most DoS attacks are ***distributed denial-of-service (DDoS)*** attacks, whereby multiple computers are used to overwhelm the network resource.

CERTIFICATION READY
Can you describe a DoS attack?
Objective 3.3

When a DoS attack occurs, the usual symptoms include the following:

- Unusually slow network performance, including when opening files or accessing websites
- Unavailability of any or all websites
- Dramatic increase in the number of spam emails received

There are three general types of DDoS attacks:

- **Volume-based attacks:** Saturates the bandwidth of an attack site or system by flooding the site or system with UPP packets, ICMP packets, or other spoofed packets.
- **Protocol attacks:** Consumes resources of server or communication devices, such as firewalls and load balancers. They include SYN floods, fragmented packet attacks, ping of death attacks, and Smurf DDoS.
- **Application-layer attacks:** Uses system or device vulnerabilities to crash the server or communication device. They include low-and-slow attacks and GET/POST floods.

Some popular and dangerous types of DDoS attacks include:

- ***UDP flood*:** Uses ***User Datagram Protocol (UDP)***, which is a connectionless networking protocol, to flood random ports on a remote host with numerous UDP packets. When the server repeatedly checks for the application listening at that port—to the point at which the system utilizes all its resources responding to it—the system becomes inaccessible.
- ***ICMP (ping) flood*:** Uses ICMP packets to flood systems. This type of attack can consume both outgoing and incoming bandwidth because the victim's servers often attempt to respond with ICMP Echo Reply packets.
- ***SYN flood*:** Many TCP protocols use a three-way handshake in which a SYN Request is used to initiate a TCP connection. The host responds with a SYN-ACK response, which is confirmed with an ACK response from the requester. In a SYN flood, the attacker sends multiple SYN requests but does not respond to the SYN-ACK responses; or the attacker sends the SYN request from a spoofed IP address. Too many SYN requests lead to the system not accepting any new connections.
- ***Ping of death*:** An attack that sends multiple malformed or malicious pings to a computer. The IP package, including the header, is 65,535 bytes in length, and many computer systems were never designed to properly handle ping packets larger than this, because it violates the Internet Protocol. By sending IP fragments with oversized Fragment Offsets, attackers can cause the IP packets, which were split into smaller sizes for travel, to form packets larger than 65,535 bytes after reassembly at the receiver, overflowing the memory buffers. Thus, important memory areas are overwritten, causing denial-of-service for legitimate packets.
- ***HTTP flood*:** Uses many HTTP GET or POST requests to attack a web server or application. This attack is most effective when it forces the server or application to allocate the maximum resources possible in response to each single request.

- ***Email denial-of-service attack*:** Sends so many emails to a user or domain, the server becomes overwhelmed.
- ***Zero-day attacks*:** These attacks are based on using unknown or recently announced vulnerabilities.

To protect against a DoS attack, use a combination of attack detection, traffic classification, and response tools that can identify and block illegitimate traffic. This includes intruder prevention systems (IPS) and security options available in firewalls, routers, and switches that reduce the impact of flooding. Also, check the documentation for best practices in hardening the server and network equipment. Make sure the servers and network equipment are equipped with the latest security patches.

The final component of network security to be aware of is wireless security.

■ Securing the Wireless Network

THE BOTTOM LINE

Wireless LANs have become one of the most popular forms of network access, rapidly spreading through homes, to businesses, to public access wireless hotspots like those found in Starbucks or McDonalds. The convenience of wireless networks must be balanced against the security implications of a network that is not contained by the walls of your building. In this section, we will discuss those security implications, and some of the techniques that can be used to secure a wireless network, including encryption keys, SSID, and MAC address filters.

