1

Defining the Scope

Growing up you may have played the Milton Bradley game *Stratego*. The concept of the game is fairly simple: There are two players, each player has a flag, and the goal is to defend your flag while trying to capture the enemy's flag.

Stratego is a good analogy for network security. When designing a network security strategy your ultimate goal is to protect your corporate infrastructure, your flag, from attackers—both internal and external.

This book is designed to help you develop your network security strategy. Because each company has a different approach to network security, it does not attempt to define a specific strategy for you. Instead, the book concentrates on ways to help you reach your ultimate goal in network security: protecting your network infrastructure.

This chapter defines the scope of network security, as covered in this book, discusses various types of network security, and outlines some of the costs associated with lax security policies.

In addition to the theoretical aspects of network security, it covers common network security mistakes, and describes a typical corporate network.

Pay close attention to the diagrams in the last section of this chapter. They will be referred to throughout the book, as potential security holes are plugged. The final chapter will present a redesigned and more secure network.

1.1 What is Network Security?

The first step in any discussion about network security is defining network security. If you ask 10 different administrators for a definition of network security you will probably get 10 different answers.

For the purpose of this book the definition that originated from the United States National Security Agency will be sufficient: Network security is the protection of networks and their services from unauthorized modification, destruction, or disclosure. It assures that the network performs its critical functions correctly and there are no harmful side effects.

This is, admittedly, a very broad definition, but a general definition better prepares network administrators to deal with new types of attacks. If a network security plan is broad in scope, then the tools will be in place that help deal with new types of attacks. Some security incidents are obviously network problems. A distributed denial of service attack (DDoS) is an obvious network issue. A DDoS attack occurs when multiple systems flood a network, or network device, with traffic (such as Ping floods), rendering it unusable by legitimate users. DDoS attacks have to be stopped before they reach the server; in other words, at the network level.

On the other hand, e-mail worms are an example of an attack that is more of a gray area. An e-mail worm is a file that is sent as part of an e-mail. The file exploits security holes in popular e-mail programs to cause damage to a machine's file system, and then sends itself to other people via the address book, continuing to wreak havoc. At first glance an e-mail worm might be considered a problem for server administrators to deal with, but worms, in addition to flooding servers, also clog the network, and, in extreme cases, may force you to remove your network from the Internet, while the worm is being dealt with.

1.1.1 Network Security and Compromise

As with all security, network security is about compromise. As shown earlier, even defining the scope of network security involves compromise. A network security policy is not developed in a vacuum. Network administrators have to work with other departments, especially a company's legal department, and

within the confines of a limited budget to determine the scope of an organization's network security policy.

Unfortunately, compromise often leaves a network administrator in the position of being damned if you do, and damned if you don't. Network administrators often find themselves in the hot seat for incidents that might have been prevented if the requested budget had been allocated.

Network security compromise is often a combination of education and risk management. Security personnel have to remain abreast of the latest security vulnerabilities, and communicate new information to others in their group, and often to the chief information officer (CIO), directly or through the normal chain of command. The CIO is then responsible for communicating the information to the rest of the organization.

NOTE

Throughout this chapter you will see references to the CIO. Depending on the size and structure of your company, the duties described may be handled by a chief technology officer (CTO) or an information technology (IT) manager.

When communicating security information to others in the company, it is often necessary to act like a salesperson. Security problems should be explained in terms of benefits not features—explain what can happen rather than the technical aspects of an attack. If a new security hole may allow DDoS attacks against a server, don't discuss the minutiae of the ISO OSI reference model or the Transmission Control Protocol (TCP). Instead focus on the fact that if this security hole is exploited it may cause your website to become unreachable by legitimate users.

Another tactic is to explain problems in terms of cost. If bandwidth is billed using a burst model (e.g., you have a 10-megabit connection, but can use up to 45 megabits) a DDoS attack can cause the organization to use its fully allotted bandwidth, thereby incurring a quantifiable additional expense.

