CHAPTER

# 14

# Network Security

***Uncertainty is the only certainty there is, and knowing how to live with insecurity is the only real security.***

**— John Allen Paulos**

We've all heard the sensational stories of large databases being hacked and people's medical records, charge cards, banking account information, and other sensitive data being compromised and released for anyone with an interest in using that information for some unethical purpose. However, there are breaches of network security that happen daily and may not make the news outlets for all to hear. These are those little quiet events that happen to individuals, listed under the heading of ''identity theft.'' With the Internet, large amounts of data can be collected on any person. The insidious thing about it is there is no warning that you are being tracked or spied upon. There is no way for the individual under scrutiny to know that someone has an interest in who they are and any other information about them that can be garnered from searches on the Internet. With such information, an unscrupulous person can set up a parallel identity and begin assuming that unsuspecting person's life. The stories that eventually come to light in a case of identity theft are when the victims of such a crime have had their lives totally ruined.

Network threats are real and constant. It is unfortunate that so many individuals seek to prey on innocent and trusting people, but it falls upon those entrusted with that information to safeguard it as if it were their own personal data. This chapter explores the various aspects of network security and what it entails.

# 14.1 Elements of Network Security

You cannot become complacent that your network is secure. It takes diligence to ensure that the information entrusted to the network, either stored on mass media or as it is being transmitted over the network, is protected from any compromise. Network security has become a field unto itself, with professional organizations dedicated to maintaining data integrity and security within the computer network environment. These organizations help develop and set standards for the main elements of network security — policies, access control, data integrity, monitoring, and assurance.

**RANDOM BONUS DEFINITION**

SANS (SysAdmin, Audit, Network, Security) — SANS Institute is a research and educational organization dedicated to training and sharing information with security professionals around the globe.

## 14.1.1 Network Security Policies

''Policy'' is a pretty broad term in the networking environment. For larger organizations, it is imperative that these policies be written out and maintained in a network usage and security handbook, which all those requiring network access to the organization's network should be aware of. The organization's size, structure, and type of information handled will determine how rigorous and how vigorously applied these network policies are. Policies are usually tailored to fit the needs and functions of the organization for network resource usage.

Consider a simple example of possible network policies using a small family model of Dad, Mom, #1 Son, and #2 Son. #1 Son is a college student who is home only for weekends. His computer is a laptop that was purchased when he entered college. #2 Son is a high school student and needs to be monitored for computer usage, since he tends to let his homework slide. Mom and Dad also want to make sure that their sons are not visiting questionable Internet sites. Mom has a laptop to maintain her social organization's information since being elected its president. Dad is a businessman who has both a home office computer and a laptop. The home network consists of both wireless and wired Ethernet connections to a router connected to the Internet. Figure 14-1 shows a topological diagram of the family's network.

At a family meeting, policies were developed to control network access, monitor content, and ensure data integrity. It was unanimously agreed that each computer will have virus protection software loaded, virus definition files will be maintained on a timely basis, and each computer will be scanned for viruses on a regular basis.

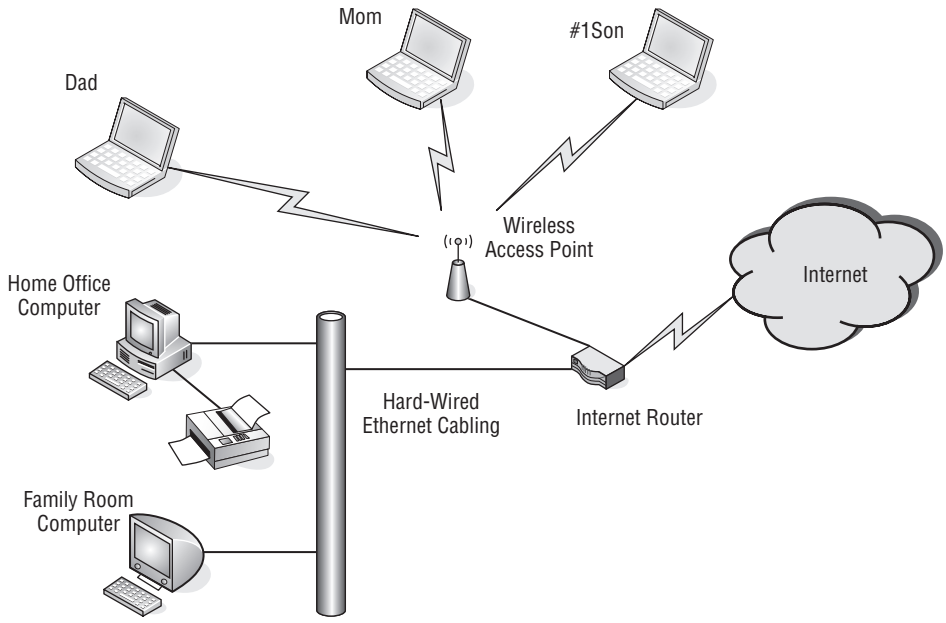


**Figure 14-1** A family's network

Computer access for Dad, Mom, and #1 Son is totally unrestricted. #2 Son has been relegated to using the family room computer, which only has access granted to connect to the Internet from 8 AM to 9 PM. At all other times, it can connect to the local network and use the local resources of the network, such as the shared printer connected to the desktop computer located in Dad's home office. In addition to an Administrator profile, each laptop has a user profile that allows the owner to log on and use the computer either locally or on the home network. The desktop computer in Dad's home office has user profiles for Administrator, Dad, and Mom, with no guest account. The desktop computer in the family room has a user profile for Administrator, Dad, Mom, #1 Son, #2 Son, Local User, and Guest. Dad delegated himself as the network Administrator for all of the household computers.

**ACRONYM ALERT**

SRF — Specifically routed frame

Basically, overall network control and security was set during the family meeting. Dad had already taken care of some of the more obvious sites he preferred that the boys not visit, but also notified them he would be monitoring their access. #2 Son's Internet connectivity would be controlled and monitored by Mom. She will keep an eye on #2 Son to make sure he does his homework prior to any recreational Internet usage. The family room desktop's Internet access is accomplished by allowing access for that computer's IP address

within the time between 8 AM and 9 PM. #2 Son's computer access is controlled by changing his password every day. He is granted a new password on completion of his homework. The Guest user profile password is kept secret by Mom and Dad and will be given to a guest only as necessary. The Local User password is known by all the members of the family, but it only allows use of the local network for access to shared resources on the network, such as the printer. The Local User profile has no Internet access rights.


The reason for the Administrator user profile on each computer is to permit a centralized entity to control every computer's configuration and access privileges to the network. The centralized network entity can enforce network policies on each computer and user. In this example, Dad is the single entity. In a larger organization, however, there may be a single department given that responsibility, with more than one staff member tasked to perform the enforcement of the organization's network usage and access policies. The larger the organization, the greater the need to formalize and document every policy to avoid confusion and to have uniformity across the organization's network infrastructure.

In larger corporate networks, network policies can be enforced by ''pushing''[1] policies down to the computers as they log on to the network. Computers in a corporate environment are usually standardized with as few variations as possible to aid in the supportability of the company's user computer base. It is easier to cookie cutter[2] computer systems than to tailor each one individually. Usually there is an initial software suite of corporate applications that is installed by the IT staff. Depending on the organization and how they enforce their computer and network usage policies, the users of these computers may not be allowed to load additional software without prior authorization. Organizations frown on the loading of rogue[3] software, so if employees require additional software applications beyond the organization's standard software suite, they must seek prior approval.

In our family example, network access is controlled by the computer's user password. All members of the family except #2 Son have control over their own passwords. The homework is #2 Son's token to gain network access for Internet usage. If he needs local access to write reports or print homework,

[1]There is no real pushing involved. Network administrators like to use terms that give a sense of direction and control. The idea of pushing is having an application run on a computer that is automatically spawned each time a user logs on to a network. The application connects with a policy server, which sends the updated policies to the computer. In reality, this is two-way communication, so ''pushing'' is merely a euphuism for the ability to enforce policies remotely on a computer.

[2]*Cookie cutter* is a term used by computer and network administrators to illustrate that it is much easier to replicate the same thing over and over again.

[3]The term *rogue software* refers to unauthorized software.

he can do so by logging on to the shared computer in the family room under the Local User ID and its password. If #2 Son attempts unauthorized Internet access, he will not be granted Internet access for a week.

Dad is both the network and computer administrator for the family and has the ability to program the Internet router to block certain objectionable sites. He had already added objectionable sites to the block access list within the router and has the ability to monitor the sites that are being visited. Periodically, Dad logs on to the router and peruses the accessed site log, looking for URL names that may be for sites with objectionable content. This logging provides Dad a means of monitoring access to the network as well as the address of the computer that is making the connection. From this information Dad can determine if any of the household members are violating the computer and network usage policies.

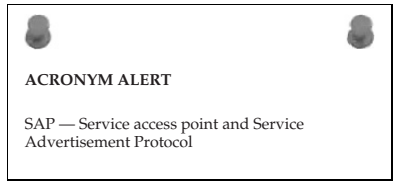


Because all the laptops use wireless NIC cards for network connectivity, Dad has added a WEP[4] key to the wireless router and on the wireless NIC card configuration for each laptop computer. This is basically a software configuration that identifies a user's laptop to the wireless component of the Internet router by the use of a shared key. The WEP key, Internet router password, and the Administrator passwords are only known to Dad. The only time he would relinquish a password to Mom is when there is a network issue and he is not available to perform the task requiring the password. On Dad's return, he will change the password in order to maintain a restriction to the network's resources by using a password revision control. If there is a need to document passwords, this must be safeguarded under lock and key to prevent inadvertent compromise of the network resources protected by these passwords.