CERTIFICATION READY
Which methods can be used to secure a wireless network?
Objective 1.4

A wireless LAN (WLAN) allows users to connect to a network while allowing them to remain mobile. While this allows users easy access to the network from areas like conference rooms, offices, lunch rooms, and other such areas that a wired connection wouldn't allow, this also allows potential attackers a similarly easy access to the network. Many corporate wireless networks can be accessed by anyone with a laptop and wireless card. When using a wireless connection in a neighborhood, it is common for your computer to see wireless networks other than your own. Businesses have the same issues as your neighbors. They are broadcasting their network to anyone within range. In fact, with specialized antennas, wireless networks can be accessed from surprisingly long distances, and the access can occur without your knowledge.

In the early days of wireless networking, implementing it was easy, but securing it was not. As a result, there were battles between users, who wanted the ease of access and mobility that wireless promised, and security departments, who were acutely aware of the risks wireless introduced to the environment. As a result, most corporations had strict policies prohibiting the use of wireless to access the internal network directly, frequently requiring users to use VPNs to connect from the production wireless network to the internal network. As a result, users would install wireless access points under their desks, and hope that no one from security would notice. Attackers would drive around office parks looking for these unsecured access points, so they could breach the perimeters of corporate networks and attack the unprotected internal networks. Corporate security organizations would also perform similar exercises, in the hopes of finding rogue wireless connections before the attackers found them. With some of the new security capabilities available with wireless networks, it is now possible to offer reasonably secure wireless access to internal networks, reducing both the frequency of rogue access points and the resources being used to try to find and shut down those access points.

Another capability sometimes discussed when looking at deploying wireless networks is ensuring that wireless access point radio strength is tuned appropriately. While there is the possibility of tuning the wireless signal to reduce the risk of unauthorized users, it is not a good idea to rely on that as a first line of defense when trying to keep a wireless network secure. Frequently, this capability reduces usability far more than it improves security.

Understanding Service Set IDentifier (SSID)

The most basic component of the wireless network is the SSID (Service Set IDentifier). The SSID is defined in the IEEE 802.11 standard as a name for the WLAN. It does not provide any inherent security capabilities, although specifying the SSID name of the WLAN you want to connect to will ensure that you connect to the correct WLAN.

While there aren't any specific security capabilities associated with the SSID, there are some security considerations that should be taken into account:

- **Choose your own SSID:** The first thing to do when setting up a WLAN is to set a unique SSID. Each WLAN access point will come with a default SSID set. If you use the default, there is a risk that one of your neighbors will also use the default, causing confusion and conflicts. So be sure to select a unique, yet easy to remember name for your own SSID.
- **Follow naming conventions:** After choosing an SSID name, there are some measures to take to make it a little more challenging for an attacker to identify the owner of a WLAN. It is generally not a good idea for corporations to broadcast the fact that they are the owner of a specific wireless network. Selecting SSIDs based on company name, company product lines, or anything else that might allow an attacker to confirm who owns the WLAN should be avoided. Select an SSID that the employees can remember, but which doesn't invite attacks. Choosing things like city names, sports, mythological characters, or other generic SSID names are generally safe choices.
- **Turn off your SSID:** The SSID is used to identify your WLAN, and permit computers to connect to it. If this information is broadcasted, then the client systems can search for available wireless networks, and the name of your WLAN will appear in the list. A few clicks and someone can connect to the WLAN. While extremely convenient for an authorized user, broadcasting an SSID makes it equally easy for an attacker to connect in the same way. It is possible to turn off the SSID broadcast for your network, rendering it essentially invisible to casual wireless network browsers. The problem with this idea is two-fold. First, it makes getting authorized users connected to the network more difficult, and second, any attacker trying to get in through your WLAN will most likely have a wireless sniffer, which will show them the SSID of your WLAN whether it's broadcast or not, because that information is in the wireless packets. In this case, it's generally wise to select ease of use over hiding your SSID—in other words, usability over obscurity.

Now let's look at some techniques for securing a WLAN.

Understanding Keys

The best available security mechanism for securing a WLAN is to use authentication and encryption. WLANs provide three key-based security mechanisms to provide that security.