In fact, the more often a security risk can be quantified, the easier it is to convince others to approve, or facilitate, the ability to act.

1.1.2 Risk Management

Quantification of network problems also allows network administrators to better handle risk management. Risk management is the process of assessing the potential threat from a security risk.

Risk management also means understanding when cost is not a factor. While this section largely focuses on determining the true cost of implementing security solutions, it is important to remember that there are some solutions that are so important they need to be implemented no matter what the cost.

Effective risk management requires an understanding of the full impact of every security threat. Full understanding of a risk gives network administrators the ability to weigh the true costs involved in not fixing a security hole. For instance, if mail servers are left unsecured, so anyone can send a message through them, there is a potential security hole that has a high risk of being exploited. Risk management involves looking at the costs of fixing the server versus not fixing it. The cost of fixing the mail server is relatively minor: Simply do not allow anyone outside the local network to relay through the server, or, if an organization has many remote users, implement a security system that requires people to authenticate before they can send mail. The cost of not fixing it is great. There is the obvious cost of someone using your server, and network connectivity to send mail to millions of people. But there are also administrative costs involved in a situation like this: angry e-mail from people who received the mail, losing the ability to send mail to some people because your mail server is blacklisted, and having to restrict access to the mail server anyway.

In April 2002, the FBI and the Computer Security Institute released the results of their "2002 Computer Crime and Security Survey." The survey, which collects data about security practices from randomly selected companies, provides information about the frequency of common network attacks. Table 1.1 lists the percentage of companies that reported successful attacks.

Computer worms are by far the most common type of network attack detected[1] and reported. The operative words are *detected* and *reported*. Obviously,

1. The second most common type of attack, not listed in the chart, is one originated internally by an employee, or group of employees.

Table 1.1 Reported Network Attacks

Type of Attack	Percentage Reporting Successful Attacks
Computer virus/worm	85
System penetration	40
Denial of Service attacks	40
Web server penetration	38

not all attacks are detected—and even some that are go unreported—so the percentages may not reflect the true number of network attacks experienced by these businesses.

Some companies feel that there is a stigma associated with network attacks and, despite the fact that network attacks are a common occurrence, the Computer Crime and Security Survey continually suffers from underreporting by companies.

Why is data like this important? It helps to give network administrators an idea of how an organization's resources should be distributed when developing a network security strategy. If it is known that a company is twice as likely to be the victim of a virus or worm than any other type of attack, server administrators can plan appropriately.

Many companies use a form of risk profiling to determine the cost of implementing a network security policy. Risk profiling involves evaluating a security risk from four perspectives and using the number gained to assign a priority to each threat.

As with the other aspects of risk management risk profiling has to be handled by a security group, and needs direct involvement from the CIO, senior management, and the legal department.

The risk profiling method developed by the National Institute of Standards and Technology involves creating a matrix that evaluates the threat, visibility, consequences, and sensitivity of a potential threat. This type of risk assessment fits well into most network security models, as discussed in the next chapter.

To create a risk profile, first create two charts (Table 1.2 and Table 1.3).

Table 1.2 Risk Profiling: Threats and Visibility

Threat	Rating
No currently identified threats	1
Unknown, or multiple exposures	3
Active threats, multiple exposures	5
Visibility	**Rating**
Very low profile, no publicity	1
Occasional publicity	3
Active publicity	5
Multiply threat value by the visibility value.	

Table 1.3 Risk Profiling: Consequences and Sensitivity

Consequences	Rating
Consequences have no cost, are within budget, or the risk can be transferred.	1
May impact internal functions, cause budget overruns, or there may be opportunity costs.	3
External functions may be impacted, and revenue loss will occur.	5
Sensitivity	**Rating**
Part of the cost of doing business, no organizational impact.	1
Unacceptable impact for a specific business unit and good-will costs.	3
Unacceptable management costs, and business relationships affected.	5
Multiply consequence value by the sensitivity value.	