**POP QUIZ**

Think about your own personal computer network. Do you have any policies in place? Evaluate your current network and find where you think you may need to set up some sort of policy. Think about your place of employment. Do they have a computer network, and if so, do they have policies for its usage? If they have policies, are they adequate in your estimation? Do you see a need for additional policies? If so, list them and why you think these policies are needed.

[4]WEP is an acronym for Wired Equivalent Privacy. This is used by the wireless components within the network for authentication using a shared key. It is used to prevent unintended use of the network.

**RANDOM BONUS DEFINITION**

antivirus (AV) software — A computer application that scans computer systems to detect and eliminate computer viruses.

This simple example illustrates what network security policies entail. They contain every aspect of a network's configuration and maintenance. They also control user access and usage, as well as the content users are allowed to access. Companies with a large number of users may have a generic user policy that covers the more general terms of computer use and network access. However, there may be further policies that depend on locale, right to know, job function, and other variables designated by the company.

## 14.1.2 Network Access Control

**RANDOM BONUS DEFINITION**

unified threat management (UTM) — A device operating on Layers 2 through 7 that is capable of providing firewall protection as well as filtering content.

Network access control is another wide area of concern. The first thought that most people would come up with when asked ''What is network access control?'' is that it's the use of passwords. But it goes beyond that. When we discuss controlling network access, we are also concerned with the ability of anyone who is unauthorized to have access to any network elements — and this includes physical as well as intellectual information — to alter a network's operation or performance. This means restricting access to where network components may be vulnerable to tampering. It means the security of a facility in its entirety, including the portions of the premises that contain the network components.

### *14.1.2.1 Network Premises Access Security*

There are various ways of securing an organization's premises. They depend on the scope and need of the network to be protected. If a network operation center occupies an entire building, the entire building must be secured to prevent either unintentional or a direct intended act to compromise the network. Since most network breaches are performed by an insider, access to network elements should only be permitted on an as-needed basis. For the most part, restricted access means closed and locked doors, with only those who need access given the right of passage into that network strategic area.

All areas of the network require protection. This is especially true where networking equipment is placed in areas that are mostly unmanned, such

as wiring closets and central network distribution points. It is very easy for someone to sit in a wiring closet with a network analyzer and capture traffic from a particular user and lift passwords and other sensitive information. For this reason alone, these areas should be kept under lock and key.

Using actual locks and keys can be cumbersome, but if that is the only means, then it needs to be accomplished until another alternative method to secure the area is realized. If it is feasible, areas should be secured with a combination badge reader and lock release mechanism. This allows for the greatest flexibility while having a means of logging who has entered a restricted area, as well as the time and date they entered. This information can be used if there is ever an investigation of a network event in that particular area. If your organization is unable to have these types of logging and locking devices and the restricted area is under simple lock and key, there needs to be a person in charge of giving out a key while logging in a ledger who borrowed the key, the date and time periods the key was on loan, along with the reason why access was granted.

Network operations should be in an area with restricted access at all times. The only people allowed in that area unescorted are the staff members tasked with the network operation. A badge key system is ideal for securing the area as well as logging who was present at any particular time of the day if the need for that information is ever required. If the network operations area is not a 24/7 facility, the area needs to be secured and monitored during off-shift hours. Monitoring can usually be accomplished as part of the overall security guard activity that goes on after hours.

Most of the security measures are to keep employees not working in the network operations area from tampering with network resources. The intent can be innocent or malicious, but it does not matter which — damage to network elements can have an impact on the ability of a company to do business. It is better to err on the side of being overly cautious or perhaps a bit paranoid.

Network access premises security is not strictly a function of the IT/network operations staff. It partly falls under the facility's management and maintenance, since they are in charge of the building, and the security department, which monitors company assets. The IT/network operations staff may oversee the overall network security, but safeguarding the network and all its elements requires cooperation and assistance from these other departments.

**POP QUIZ**

While you are at work or at school where there is a network infrastructure, note how their areas are controlled as far as control of access. Do you notice any deficiencies? Do you see areas for improvement? Would you change how the network premises access security is performed?

Although this section is primarily concerned with the physical aspects of the network premises, there needs to be attention on the information maintained within the network operations area. Computer networks can be easily compromised if someone is allowed into sensitive network areas. However, information about the network could point someone with malicious intent toward the areas they are interested in and assist them in targeting those areas. Network diagrams with network addressing should be kept securely and under document control to prevent copies from being taken to plan an attack on the network. Any shared passwords should also be kept under lock and key. A locked file cabinet goes a long way toward preventing easy access to the network.

### 14.1.2.2 Network Access Security and Control

In the earlier example of the family network, network access was primarily controlled by knowing the password to gain access using a user profile. For home and small offices, this type of network authentication may suffice. However, in large networks with a wide range of network services and resources, there is a need to restrict some users to only portions of the network that are required by their function. There is the possibility of multiple authentication services within the same organization. There may be servers within the network that do not rely on network authentication and request a user ID and password from users when they try to gain access to that server.

Figure 14-2 illustrates a small network with an authentication server.

When network users turn on their computers, they are presented with a login dialog window. This login process is twofold. It identifies a user as a particular user with a set user profile for that computer. It also authorizes a user to use the network authentication server for use of the network and the resources connected to it. The services available to the user, after being authenticated that he or she possesses the proper credentials for this network, are print services[5] and access to the Internet.

All network users in the organization have access to e-mail. The e-mail server requires them to log in to their account with a user ID and password. Some users as part of their job function within the organization have access to the application database server, which also requires a user ID/password combination. Keeping separate accounts with separate passwords for each service that requires a login is cumbersome at best. However, there are

[5] *Print service* is a general term to express the use of print servers to output print jobs to a bank of network-enabled printers. A print server is a combination queuing/spooling device that stores (buffers) a print job and directs it to a printer designated by the user. Generally, print jobs are serviced in a FIFO (first in, first out) manner. The print server notifies the user on the progress of the print job, displaying error messages if there are any problems with the selected printer.

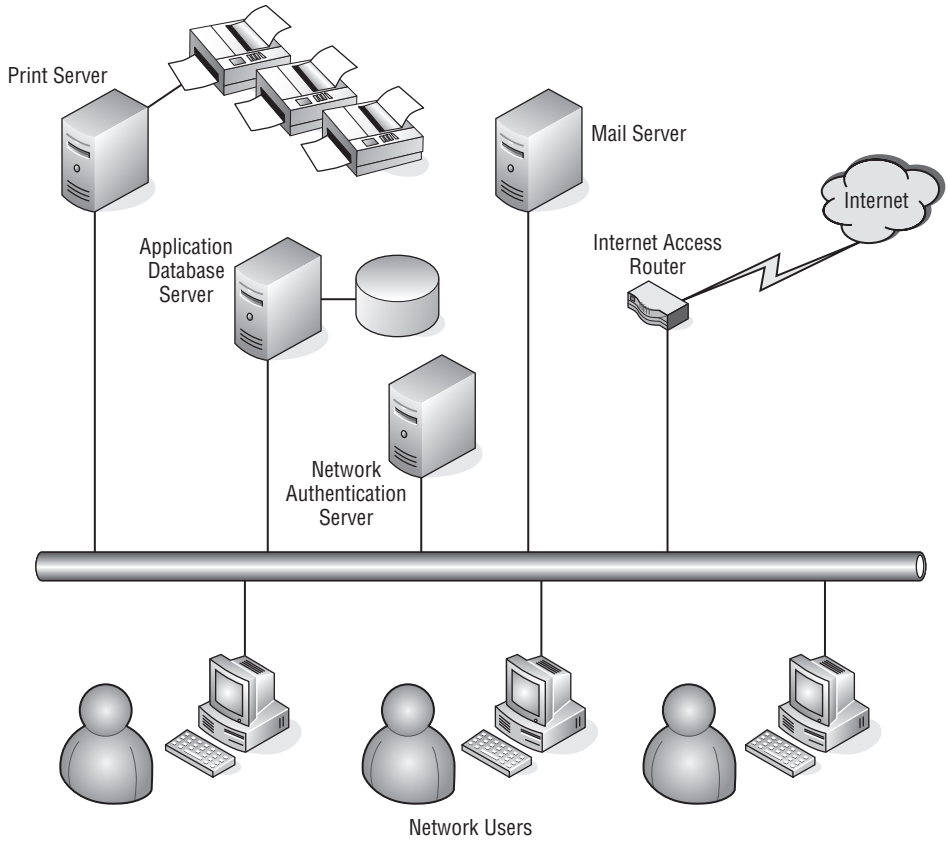


**Figure 14-2** A small network with an authentication server

network authentication services that synchronize user passwords with each service they are permitted to use on the network. If access privileges need to be tailored for each separate user, it would become a logistical nightmare to maintain for a network administrator. So a hierarchical approach is used to determine the privileges a user is allowed. Figure 14-3 illustrates a possible hierarchical map.

**RANDOM BONUS DEFINITION**

network access server (NAS) — A gateway device that protects access to a protected network resource. An additional definition of NAS is *network attached storage*, network-enabled mass storage devices (e.g., hard disks) either singularly or in arrays, that provide increased data reliability through redundancy.

In this figure, users assigned to particular tasks within the organization are placed in groups, with those groups having set privileges for access to certain

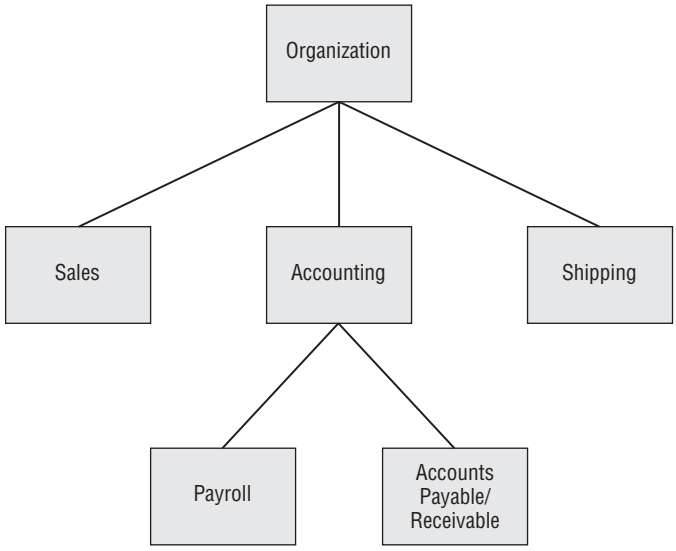


**Figure 14-3** A hierarchical authentication schema

network resources. Under the organization umbrella, all users have access to e-mail and print services over the network. However, differentiation between users begins at the group level, where tasks differ for each user.