WEP (WIRED EQUIVALENT PRIVACY)

The very first security capability available to WLAN users was ***WEP (Wired Equivalent Privacy)***. WEP was included as part of the original IEEE 802.11 standard and was intended to provide privacy. Widely recommended in the early days of WLAN use, WEP rapidly fell out of favor when a flaw with the encryption mechanism was found.

The flaw in WEP makes it relatively easy for an attacker to crack the encryption and access the wireless network, so it is generally only used if no other solution is available (WEP is better than nothing) or the WLAN is being used with older devices, or devices like PDAs or handheld games that require the use of WEP.

One of the other challenges with WEP was the confusing mix of keys vendors used. Some vendors implemented the keys in HEX, some used ASCII characters, and some just used passphrases. Depending on the version of WEP, the length of the keys could also vary. This was particularly problematic for home users who wanted to use equipment from multiple vendors. Consumers ended up with equipment that wouldn't support WEP in the same way.

WPA (WI-FI PROTECTED ACCESS)/WPA2 (WI-FI PROTECTED ACCESS VERSION 2)

WPA (Wi-Fi Protected Access) was designed as the interim successor to WEP. The WPA protocol implements most of the IEEE 802.11i standard. This was included in the updated WLAN standard. IEEE 802.11i addressed a number of the issues inherent to the original IEEE 802.11 standard. WPA included a new security protocol, Temporal Key Integrity Protocol (TKIP), which, while related to WEP to ensure backwards compatibility, adds new features to help address the issues associated with WEP. Unfortunately, because TKIP uses the same underlying mechanism as WEP, it is also vulnerable to similar attacks. The number of attacks is significantly less than with WEP, but they still exist.

WPA2 (Wi-Fi Protected Access version 2) is the standards-based version of WPA, except WPA2 implements all the IEEE 802.11i standards.

WPA/WPA2 functions in two modes:

- **Shared-key WPA:** In Shared-key WPA, a passphrase is configured that is entered on both the client and the wireless network. This is similar to how WEP works, but the protection of that passphrase is much more secure due to the use of strong encryption with automatic rekeying. This mode is generally meant to be used by home users.
- ***IEEE 802.1x*:** In 802.1x mode, WPA/WPA2 uses an external authentication server coupled with the EAP (Extensible Authentication Protocol) standard to enable strong authentication for connection to the WLAN. The typical authentication process includes:

  1. **Initialization:** On detection of a host, the port on the switch is enabled and set to the "unauthorized" state. Only 802.1x traffic is allowed while the port is in this state.
  2. **Initiation:** The host trying to connect to the WLAN transmits EAP-Request Identity frames to a special Layer 2 address on the local network segment. This is known as the authenticator. The authenticator then forwards the packets to a RADIUS authentication server.
  3. **Negotiation:** The authentication server sends a reply to the authenticator. The authenticator then transmits the packets to the connecting host. These packets are used to negotiate the EAP authentication method.
  4. **Authentication:** If the authentication server and connecting host agree on the EAP authentication method, the connecting host is then authenticated. If the authentication is successful, the authenticator sets the port to the "authorized" state and normal traffic is allowed. If it is unsuccessful, the port remains in the "unauthorized" state and the host will not be able to connect.

The use of 802.1x authentication to secure a WLAN is generally reserved for large corporate environments, where there are sufficient resources to support the additional servers and support required by this mode of operation. 802.1x authentication, particularly when used in conjunction with a token-based authentication solution, permits a very secure WLAN implementation.

Understanding MAC Filters

As we discussed earlier in the lesson, a MAC address is the unique hardware address of a network adapter. This information can be used to control what systems are able to connect to a WLAN through the use of MAC filters. By turning MAC filtering on, network access can be limited to only permitted systems by entering the MAC address information into the MAC filters. The table of permitted MAC addresses is maintained by the wireless access points.