Apply all four measures to a risk; multiply the threat and visibility values. Multiply the consequences and sensitivity values. Add the two results, and you have a risk profile.

After the measures have been applied to a risk it can then be assigned to one of the three categories in Table 1.4.

Table 1.4 Risk Profile: Final Assessment

COMBINED VALUE	RISK PROFILE
2-10	Low
11-29	Medium
30-50	High

Also note that this type of risk profiling should be included as part of any new network project completed. Analyzing and understanding security risks inherent in a new project is important to minimize future security risks to your company.

1.2 What Types of Network Security Are Important?

When a company first sets out to create a network security plan, there are usually two questions asked: Where should we start, and what is the most important part of the network? The answers depend on many factors, and the answers are different for every network.

Generally speaking, one person, or department, will not be able to answer both of these questions and one department should not develop the network security policy. The network security policy, as all security policies, should be disseminated through the CIO, and should be approved by the legal department and signed off on by the heads of all other departments. Network and server administrators may be called on to develop the first draft of the policy, but it is up to senior management to finalize, implement, and enforce the network security policy.

There are some questions administrators can ask to begin the development of the corporate network security policy.

1.2.1 How Sensitive Is the Data?

Any business has confidential data. Whether it is the customer database, proprietary software, a product design, or some other sensitive data, there is undoubtedly something that has to be protected. Such data should always be your first priority when developing a security strategy. In some cases, especially for companies that deal with medical or financial records, there are legal ramifications for not properly securing this data.

Of course, core data is useless if no one can access it. Second to protecting the core data is protecting the means by which people within an organization, or customers, access that data. The lines of communication to data—the network—have to be kept available.

In addition, employee phone lists or human resource records, important data but not as critical, need to be protected. The protection for this information does not need to be as draconian as the measures you should take for your core data, but it absolutely must be in place.

The involvement of the CIO and other groups is necessary at all levels of network security. One group cannot be sure how to rank the various databases within an organization. Someone from senior management will need to assign ranks to all data sources, so it can be determined how limited resources should be deployed.

Of course the less sensitive the information is, the more difficulty there is in securing it. Employee phone lists generally need to be accessed by other people within the company, and an internal website is probably available to everyone.

In some ways, the more available the data, the harder it is to secure. It is easy to prevent anyone from accessing information. It is harder to allow only certain people to access information, and enforce those access restrictions.

1.2.2 Secure Your Servers

The first step in securing your corporate data is to secure the servers where the data is stored.

How you go about securing a server depends largely on what operating system you are running. There are some guidelines, however, you can follow that apply to any operating system and any server, no matter what its function. These

steps are discussed in greater detail in Chapter 12, but this should give you a good overview.

There are two levels of server security: access to the server and environmental control. Access covers who can access the server and how they can do it. Environmental control covers the level of access that users can have—what they can do once they are on the server. These two types of server security are intertwined. If good access policies are enforced, but all users are allowed access to system files after they have logged onto the server, a security breach is waiting to happen. Should an attacker gain access he or she would have no limitations on what he or she could do to the server.

A server access policy should:

- Control who can log into your servers.
- Never send clear text passwords.[2]
- Force minimum password lengths.
- Impose character restrictions on passwords (mixed case, numbers, and punctuation).
- Force passwords to be changed at regular intervals.
- Set a maximum number of login tries before locking out an account.

Once a user has access to a server, there should be environmental limits that prevent users from gaining unauthorized access to system files or secured data. A good environmental control policy will include:

- Running virus scanners on all servers, especially e-mail servers. If a virus never makes it to an end-user's system it can't spread.
- Using, whenever possible, single-function servers (e.g., don't use the same server for mail and web services).
- Not storing proprietary information on public servers (e.g., do not put your customer database on your web server).
- Disabling all unused services, and if possible uninstalling those services.
- Closing all ports not being used.
- Changing all default passwords.
- Deleting unnecessary user accounts.