In this example, the group level has been divided into Sales, Accounting, and Shipping, each with separate functions and particular informational needs, although they all work in the same organization. The Sales group needs to be able to enter sales and track the progress of orders. The Accounting group is responsible for checking customer credit to permit the continuation of the order process and approve orders to be shipped. The Shipping department processes orders, and when shipped, notes that the orders are completed so that Accounting can process the billing. Although all these groups interact on a particular transaction, they each have separate functions. These functions may be broken out into a particular group of permissions a user has to perform that particular function. The functional separation provides checks and balances throughout the transaction cycle for a particular transaction. It provides accountability for each department and does not allow any one particular user the capability to force a transaction through without assistance from other users in a different department.

**ACRONYM ALERT**

RSVP — Resource reSerVation Protocol

In this example of the group hierarchical schema, there are two subgroups within the Accounting group: Payroll and Accounts Payable/Receivable. The reason for this is that payroll is a very important function within any particular

organization, and the data handled by that group is confidential and requires more security than other functions within the Accounting department.

It is evident from this example that using the hierarchical approach of groups and users not only organizes users by job function, but also aids in network security by permitting only the services required for each user by their group association. If network permissions need to be changed, they can be performed at the appropriate group level to enforce that change on all users who are members of that particular group. Further details of network authentication methods will be discussed later in this chapter.

**RANDOM BONUS DEFINITION**

hacker — A person who has evaded network security with the intention of modifying computer software, hardware configuration, and other security measures to either damage their effective operation or compromise them to the point where data theft can be accomplished without the offenders being detected.

### 14.1.2.3 Restricting Network Access

Restricting network access for unauthorized users can go a long way toward preventing both theft of services and malicious intent. Malicious intent can fall under a number of different categories. The basic elements include hacking into a network with the intent of making it unusable, altering data to one's advantage, or stealing information to get a competitive edge over an organization. The threat of a network attack is not just from outside but often from within. It is not within the scope of this book to go into the psychological implications or fathom the reasons why a member of an organization would use the network to commit crimes against the organization, but it does happen and not all that infrequently.

Recalling the use of a hierarchical schema to give permissions on a network, a network can be segmented to isolate critical areas and prevent access by those who are unauthorized to use those services. This goes beyond authentication and actually will restrict network access based on network address and type of service being requested. An example of this is illustrated in Figure 14-4.

In this example, the organization has an intranet web server to allow all members of the organization to view information about the organization and the products it offers. To prevent unauthorized changes to this important content, it has been decided to isolate the administration of this server on a network segment of its own. Even though a user ID and password are required to log on to the server

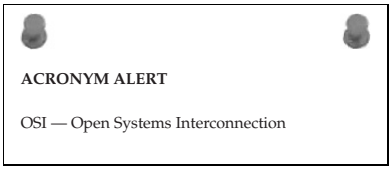

**ACRONYM ALERT**

OSI — Open Systems Interconnection

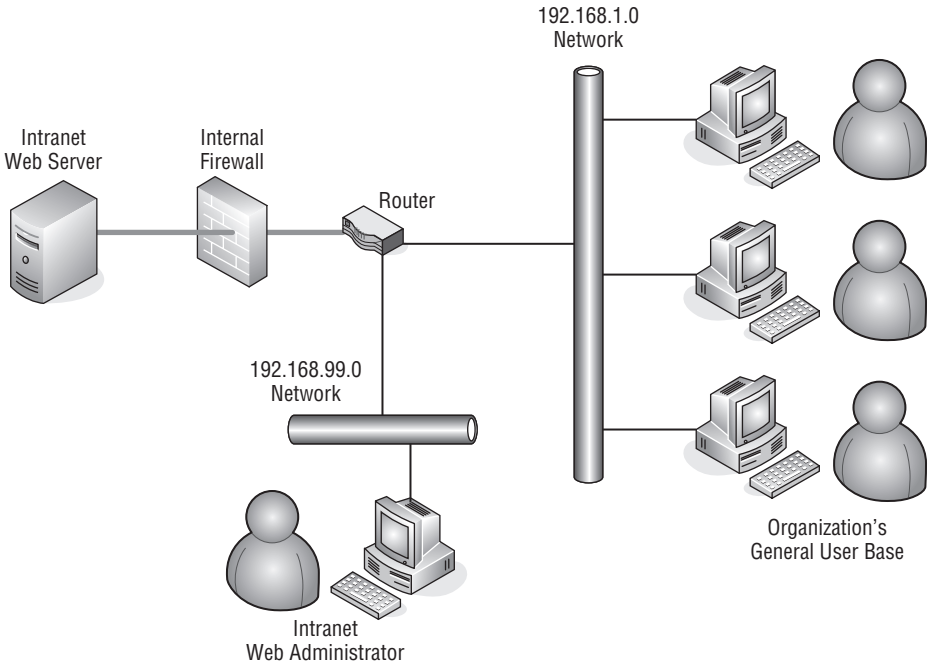


**Figure 14-4** Restricting internal network access

for administration purposes, it was thought that extra measures were needed if ever the user ID and password information got into the hands of an unauthorized user. Although it is not illustrated in this figure, the web server is in a secured location where only authorized personnel are allowed. If an unauthorized user is able to gain access to the server room and knows the user ID and password for the server, he or she would be able to access the administration applications on the web server and alter their contents.

A firewall is placed in the path from the intranet web server to analyze the network traffic that is being directed toward it. All users are permitted HTTP port 80 access through the firewall. As long as the traffic is intended for port 80, it will not be restricted. However, if the traffic is requesting a different service from that which is allowed, then those network packets will be discarded[6] by the firewall and never forwarded on to the intranet web server. Notice that the network segments are different for the web administrator and the rest of the organization's user base. For the web server administrator, a policy has been added in

**ACRONYM ALERT**

VPID — VLAN protocol identifier

[6]''Discarded'' is more appropriate than ''dropped'' when referring to data packets that are not forwarded on by a network-forwarding device. ''Dropped'' implies an action, when none is really taken.

the firewall to permit any traffic from that safeguarded network to be passed through to the web server. The operative word here is ''safeguarded,'' which implies the web administrator is able to secure not only his or her user ID and password for the web server, but also to prevent unauthorized access to his or her network connection by locking the office when he or she leaves for the day. This prevents access even if his or her user ID and password have been compromised. All precautions are required when it comes to restricting unauthorized network use.

### 14.1.3 Network Data Integrity

Throughout this book, emphasis has been placed on the ability to pass data over the network without error. So what is network data integrity? It is the guarantee that data sent to an intended network node arrives without alternation. The other part is that if the data being sent is of a sensitive nature, it needs to be safeguarded and not ''eavesdropped'' upon as it traverses over the network.

One method, as previously mentioned, is to secure the premises and not allow unauthorized personnel into areas where the network could be easily snooped. The other method is to safeguard data with encryption. Encryption encodes the data being transmitted over the network in such a manner that only the two endpoints of the network connection are able to decrypt the data. This was initially performed by sharing a known key[7] between the two endpoints, which was used to encrypt the data that was sent and decrypt the received data to make it readable. Figure 14-5 illustrates an example of sharing a key between two endpoint network nodes.

Rob and Jack want to share sensitive information over the network. The network can include the Internet and any other portions of network, whether public or private. To accomplish this, they decide to encrypt the data before sending it over the network. They settle on an application that allows them to use a shared key between them. This is secure only if they are the only two people who know what the key contains. However, they should not send the key over the same network link. If Rob and Jack are unable to meet privately and the key needs to be carried over a public network, they need to be cautious not to give the people who may be snooping their communications the key along with the encrypted document. It would be better for them to mail the key through the post office than to e-mail it. Or they could simply call each other to pass along the key. The only caveat is that with the convergence of voice and data on the organization's network, if the phone service is one and the same as the data network, then there is no security for Rob calling Jack

[7] A shared key is a string of ASCII characters used to encrypt and decrypt data. The key must be known by the entities on both ends of the network connection.

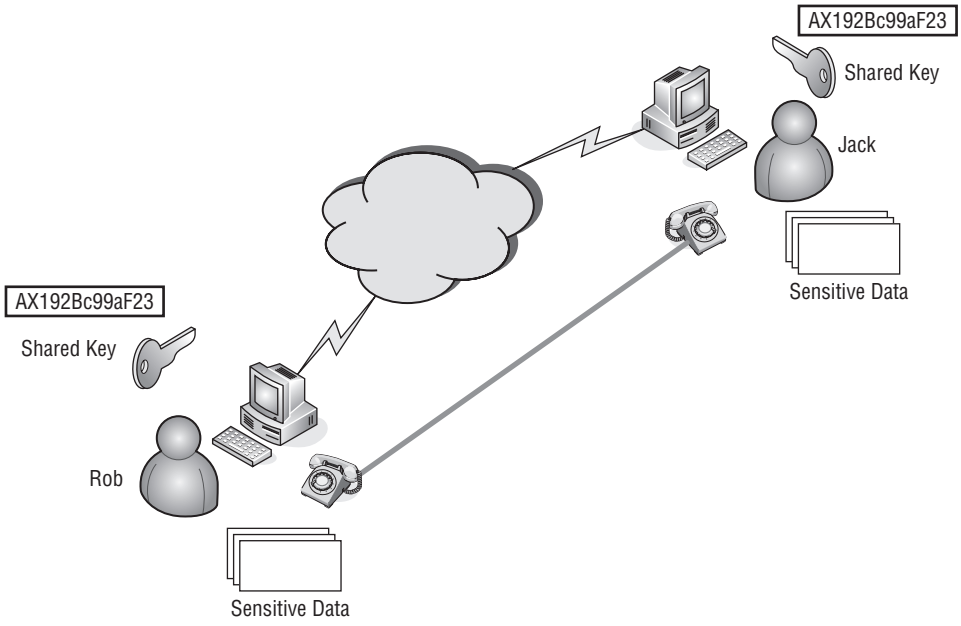


**Figure 14-5** Endpoint-to-endpoint encryption using a shared key

over that link to pass along the key. Or Rob and Jack could meet at a bar and pass the key on the back of a cocktail napkin, like a 1940s spy movie.