Understanding the Advantages and Disadvantages of Specific Security Types

Now that we have discussed the different security mechanisms available when working with WLANs, we will discuss some of the advantages and disadvantages of each type:

- **WEP:** WEP is a solution that, while better than no security at all, is not particularly secure. The vulnerabilities within the WEP protocols encryption scheme make it very easy to crack. WEP will keep neighbors from connecting to a home WLAN, but will not slow a determined attacker very much.
- **WPA/WPA2:** WPA/WPA2 is the best security method for both home and corporate WLAN security. In pre-shared key mode, WPA/WPA2 can secure the WLAN with a passcode that is shared by the clients and the wireless access points. As long as a secure passcode is selected, this is a very secure solution for small networks. For corporate networks, where additional authentication infrastructure can be purchased, the 802.1x security available within WPA/WPA2 permits a highly secure WLAN implementation. The downside with this approach is that it is more expensive, and significantly more complex than the other solutions. This complexity requires significantly higher support, as user accounts need to be maintained, additional servers will need to be supported, and troubleshooting becomes more challenging. These challenges can be overcome with a well-designed, redundant architecture for the WLAN.
- **MAC address filtering:** MAC address filtering is a good solution for a home or small office environment, but has significant challenges as the number of permitted devices grows. Manually maintaining a table of MAC addresses becomes a significant challenge when there are more than 10–20 devices, especially in dynamic environments where systems are being purchased and decommissioned regularly. Any changes to the list of permitted devices requires someone updating the MAC address filtering table, which is generally a manual process. Another issue with MAC address filtering is that MAC addresses can be "spoofed" by someone with sufficient knowledge, or with the ability to perform an internet search for a tool to change a MAC address. If they are able to get the MAC address of an authorized system, they can reset their MAC address to the authorized address, and thus gain access to the WLAN. MAC address filters are a good solution for small, static environments like a home or a small office. While they will not stop a determined attacker, they are one more impediment to ensure that only a truly motivated attacker will try to bypass them.

The good news when reviewing the available security mechanisms for wireless networks is that there are solutions available for just about any situation. In the early days of wireless, WLANs offered great convenience for users, but no security for protecting the company's network. Deploying wireless access was as easy as buying an inexpensive wireless access point and plugging it into the network. As a result, security departments were forced to dedicate resources to track down rogue wireless access points. Fortunately, there are multiple tools available today that can be used to identify rogue access points. While, the issue certainly still exists, it is not as prevalent as it was in years past.

SKILL SUMMARY

In this lesson, you learned:

- A firewall is a system that is designed to protect a computer or a computer network from network-based attacks. A firewall does this by filtering the data packets traversing the network.
- Firewalls based on packet filtering inspect the data packets as they attempt to traverse the firewall, and based on rudimentary rules, such as permitting all outbound traffic while denying all inbound traffic, or blocking specific protocols from passing through the router, like telnet or ftp.
- Instead of analyzing each individual packet, a circuit-level firewall monitors TCP/IP sessions by monitoring the TCP handshaking between packets to validate the session.
- Application-level firewalls (also known as proxy servers) work by performing a deep inspection of application data as it traverses the firewall. Rules are set based on analyzing client requests and application responses, then enforcing correct application behavior.
- Stateful multi-level firewalls are designed to provide the best features of both packet-filtering and application-level firewalls.
- Virtual LANs (VLANs) were developed as an alternate solution to deploying multiple routers. VLANs are logical network segments used to create separate broadcast domains, but still allow the devices on the VLANs to communicate at Layer 2, without requiring a router.
- Intrusion detection systems (IDS) are designed to detect unauthorized user activities, attacks, and network compromises.
- An intrusion prevention system (IPS) is very similar to an IDS, except that, in addition to detecting and alerting, an IPS can also take action to prevent a breach from occurring.
- Honeypots, honey nets, and padded cells are complementary technologies to IDS/IPS deployments. A honeypot is a trap for hackers.
- A DMZ is a firewall configuration used to secure hosts on a network segment. In most DMZs, the hosts on the DMZ are connected behind a firewall which is also connected to a public network like the internet.
- Network Address Translation (NAT) is a technique used to modify the network address information of a host while traffic is traversing a router or firewall. This technique is used to hide the network information of a private network while allowing traffic to be transferred across a public network like the internet.
- DNS Security Extensions (DNSSEC) adds security provisions to DNS so that computers can verify that they have been directed to proper servers.
- Protocol spoofing is the misuse of a network protocol to perpetrate a hoax on a host or a network device.
- The denial-of-service (DoS) attack floods the network being attacked with overwhelming amounts of traffic, shutting down the network infrastructure like a router or firewall.
- A man-in-the-middle attack is a type of attack where the attacker breaks into the communication between the endpoints of a network connection. Once the attacker has broken into the communication stream, he can intercept data being transferred, or even inject false information into the data stream.
- Backdoor attacks are attacks against an opening left in a functional piece of software that allows access into a system or software application without the owner's knowledge.
- A DNS poisoning attack is an attack against the cached information on a DNS server.