2. Expect to see this comment about 30 times throughout the book.

- Limiting users who have administrative access to the server.
- Deleting any sample files that ship with installed programs.
- Storing user files separate from administrative files (either on a separate partition or file system).
- Logging all movements by administrative users.
- Updating the system frequently with vendor security patches.

These steps are a good start toward securing your server, and protecting the data on those servers.

1.2.3 Secure the Network

Of course, the sooner you can stop a potential intruder, the better. This is especially true when dealing with server attacks. Ideally, you would like to prevent a potential intruder from ever reaching your server. Later parts of this book discuss strategies for securing your network in detail. Here are some useful guidelines that should be implemented on any network to help stop attacks:

- All machines in the network, except for the edge routers, should be behind a firewall.
- Authenticate all network protocols in use on the network (BGP, OSPF, VRRP, etc.).
- Restrict access to secure parts of the network by Media Access Control (MAC) address.
- Do not allow external traffic into the secure network areas.
- Use virtual local area networks (VLANs) for added levels of switch security.
- Change default passwords.[3]
- Use virtual private networks (VPNs) for employees who need to access sensitive information remotely.

These are general guidelines that should help administrators start forming a network security policy that works for an organization. As the book progresses, the policy can be refined.

3. This is another comment you can expect to see repeated.

1.2.4 Monitor it All

Never be complacent when it comes to network security. No matter how great the security measures taken, the fact is that a skilled and determined hacker will probably find a way into your network.

If that does happen, it is best to know about it quickly, and be prepared to stop it. To do that, monitor everything on the network. Anything that may be deemed as suspicious has to be brought to your attention. Monitoring is discussed in detail in Chapter 16.

In addition to monitoring, extensive logging of network activity should take place. It is unrealistic to expect the administrator's staff to have the time to scour hours of log files every day, but if an incident does occur, good, uncorrupted log files will be essential in tracking down how security measures were breached, and in trying to track down the attacker. At that point, you will be grateful for extensive logging.

A good monitoring strategy involves collecting a lot of data, and recognizing patterns within that data that may resemble attacks. These patterns generate an alarm, which will allow administrators to manually investigate the network or servers, and determine if there really is an intruder, or if it is simply a logging anomaly.

Some security experts advocate the use of honeypots as part of a monitoring strategy. A honeypot is a system that is intentionally left open to attract potential intruders. An attacker takes the bait and tries to break into the system. All interaction with the system is extensively monitored, and the honeypot becomes a tool to help network administrators learn more about security flaws in their system.

1.3 What Is the Cost of Lax Security Policies?

There are really two costs involved with lax network security: quantitative and qualitative. Quantitative costs, the ones most often discussed, are those that have the most immediate impact on the corporate bottom line, but qualitative costs can be just as important to a company in the long run.

According to The Yankee Group, network attacks accounted for $1.2 billion in lost revenue in 2000. That number doubled in 2001, and is expected to double again in 2002. Lost revenue is an example of a quantifiable cost of a security incident.

There is no universal formula to calculate the quantifiable costs of a network attack. There are, however, some commonalities that you can use to help develop your own, internal, formulas.

Some of the costs are easy. If you have an e-commerce site that is interrupted by a DDoS, or an attacker manages to gain entrance to one of the servers, forcing you to take your website offline for *X* number of hours, then one of your quantifiable costs will be the amount of revenue lost during that time. If your site normally generates $100,000 an hour, and it was offline for six hours, then one of your costs was $600,000.

Loss of revenue is not the only quantifiable cost. If it took you six hours to restore the website from backup and rebuild the database, then time becomes a quantifiable cost, as does the time spent researching the incident and reporting it to the proper authorities. There is also the cost involved in implementing a security fix, so a repeat attack cannot happen.

In addition to time, it is necessary to calculate the lost productivity of other groups within your company. If a design team made changes to the site after the last backup, then their changes will all have to be redone; their time is another quantifiable cost.