The advent of VPN technology has removed some of the melodramatic antics about passing a shared key. A VPN can create a virtual data tunnel over a network between endpoints that are secured with encryption while data is being passed between them. The primary reason VPN technology evolved was to use the Internet as a point-to-point carrier for two remote locations. Figure 14-6 illustrates the connection of two networks over a network cloud.

**ACRONYM ALERT**

MSTP — Multiple Spanning Tree Protocol

In this example, the network cloud can be an internal network, the Internet, or a combination of both, depending on where the VPN routers are placed in the network. The VPN routers establish an encrypted tunnel between them with the use of a pre-shared key and a private key that is known only to the endpoint it belongs to. When the tunnel is established between the two VPN routers, the keys are used to encrypt the data entering the tunnel and decrypt the data as it exits the tunnel. The data is safe between the two VPN routers because of the heavy encryption,

**ACRONYM ALERT**

SNMPv2 — Simple Network Management Protocol version 2

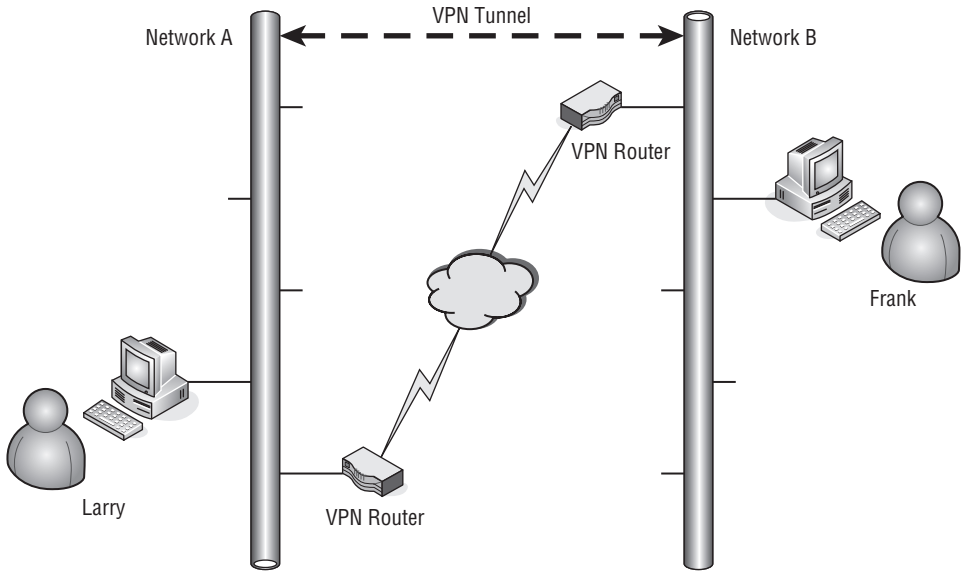


**Figure 14-6** A VPN tunnel connecting two networks

but there is still vulnerability once the data has been decrypted and is traveling over the private network.

## 14.1.4 Network Security Monitoring

Depending on the size of the network, constantly monitoring every node can be overwhelming. However, continuous checking is possible with random checks of certain key areas. The first red flag that should go up is if any of the server logs shows a larger than normal amount of authentication failures. This could be a good clue that someone is attempting to hack into the network. Rest assured that if authorized users forget their password, they will not wait too long before calling the network operations help desk.

Most network administrators generally permit only a single login per user. If a user who has only a single login account tells the help desk that the network authentication server is saying they already have an active session, it is recommended to investigate the reason why before ever increasing a user's account to allow them to have more than one simultaneous active user session. There is a possibility that their user ID and password may have been compromised and another unauthorized user is using their account to gain access to the network.

Administrators should review login data on all servers. They should review the logins to the network authentication server to look for any unusual logins — for instance, if a daytime user is seen to be logging into the network late at night when they have never done so in the past. This can indicate the

possibility that a user ID and password have been compromised and that an unauthorized user is gaining access to the network using that account. Unusual activity even by an authorized user may indicate they are attempting to use the network for covert activity.

Calls to the network operations center can also be an indicator that the network is under attack. Of course, there is a possibility that an actual hardware failure may have caused the issue. However, those are usually total outages of portions of the network. Calls such as slow performance or sluggish response from some application servers need to be investigated not as a possible hardware issue but as possible unauthorized network usage. It could be a deliberate act by an individual, or it could be inadvertent and caused by a user receiving a virus on their computer. In any case, a network analyzer should monitor the network traffic on the affected network segments.

Look for what appears to be heavy network traffic flow and determine if it is from a single source or the whole segment. It is possible that a virus has infected multiple computers on a particular segment. There is a possibility that there is a hardware cause, such as a chattering[8] NIC. You may need to take a divide-and-conquer approach, where you isolate the offending network segment until you can isolate the source of the issue. Sudden changes in performance without any configuration changes on a network could be an indicator of component failure, but usually the issue is not hardware-related. Never ignore users' complaints about network slowness unless you already know the cause. If not, the matter needs to be investigated as quickly as possible to avoid a more catastrophic network event.

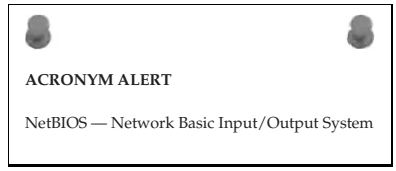


## 14.1.5 Network Security Assurance

A large network needs to be monitored closely on a daily basis. However, assurance is more than investigating issues when they occur. Assurance is a recurring, proactive activity that is accomplished at fixed intervals. The size of the organization's network will determine how often a review of the entire network needs to be done. In the case of multiple installations in various locales, a network security audit can be completed on a rotating basis between the different network areas.

Assuring the network from a security perspective requires full documentation as the network currently exists. The documentation should include

[8]*Chattering* is a term used to express that a network card is broadcasting when it should not be. It is constantly beaconing and causing unnecessary traffic flow on the network segment it is connected to.

network diagrams with network address schemes and the physical locations of the equipment and cabling being used. Any deviation from what has been documented as part of the network needs to be investigated. As noted earlier, many if not most network security breaches are caused by personnel employed by the company.

A network assurance security audit should entail the following:

- A list of equipment located in the network segment under audit.
- A network topology diagram showing the network addressing scheme.
- A password list for network servers in the segment.
  - Ensure that all user IDs are active.
  - Ensure that unauthorized users are not on the list.
  - Verify that staff members who have left the company are not on the list.
- Server access logs.
  - Look for suspicious and repetitive logins.
  - Examine what appear to be frequent or out of the ordinary login times.
  - Look for repeated login denials. Are there user IDs that suggest an unauthorized person is trying various user ID combinations? Are these coming from a particular network address?
- Traffic patterns on the network segment under audit.
  - Look for what appears to be unusual traffic flow. Is this traffic legitimate? If so, it may indicate the need to redesign portions of the network segment.

## 14.2 Network Security Methodologies

There are various methods for ensuring authorized users are permitted access to only those resources they need to perform their function within the organization. This section reviews some of the most widely used authorization methods for network access. The primary authentication methods discussed are Lightweight Directory Access Protocol (LDAP), Remote Authentication Dial-In User Service (RADIUS), and certificates.

**POP QUIZ**

How often should a network administrator think about their network's security? Give a reason for your selection.

**RANDOM BONUS DEFINITION**

managed security service provider (MSSP) — An ISP that provides additional network security management, which may include virus scanning, intrusion detection, and firewall capabilities.

Network protection requires more than just access control; it also includes data integrity. Data traveling to the trusted portions of the network, whether from internal or external sources, can be protected using encryption.[9] The use of data encryption provides the capability for secured tunneling for the creation of virtual private networks (VPNs). The tunnel is created between network nodes that can encode and decode the data that is passed between them.

## 14.2.1 Authentication

The process of network authentication can be as simple as a user ID[10] and password. However, even with that, users can be restricted by their group association to certain locations or resources within the network. It is the responsibility of the client to properly identify itself to the authentication servers on the network.

### *14.2.1.1 Lightweight Directory Access Protocol*

The use of LDAP emerged from the X.500 directory service. It has gained in popularity as a means of authenticating users for a wide range of network services. It is the model being used for directory services on the Internet. X.500 is the International Standards Organization (ISO) and International Telecommunications Union (ITU) standard that defines how global directories are to be structured. LDAP is used by many suppliers of software for their directory services strategy and has widespread acceptability among the network user base. The standard uses a hierarchical directory structure that is parsed on different levels of categorized information. The customary information used for these categories are elements such as country, state, city, and other locale information.

LDAP uses an Internet identity schema that defines common attributes to define the objects contained within it. Many levels of an LDAP can be defined, and authenticating users depends on the granularity required by the network site.

[9]Encryption is using cryptographic means of using ''crypto'' keys to encode the data so it is not easily readable by anyone that does not have possession of the key.
[10]User ID (identification) is synonymous with username or the prompt that is displayed on some systems as ''username.''