- A replay attack occurs when an attacker is able to capture an intact data stream from the network using a network sniffer, modify certain components of the data stream, and then replay the traffic back to the network to complete their attack.
- A buffer overflow attack exploits poorly written code by injecting data into variable fields and leveraging the response to access information in the application.
- SQL injection attacks are one of the oldest attacks against web applications using the SQL Server database application.
- A wireless LAN (WLAN) allows users to connect to a network while allowing them to remain mobile.
- The SSID (Service Set IDentifier) is the name for the WLAN. A connecting host must know the SSID to connect.
- WEP (Wired Equivalent Privacy) is an older wireless encryption protocol, which rapidly fell out of favor when a flaw with the encryption mechanism was found.
- WPA (Wi-Fi Protected Access) was designed as the interim successor to WEP.
- WPA2 (Wi-Fi Protected Access version 2) is the standards-based version of WPA, except WPA2 implements all the IEEE 802.11i standards.
- A MAC address is the unique hardware address of a network adapter.
- By turning MAC filtering on, network access can be limited to only permitted systems by entering the MAC address information into the MAC filters.

■ Knowledge Assessment

Multiple Choice

Select the correct answer(s) for each of the following questions.

1. Which of the following should be considered when deciding whether to use a software or hardware firewall? (Choose all that apply.)
   - **a.** Host operating system
   - **b.** Application conflicts
   - **c.** Operating system version
   - **d.** Firewall service efficiency
   - **e.** Stability
2. Which of the following are layers of the OSI model? (Choose all that apply.)
   - **a.** Physical
   - **b.** Control
   - **c.** Application
   - **d.** Network
   - **e.** Encryption
3. Routing occurs at which layer of the OSI model?
   - **a.** Physical
   - **b.** Data-link
   - **c.** Transport
   - **d.** Session
   - **e.** Network

4. Which of the following are valid firewall types? (Choose all that apply.)
   a. Virtual
   b. Network
   c. Packet filtering
   d. IPsec
   e. Application

5. Which of the following are typically examined by a stateful inspection firewall? (Choose all that apply.)
   a. IP address of the sending host
   b. IP address of the receiving host
   c. IP address of the router
   d. Data packet type
   e. Data packet size

6. Which of the following is an attack that relies on having a user execute a malicious script embedded in a web page? (Choose the best answer.)
   a. Man-in-the-middle
   b. Brute force
   c. Cross-site scripting
   d. SQL injection

7. A small business owner has purchased a new wireless access point and wants to ensure that only his systems are able to connect to the wireless. He enables MAC address filtering and puts the MAC addresses for all of his computers in the permitted table.