The most common used elements in LDAP are:

- Users
- Groups
- Filters
- Services


An LDAP directory service is based on a collection of attributes used to define a distinguished name (DN). The intended use of a DN is to define an entry without any ambiguity so that each entry is unique. The entry is defined by a series of attributes, which are commonly mnemonic strings, such as cn for common name, mail for an e-mail address, etc.

Being a hierarchical, tree-like structure, an LDAP directory's entries are arranged to reflect boundaries that are categorized by geographical, political, or organizational descriptions. An example of this tree-like structure would be starting at the top with the largest entity as country followed by each subgroup (e.g., country/state/organization/department/user).

LDAP uses a special attribute called objectclass to control the required attributes to define an LDAP entry. The objectclass defines the schema rules used for interrogating, maintaining, and updating the LDAP. The primary function of an LDAP server is to respond to service requests for an inquiry by searching the contents of its directory. Since for the most part LDAP data is stored in cleartext,[11] an LDAP server requires the client requesting service to authenticate itself prior to responding to the request for information. Usually the authentication scheme used between an LDAP client and an LDAP server is Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP). PAP is a point-to-point protocol that uses a simple username and password sent over the network in cleartext for the authentication of an LDAP client requesting LDAP services from the LDAP server. To increase security between an LDAP client and an LDAP server, CHAP can be used. CHAP is also a point-to-point protocol that uses a three-way handshake to validate the identity of the remote client. Both client and server use a hashing algorithm using a shared secret to ensure the validity of the connection. CHAP has security advantages that are more desirable than what PAP offers. However, both client and server must be capable of using that protocol.

Figure 14-7 illustrates a model used between an LDAP client and server.

[11] Cleartext is text that is clearly readable and is not encrypted. These files can be usually examined with any text editor.

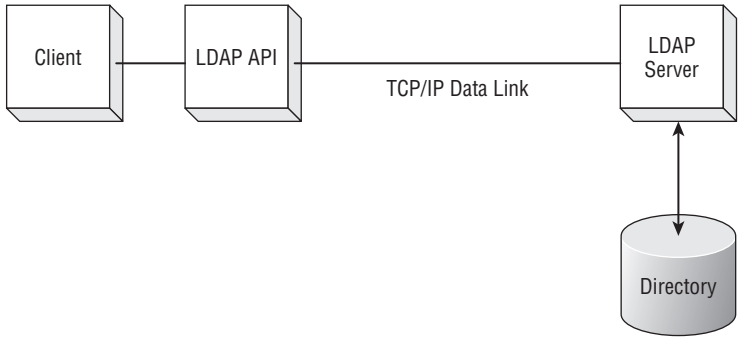


**Figure 14-7** An LDAP model

The client can be any device that needs to authenticate users to permit the use of services over a network. The client loads an LDAP application program interface (API) allowing it to open a TCP/IP socket with the LDAP server. Once the server authenticates the client, it permits it to access its directory. A real-world scenario may be similar to what is illustrated in Figure 14-8.

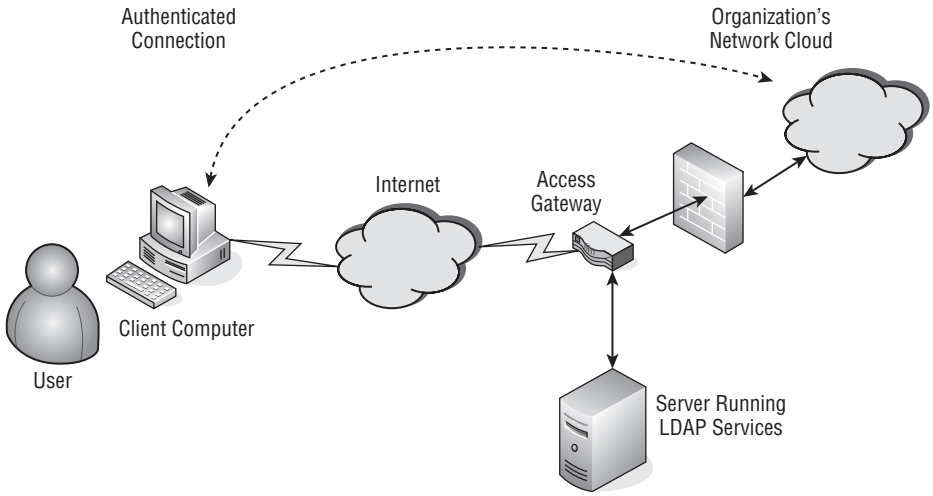


**Figure 14-8** User authentication using LDAP

In this figure, a user is attempting to gain access to the organizational network of which he is a member. When the user became a member of the organization, the IT staff responsible for maintaining the LDAP server entered the information regarding this particular user. The information can contain some or all of the following, depending on the schema that is in use: username, full name, department, e-mail address, and filters used to determine access privileges. The user initiates a connection to the device acting as the access gateway.

Upon receipt of the initial contact, the user is required to enter his or her credentials for access. The access gateway device acting as an LDAP client sends the request to the LDAP server it is connected to over the local network. If the access credentials match the database, a response is returned granting access as well as the level of rights that this user is to have on the network. On establishment of the authenticated connection, the remote user is virtually on the organization's network. This illustration is just showing LDAP being used for remote access, but in reality it also can be used for internal services on the network. Multiple servers can have the capability to use the common LDAP database for authentication to the particular services they offer on the network, as shown in Figure 14-9.

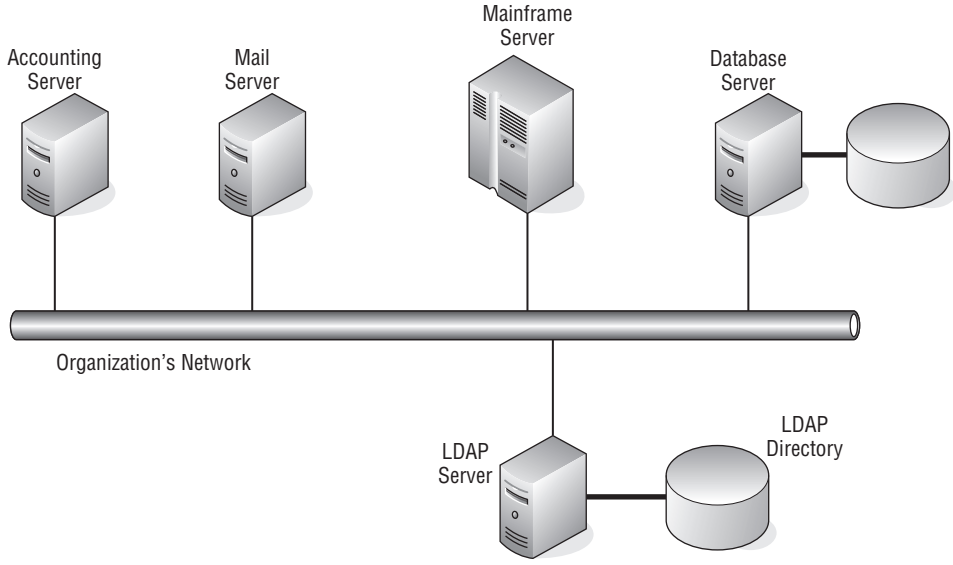


**Figure 14-9** An LDAP server servicing multiple clients

As illustrated, there are a number of servers on the network, each of which requires authentication of a user in order to use the services provided by that server. Many servers offer some sort of authentication scheme, but if each server had its own user database, administrators of those servers would be busy ensuring that a particular user has rights to use that server and the level of permissions they are to have on that server. The commonality of having a single LDAP server performing the authentication function service for the network does alleviate the headache of maintaining so many servers. However, the caveat is that with a single LDAP server, there is a single point of failure if that server should go down. Depending on the number of users a site may have, it may be prudent to have an alternative authentication method. One scheme would be to have redundant LDAP servers that synchronize their

databases to ensure that user credentials on both LDAP servers are current. If the primary LDAP server were to go down or become unavailable for any reason, the devices requiring LDAP services would switch to the backup LDAP server.

On a smaller network or perhaps due to cost restraints, a redundant LDAP server is not possible. A possible workaround can be accomplished using the internal authentication services of the servers themselves to act as a backup to the LDAP server. This does require additional work on the part of the server administrators, but if needed it can be used if the external LDAP server experiences a failure. The configuration that would be used is to program the servers to first search the external LDAP for authentication requests and if no response is received from that LDAP server to then search the internal authentication database. Using this method takes more effort to maintain, but it would allow the network users to still gain authorized access to the network.

**ACRONYM ALERT**

GVRP — GARP VLAN Registration Protocol

Figure 14-10 illustrates the flow between the LDAP request and the information returned.

As illustrated, the LDAP client opens a connection over TCP/IP to the LDAP server. After the client is identified and authorized by the LDAP server, the client is bound to the LDAP server and submits an LDAP query. The server continues to return data to the client until the client's query has been satisfied. Upon completion of the query, the LDAP client unbinds from the LDAP server and the TCP/IP connection is closed, indicating the completion of the LDAP transaction.

The following are the basic responses an LDAP client can receive:

- Authenticated
- Denied
- Timeout

When a user is authenticated, the LDAP server returns information about the user, such as group membership, to the client that is requesting the information. Another valid response is that the user is not in the LDAP directory. In the event that a client receives no response, a timeout condition is reached and the LDAP client has two options available: to seek another means of authentication or to deny the user access. A

**RANDOM BONUS DEFINITION**

firewall — A function using hardware, software, or a combination of the two to detect and prevent unauthorized access to a network.

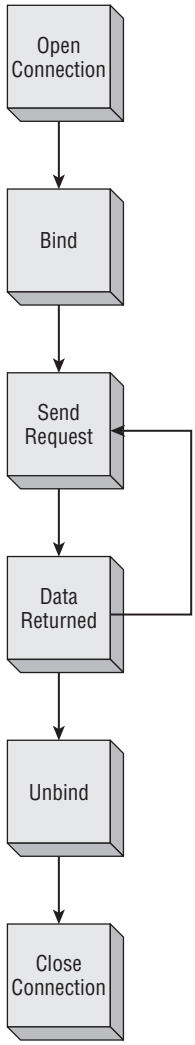


**Figure 14-10** The flow of an LDAP request

timeout condition can be reached for many reasons, such as network congestion, a busy server, or simply that the LDAP service on the server has been shut down. This is the primary reason for having multiple authentication servers available. Some larger organizations go as far as having redundant sites that are interconnected but located geographically distant from each other. Redundancy schemes depend on the size of an organization, but no matter how small the organization, an administrator needs to consider the alternatives in order to have high availability of the network and its resources for the user base.

### 14.2.1.2 RADIUS

Like LDAP, RADIUS performs user authentication but also has an accounting component (a RADIUS accounting server) associated with it. Initially, RADIUS authentication and accounting was used by telecommunications companies to authorize subscribers to the network services being offered and as a means of tracking usage for billing purposes.[12] However, RADIUS is not just used by older ''legacy'' network systems. It has found wide usage in the industry primarily for authentication purposes, although there are organizations that sell information on their online databases based on the amount of time a user remains connected to their service. These organizations use both the authentication and accounting components of a RADIUS server. Interestingly, RADIUS servers or their derivative will remain in use by telecommunications companies for a while to come. They are being used in the cell phone industry to authenticate your calls as well as to keep track of your minutes.

As with an external LDAP server, a RADIUS server provides a centralized user administration service. A RADIUS server can provide authentication services for multiple network devices. These network devices have a RADIUS client embedded within them that can be configured to establish a secure connection with a RADIUS server. The client and server use a shared secret to hash the user password that is being passed from the client to the server. This protects the identity of the user; however, the connection between the client and server is not as secure. The initial connection between the RADIUS client and server is established using either PAP or CHAP, which is less secure since the password being passed between client and server is in cleartext.

When a RADIUS client passes a user's authentication credentials to the RADIUS server, the server will respond with one of the following responses:

- **Access Reject** — The user is denied access to all network resources. Reasons may be invalid credentials, no account on the server, or an account that has been deactivated.
- **Access Challenge** — The user needs to provide additional information. Information requested may be a secondary password, PIN, or token.
- **Access Accept** — The user is granted access.

Once a user is granted access, the RADIUS server returns attributes that have been determined by the information in the user's record contained in the RADIUS server's database.

[12] In today's high-bandwidth, constantly connected network world, billing is primarily a monthly flat fee charge dependent upon the type of service that is being subscribed to. But some of us long in the tooth and gray haired guys remember the days of dialup services on 300 baud modems, which were painfully slow and billed by usage. The service type was billed in an increment of hours, such as 25 hours per month, with additional charges for each minute over, which could get pretty expensive. So, yes, the ability to tie accounting to network access was extremely important for those whose revenue was dependent upon it.

The user attributes that can be returned are as follows:

- **Assigned user IP address** — A statically assigned address may be given to a particular user.[13]
- **Assigned IP address pool for the user** — A dynamically assigned address from the address pool this user is assigned to.
- **Maximum connection time** — For connections that are limited by the amount of time they can remain connected.[14]
- **Service level** — These may be in the form of permissions or restrictions on the user's ability to use particular resources on the network.

RADIUS authentication and accounting not only has widespread use among legacy systems, but it can also be found in many installations requiring centralized administration for authentication and accounting purposes. This makes RADIUS the current de facto standard for authentication systems.

**RANDOM BONUS DEFINITION**

disaster recovery (DR) — A plan to maintain or restore network services after a catastrophic event.

### 14.2.1.3 Certificates

Digitally signed certificates came into use as a security method of ensuring that the two parties on either end of a connection are who they claim to be. This is to prevent unauthorized (spoofing) users from gaining access to a server pretending to be someone else. This is especially important with the establishment of e-commerce over the Internet. The process binds a user's identity to a publicly encrypted signed key that has been verified and validated by a trusted third-party called a certification authority (CA). The CA registers the certificate, ensuring that the certificate and the relationship between it and the individual user are accurate. Figure 14-11 illustrates the relationship between a user, server, and CA.

[13]Particular care is required to avoid duplication when statically assigning user IP addresses. Duplicate addresses on a network can wreak havoc and can be difficult to diagnose unless you get lucky and guess that there may be two different devices responding to requests on a particular IP address. Document well when using statically assigned IP addresses.

[14]Many broadband connections don't use a maximum connection time attribute. However, for services offered by a server, there may be a limited amount of resources, and the number of connections being serviced is a determining factor of the quality of service the server is able to provide. In these instances, this attribute can be used to force log off users that are hanging on the services for long periods of time. The length of time is purely dependent upon the installation site and its operating mode. Many installations allow a maximum of a 24-hour period before a user is given a forced log off.

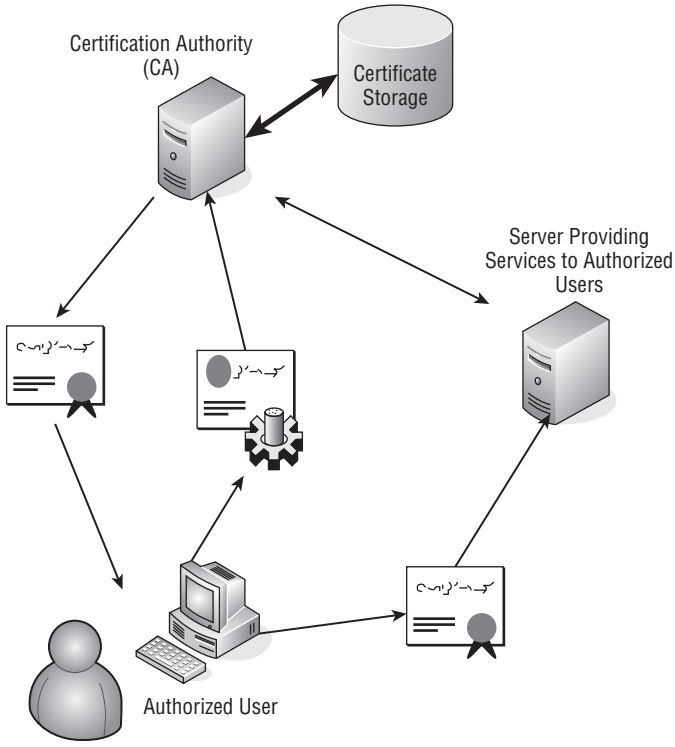


**Figure 14-11** The certificate relationship

A user is required to have a signed certificate to gain access to a server over the Internet. The user provides credentials to a CA, which confirms the user's identity and provides the user with a signed certificate. The user then presents the certificate to the server they want to gain access to. Because the certificate is from a trusted CA, the server permits access to the user. The server providing the service can confirm the validity of the certificate by communicating directly with the CA.

The CA is responsible not just for signing certificates to validate a user's identity, but also for administering the certificates it issues. The CA needs to provide storage for issued certificates to provide maintenance of the certificate's validity. If for any reason a certificate has been invalidated, the server must also maintain a certificate revocation list (CRL). The CRL is used to prevent users with invalid credentials from gaining access to a server. Certificates are usually issued with an expiration date, and this too can be a reason for a certificate to be placed on the CRL. Certificate time periods are usually for a number of years, which is set by the certificate issuer.

Certificates are heavily used in web-based applications. Most web browsers can use HTTPS (Hypertext Transfer Protocol Secure). This uses the Secure Socket Layer (SSL), which resides between the HTTP layer and the TCP

layer. The user/client creates a secure socket connection to a website/server, taking advantage of public and private key encryption with the use of a digital certificate. SSL has been recently superseded by TLS (Transport Layer Security). TLS 1.2 is defined by RFC 5246, ''The Transport Layer Security (TLS) Protocol Version 1.2.''

TLS is a client/server-based protocol that establishes a stateful connection using a handshake procedure. The handshake is used to negotiate and establish the network security that is used between the client and the server. This handshake interaction is initiated when a client requests a connection to a TLS-enabled server. The client informs the server which ciphers and hash functions it supports. The server then selects the strongest cipher and hash function it supports, and sends a response to the client. In its response, the server returns a digital certificate that verifies its identity, which usually contains its server name, the name of the trusted CA, and its public encryption key.

Upon receipt of the server's digital certificate, the client contacts the trusted CA to verify the certificate's validity before proceeding any further. If the client is satisfied with the server's certificate's authenticity, it generates a random number, which it encrypts with the server's public key, and forwards the encrypted value to the server. The server is the only entity that can decrypt this encrypted number from the client using its private key. The random number is used by both client and server to develop keys to encrypt and decrypt the data that is passed between them. With the completion of the handshake, the secure connection is established and the generated keys are used to encrypt and decrypt the data being passed, until the connection is terminated.

**POP QUIZ**

Open your browser and examine the certificates in its certificate store. Open a certificate and examine its contents. Note in particular creation and expiration dates. Note the intended uses for this certificate. Does the certificate store contain all the certificates you expected or many more? Can you think of a reason why that is? Note that each browser program may have different ways of displaying the certificate store. An example of this is Apple's Safari browser, where you select Edit ➢ Preferences ➢ Advanced, and then click the Change Settings button for Proxies. Click the Content tab and then in the Certificates section, click the Certificates button to view the certificate store. There are a number of tabs to select the various types of certificates.

If for any reason any part of the handshake process fails, the connection will not be created. The client must attempt to initiate a new handshake sequence when making further attempts to create a secure connection to that particular server.