   This filtering occurs at which layer of the OSI model?
   a. Physical layer
   b. Data-link layer
   c. Network layer
   d. Transport layer
   e. Session layer

8. A sales team for a medium-sized manufacturing company has just deployed a new e-commerce application to allow for the direct sale of products to its customers. To secure that solution, an application firewall is deployed. At which layer of the OSI model does the application firewall occur?
   a. Physical layer
   b. Data-link layer
   c. Network layer
   d. Presentation layer
   e. Application layer

9. Which of the following are password-based attacks? (Choose all that apply.)
   a. Replay
   b. Network sniffer
   c. Brute force
   d. Man-in-the-middle
   e. Dictionary

10. Which of the following is an attack that relies on the attacker being able to trick the sending host into thinking his system is the receiving host, and the receiving host into thinking his system is the sending host? (Choose the best answer.)
    a. Replay
    b. Brute force
    c. Man-in-the-middle
    d. Cross-site scripting
    e. SQL Injection

11. Which of the following are common uses for a VPN? (Choose all that apply.)
    a. Remote access
    b. Server isolation
    c. Intrusion detection
    d. Extranet connections
    e. Domain isolation
12. Which of the following are common types of routing protocols? (Choose all that apply.)
    a. Link vector
    b. Dynamic link
    c. Distance link
    d. Distance vector
    e. Link state
13. Which type of DoS attack uses large ICMP packets to cause an overflow of the memory buffers allocated for packets?
    a. SYN flood
    b. ICMP flood
    c. Ping of death
    d. HTTP flood

Fill in the Blank

Complete the following sentences by writing the correct word or words in the blanks provided.

1. A network administrator that has been put in charge of registering a company's domain name and setting up the DNS so that people on the internet can get to the website should use _____ to ensure that DNS entries are not poisoned by an attacker.
2. The two most common protocols that can be used to create the VPN are _____ and _____.
3. The three common types of protocol spoofing are _____, _____, and _____.
4. A type of attack that uses a weakness in an operating system or an application is known as a(n) _____.
5. An attack that relies on access to the physical LAN segment is known as a(n) _____ attack.
6. An attack that records a stream of data, modifies it, and then resends it is known as a(n) _____ attack.
7. The two common types of Network Address Translation are _____ and _____.
8. When setting up a WLAN in a corporate environment, and using 802.1x and a RADIUS server to secure the connections, it is necessary to use _____ or _____ keys.
9. A(n) _____ can be deployed to distract an attacker from the critical systems on a network.

■ Business Case Scenarios

Scenario 4-1: Using Windows Firewall

You are an administrator for the Contoso Corporation and you need to open the Windows Firewall console on a computer running Windows 11 and create a Windows Firewall inbound rule that allows Internet Explorer to communicate over ports 80 and 443. Describe the steps necessary to completing these tasks.

Scenario 4-2: Using a Routing Table

You are administering a computer running Windows 10. Which commands should be used to display the current routes? You want to add a route to the 10.24.57.0 network using the 192.168.50.1 gateway. Display the routes to confirm it has been added. Lastly, delete the new route.

Scenario 4-3: Using Ports

One of your organization's programs needs access to a server that is on the DMZ using the following protocols:

- Secure Shell (SSH)
- Network News Transfer Protocol
- Simple Network Management Protocol
- NetBIOS Session Service
- Network Time Protocol

Define a port and describe which ports are involved with these protocols.

Scenario 4-4: Accessing and Configuring Wireless Settings

As an administrator at the Contoso Corporation, you need to access and configure the D-Link DIR-655 emulator. Describe the steps necessary to performing these tasks.

Workplace Ready

Defense in Depth

In Lesson 1, the concept of defense in depth was covered, describing how it provides multiple layers of security to defend your assets. This ensures that if an attacker breaches a layer of your defenses, there are additional layers of defense to keep them out of the critical areas of the environment. To use access control, establish physical security so that no one can gain direct access to the servers without going through the network. Provide firewalls and routers to limit access over the network. Then, use host firewalls, User Account Control, and other components to protect the server itself.

Besides looking at access control, keep in mind the issues of authentication, authorization, and accounting. To protect the network resources, establish a system that will allow access based on authentication and authorization. Also, to ensure that a security breach has not occurred, remember to establish accounting that needs to be monitoring and reviewing regularly.