All browsers have a certificate store associated with them. Usually this can be found under the options for a browser application. An example screen is illustrated in Figure 14-12.

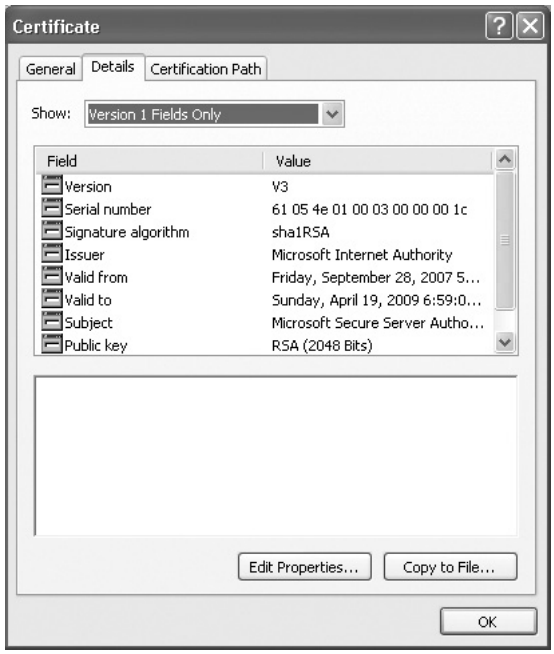


**Figure 14-12** A browser's certificate store

## 14.2.2 Data Integrity

The Internet has provided many positive things. It has spawned whole new businesses with what is called e-commerce. It has allowed people to take communications to new levels. But the flip side is that the hyper-connected world, where information about anything and anyone can be found, has brought out those who have found illicit uses for that information to prey upon unsuspecting users.

The previous section discussed SSL and TLS using certificates to provide proof of identification for authentication. They can also provide an encrypted connection to pass data safely between a user (client) and a server that are directly connected to each other. But what about a remote user who requires services located on the company's intranet? SSL/TLS are client/server-based, which means a connection is established between a user and a particular server. If a user requires many different services on a network, they need to establish a secured connection with each server providing that service.

Tunneling protocols have been developed to allow remote users to work as though they were directly connected to a local network. A conceptual illustration is shown in Figure 14-13.

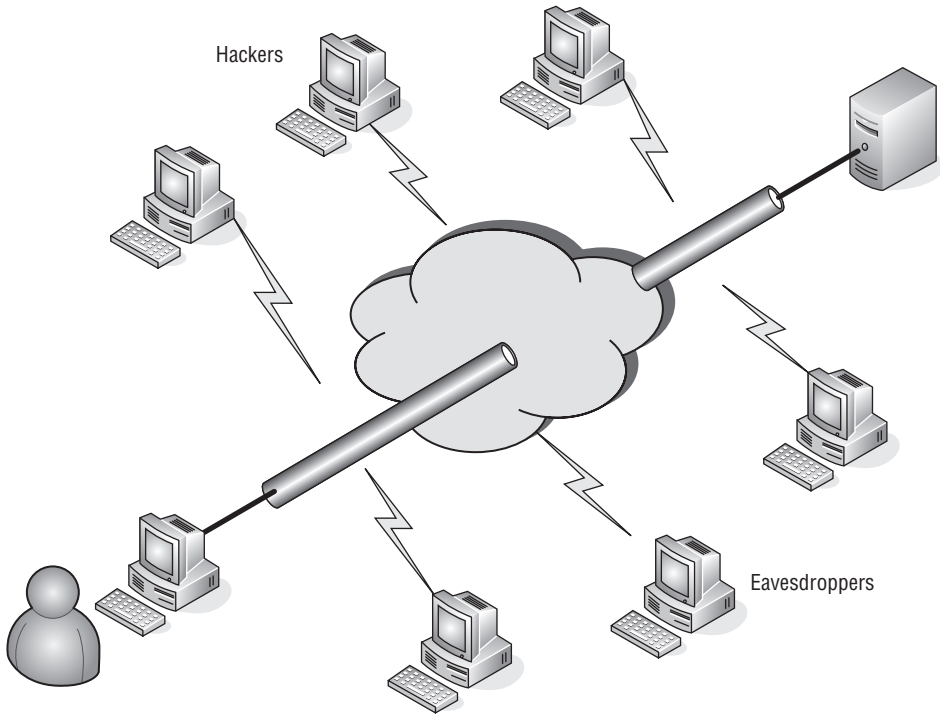


**Figure 14-13** The tunneling concept

As shown in this figure, a connection between a user and a particular network or server can be under attack not only over the Internet but locally as well. However, if a connection can be tunneled through that hostile environment, the conversation is protected from eavesdroppers or hackers trying to steal a user's identity for later use in an attack on that network or server. With the development of tunneling protocols such as PPTP, L2TP, and IPSec, the concept of the VPN was spawned.

The basic concept of a VPN is that the local network is secure. Earlier in this chapter, we discussed local network security and why it is needed. If the premise is that the endpoint networks are secure, the only way to ensure total network security is to secure the link between the locally secured networks. Tunneling protocols allow organizations with separate, geographically distant networks to connect these local private networks using the Internet as the conduit. Figure 14-14 illustrates this concept.

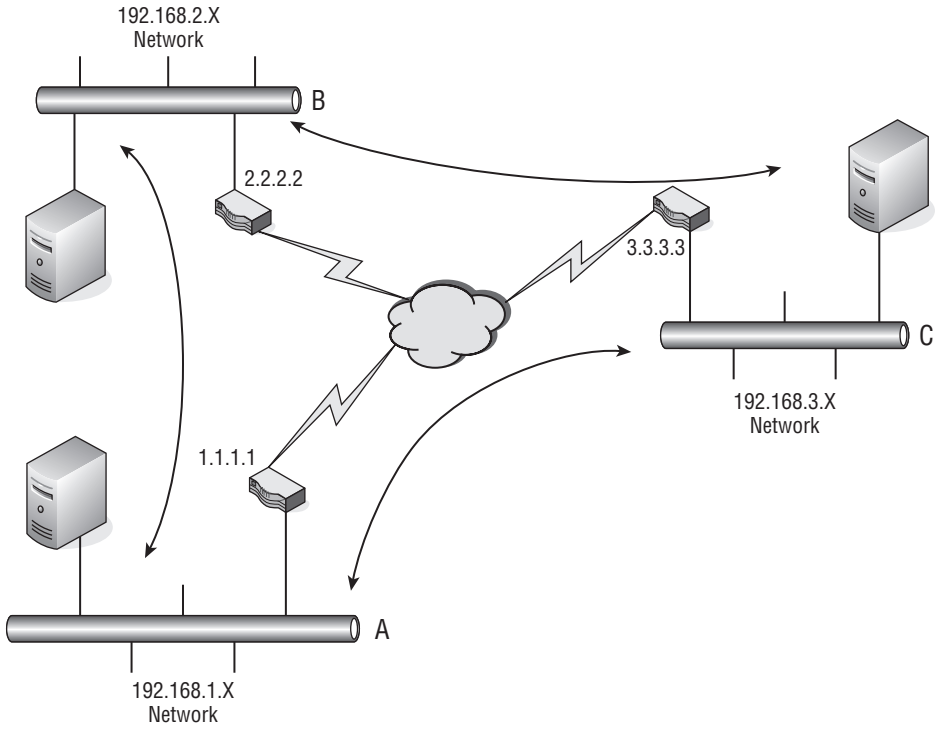


**Figure 14-14** The use of the Internet for VPN

This figure shows three private networks: 192.168.1.X, 192.168.2.X, and 192.168.3.X. These are 24-bit mask networks.[15] Each network is connected to the Internet with a VPN-enabled router. The routers require two separate tunnels for each of the two locations they are to be connected to. Although the routers have only one physical connection to the Internet, they are capable of having multiple virtual tunnel connections, as needed. The VPN-enabled router knows the networks that it's connected to, including both the physical connections and the virtual connections.

In this illustration, assume that these routers have only two interfaces: one on the public network (Internet) and the other on the private network (192.168.X.X). Network A's router knows its private network is 192.168.1.0,

[15]This is for those who either did not read the earlier chapters or may need a refresher. A subnet mask is made up of four octets having the values of 255, 252, 248, 240, 224, 192, 128, and 0. We will leave it as an exercise for the reader to figure out how we came up with just those numbers and nothing else in between. A 24-bit mask has the first three octets filled, so it would be 255.255.255.0 in decimal dot notation. If 0 is the network address and 255 is the broadcast address for the network, then a 24-bit mask network can have 254 distinct addresses from 1 to 254. If more addresses are needed, it is time to either do some creative subnetting or add routers with additional subnets. Both work and have their respective advantages and disadvantages. If you do not know the differences, it is time for a review of the earlier chapters.

so any packets it receives for this network would be passed into the network and directed to the device that responds to the ARP[16] for that particular IP address. For addresses located on either the 192.168.2.0 or 192.168.3.0 network, the VPN-enabled router knows it must use one of its virtual tunnel routes.

In order to route network traffic over the Internet, actual addresses that are routable over the network need to be used. This illustration is using private IP space addresses, which are not routed, so how does a packet with a destination that is in the private IP address space get routed? This is where encapsulation comes in. The source and destination addresses are the actual physical public addresses of the VPN-enabled routers. So, if a network node on network A wants to send network traffic to network C, the packet is encapsulated within the tunneling protocol and the destination address is set to 3.3.3.3 with a source address of 1.1.1.1. Although these are not the actual addresses of the physical devices that are sending and receiving the data, the VPN-enabled router on the sending node knows that it must encapsulate these addresses, and the receiving VPN-enabled router knows that it must de-encapsulate the packet to ensure its delivery over the private network the router is connected to.

Using encapsulation combined with strong encryption can safeguard and maintain data integrity even while passing the traffic through a hostile environment. If packets are intercepted, the information they contain will not be easily decrypted and thus will not be compromised. Tunneling does not necessarily travel over the Internet, although that is where it is used most frequently. VPN tunnels can be used within an organization's intranet, as well.

**RANDOM BONUS DEFINITION**

ISDN (integrated services digital network) — A telecommunications standard used for the transmission of voice, data, and video using digital telecommunications over ordinary telephone lines.

### 14.2.2.1 Point-to-Point Tunneling Protocol

The Point-to-Point Tunneling Protocol (PPTP) does not provide safeguarding of data or perform encryption upon the data it carries. If encryption is required, PPTP depends on the protocol of the data that is being tunneled. PPTP is a peer-to-peer PPP session with generic routing encapsulation (GRE). For the GRE session to be initiated and maintained, a second session is required on TCP port 1723. Due to the need for a second session, PPTP is difficult to pass through a firewall since it uses two separate sessions for tunnel creation.

[16]Once again, if you are having a problem with this you better go back for a review. Like, what is Address Resolution Protocol? We ain't telling — you tell us.

The popularity of PPTP, even with its issues with security and its inability to be passed through firewalls, is due to the fact that it was the first tunneling protocol supported in Microsoft's Dialup Networking, initially released with Microsoft's Windows 95 operating system. Authentication to initiate a PPTP tunnel is performed using MS-CHAP or EAP-TLS,[17] which require the use of client certificates. Using a weak password with MS-CHAP is a security risk due to the possibility of the password becoming compromised. EAP-TLS adds further security but requires that clients provide a certificate. Microsoft supports both client and server EAP-TLS implementations in its Windows operating system. The progression path from PPTP VPN tunneling is usually to L2TP or IPSec.

#### 14.2.2.2 Layer 2 Tunneling Protocol

Similar to PPTP, L2TP does not encrypt the data carried within the packet, but depends on the protocol being carried within it to provide encryption and maintain data confidentiality. Although L2TP behaves as a Data Link layer protocol (Layer 2 of the OSI model), it is in reality a Session layer protocol (Layer 5) using UDP port 1701. The entire L2TP packet, including header and data, is transmitted within a UDP datagram.

Because L2TP lacks the capability to maintain confidentiality, it is often deployed with an implementation using IPSec, referred to as *L2TP/IPSec*. L2TP/IPSec negotiates with the IPSec Security Association (ISA) through Internet Key Exchange (IKE). This is accomplished over UDP port 500 using preshared keys, public keys, or certificates for both endpoints of the tunnel. Because L2TP is encapsulated, there is no need to open port 1701 on any firewalls that may be in the path between the endpoints creating the tunnel. The L2TP packet is totally encrypted within the IPSec packet and allows for the secure transport from endpoint to endpoint. In this implementation, IPSec provides for a secure channel within which L2TP can tunnel safely.

#### 14.2.2.3 Internet Protocol Security

IPSec is really a suite of protocols for securing Internet communications, utilizing authentication and encryption within each packet of the data stream between two network nodes. Each endpoint of an IPSec connection negotiates the type of authentication to be used at the start of a session and the form of encryption to be used while the session is maintained. IPSec resides at the Internet layer of the TCP/IP model, which is comparable to the OSI Layer 3 Network layer. Upper level applications above these layers can be protected easily within IPSec, since there is no special design consideration required for its use.

[17]EAP-TLS is Extensible Authentication Protocol–Transport Layer Security

IPSec has been embedded into network edge devices such as VPN-enabled routers. Two of these devices can be configured to establish a secure tunnel between them, passing traffic from one protected private network to another. A tunnel of this type is normally referred to as a *peer-to-peer tunnel* since each endpoint of the tunnel is aware of the other endpoint's IP address. In instances where one endpoint is unable to have a static endpoint address, there is a method for tunnel establishment that is referred to as *aggressive mode tunneling*. This is similar to a client connection but allows for the passing of traffic for network addresses that have been defined within the tunnel definition of its security association (SA). At least one endpoint must have a static public address. It is not possible for both ends to be unknown since the IP address is used to provide part of the security for tunnel establishment. The dynamically assigned IP address endpoint knows the peer it is connecting to. The statically assigned IP address endpoint depends on authentication schemes to verify the identity of the aggressive mode peer requesting the connection. Because the dynamically assigned IP address endpoint is not known to the statically assigned IP address endpoint, the tunnel-initiation request has to be started from the dynamically assigned IP address endpoint. Figure 14-15 illustrates an IPSec deployment over multiple sites.

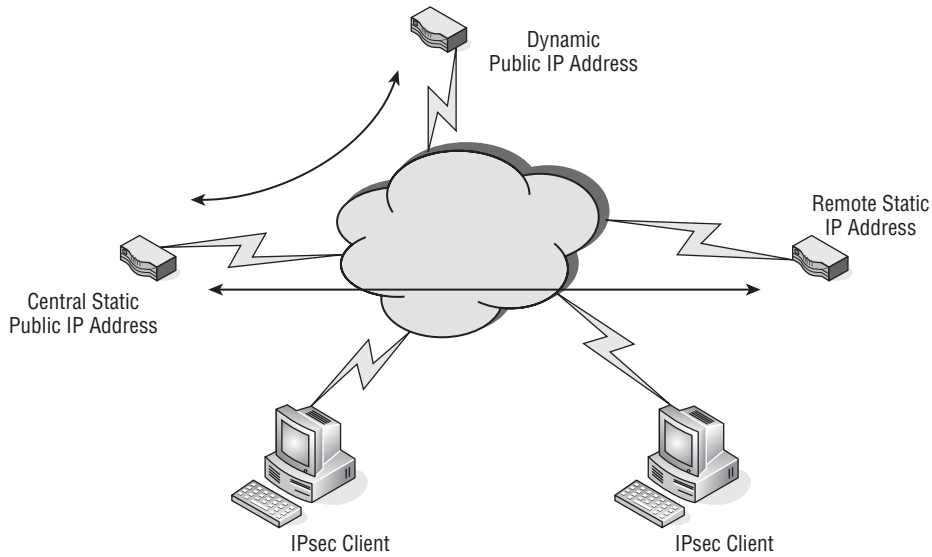


**Figure 14-15** IPSec deployment

This figure shows that there is a central site with a statically assigned public endpoint IP address and two other sites to which it has VPN connectivity. One endpoint also has a statically assigned public IP address, and the other endpoint is connected to a service that only allows it to have a dynamically

assigned public IP address. These two endpoints know each other and can reach each other because of the statically assigned public IP addresses. Since this is a peer-to-peer connection, either endpoint can initiate a tunnel when there is a need to pass traffic between the two locations. In essence, the tunnel is an on-demand connection. If there is no traffic passing between the endpoints, and if the tunnel idles for a period of time, it can be torn down by a configurable idle timeout setting. Because either end may bring up the tunnel on demand, there is no need for a keepalive[18] to maintain the tunnel in a secure operational state.

The behavior of the aggressive mode tunnel is different from a main mode, peer-to-peer tunnel[19] since the tunnel can be reestablished only from the dynamically assigned public IP address endpoint, sometimes referred to as the *remote endpoint*. If there is an idle timeout and the tunnel is dropped due to inactivity, the central site's statically assigned public IP address endpoint will not be able to bring up the tunnel. If the central site needs to get to portions of the remote network when there is no one at the remote site, a keepalive can be used to allow the tunnel to remain up even when no real traffic is being passed over the tunnel.

Remote users can be located anywhere there is an Internet connection available to connect to any VPN-enabled router on which they have an account. In Figure 14-15, if either IPSec client user has an account on both of the statically assigned public IP address–enabled routers, they can connect to that site with a secure IPSec tunnel. However, there are sometimes special conditions that need to be met if the client is connected on a private network that is using network address translation (NAT) to hide its private IP address space from the Internet. For these conditions, the VPN-enabled router would need to be able to handle ''NAT traversal'' (NAT-T).[20] There are different schemes for how this is handled and VPN-enabled routers are different in how they handle these scenarios. However, it is something to become familiar with if you do not want to be awakened at 3 AM when the CEO who traveled to Hong Kong is unable to get his e-mail from his hotel room.

[18]Keepalive is a mechanism that prevents a tunnel from being shut down due to traffic inactivity between the tunnel endpoints. If a tunnel is idle, there is no certainty that it remains secure. Keepalive traffic maintains the tunnel in an active state and both endpoints are secure in that each is connected to its secured peer.

[19]A peer-to peer tunnel is one where both endpoints have equal capability to initiate and establish a tunnel with its remote endpoint. Both endpoints are aware of the other's IP address.

[20]NAT traversal is when an IPSec client tunnel is created from a private IP address that is not routable over the Internet. How NAT is performed by the router local to the client PC determines how the VPN router that the client is attempting to connect with handles the encapsulation of the returned encrypted packets. There is no set standard on how NAT traversal is accomplished between client and VPN router; it varies from manufacturer to manufacturer.

IPSec tunneling is flexible and fits many VPN schemes. It supports a wide variety of authentication methods and has strong encryption capability. Previously, IPSec used a 56-bit data encryption standard (DES). Now triple DES (3DES) is fairly common. There are IPSec clients for many platforms and operating systems, from desktops to laptops to handheld devices. IPSec has been widely deployed and will be with us for some time into the future.

**POP QUIZ**

Think of a network that you are aware of and how a VPN solution may be beneficial to the organization. How would you implement the VPN design? What tunneling protocol would you select and why?

## 14.3 Chapter Exercises

1. How would you best protect network elements that are located in a remote area away from the network operations center?
2. Name a service that can provide not only user authentication but determine the amount of time a user has been logged in.
3. What is a digital signature associated with?
4. Which tunneling protocol was first supported with Microsoft's Windows 95 operating system?

## 14.4 Pop Quiz Answers

There are no hard-and-fast answers to these questions. You should use them as an exercise to attune your mind to what it takes to secure a network.