Qualitative losses are more difficult to measure, but can be just as important, and increase with the severity of an attack.

Using the example of an e-commerce site again, if someone were to force the website offline, in addition to the outlined quantitative costs, there are several qualitative costs. The most obvious is the loss of future customer revenue, and, depending on the severity and length of the attack, the loss of customer confidence.

If a customer cannot get to the site, he or she visits a competitor's site, has a good experience, and not only is the revenue lost, but future revenue may have been lost as that customer may continue to visit the competing site. If the attack is particularly successful, an attacker may gain access to your customer database, which is often enough for the attack to make the news. Now, on top of the potential loss of future revenue, other customers may not feel comfortable returning to the site, and potential customers may never shop at the site. There is also the added, quantifiable expense of hiring a public relations firm to deal with the problem.

A final qualitative cost is the loss, or delay, of future revenue from projects that were put aside because of the time spent dealing with an attack. If six hours is spent restoring a compromised system, that puts at least a six-hour delay on other projects. If the majority of time is spent dealing with security issues other projects may face an indefinite delay or cancellation. The revenue that would have been gained from those projects is now lost.

1.3.1 The More Severe the Attack, the Greater the Cost

It may seem like an obvious statement, but it is important to remember. The more severe an attack is—the further an attacker is able to penetrate into your network—the greater the cost, both in terms of qualitative and quantitative expenses.

A successful attack against one e-commerce website is relatively trivial, compared to more extensive attacks.

As mentioned earlier, an e-mail worm can paralyze an entire network, to the point of having to shut down e-mail servers and even force a company to disconnect from the Internet. Such an attack can cost a large company several million dollars in lost time and productivity.

Undoubtedly the most expensive attacks against a company are those that compromise data confidentiality and integrity. The compromise of confidential data, such as an e-mail system, corporate intranet, or a customer database can have long-term negative consequences. An attacker who gains access to these tools may not disrupt your network, but will have proprietary information that can be sold to competitors, or used to try to blackmail the company. If this attacker is discovered days, weeks, or even months after he or she has gained this level of access to your network, the cost to track down how the network was breached, and to find all of the security holes, can be extraordinary. Not only will you have to plug the initial security hole, but also each server and network device will need to be thoroughly audited to determine if the intruder left any trapdoors that would allow easy entry back into the network.

Data integrity attacks occur when an attacker gains access to—and modifies—confidential data. Sometimes the modifications are puerile and juvenile, such as defacing a website. Unfortunately, if an attack is targeted specifically to your company, data modifications can be more subtle, and their ramifications greater.

It is almost impossible to calculate the costs of a data integrity breach. Having to audit an entire customer database or verify the validity of confidential customer information can cost millions, not to mention the other costs normally associated with these attacks.

Data confidentiality and integrity attacks bring in the possibility of two new costs associated with security breaches: lawsuits and fines. If confidential information about the customer database or dealings with other companies is leaked, an organization may be open to a lawsuit. Even if it can be demonstrated that reasonable security measures were taken there are still legal costs associated with the lawsuit, as well as the aforementioned negative publicity and loss of customer confidence.

Depending on the type of data that is breached, a company may also be fined by the government. There are several bills before the United States Congress that would fine companies that do not meet minimum standards for network security. Some of these bills would allow companies to be fined up to $1 million if their networks are successfully breached.

1.3.2 Creating the Formula

Creating a company-specific formula that will help measure the cost of an attack is essential. If an organization is going to be able to implement a new security policy, you have to be able to show that the cost of not implementing it is greater than the cost of implementing it.

Again, it is important to keep in mind this formula should not be created by one person or group. The CIO, working in conjunction with senior managers from all departments, should develop the formula jointly.

The formula will vary depending on the type of attack for which the organization is trying to determine the cost. The best bet is to try to divide attacks into broad categories. In Chapter 2 common attacks will be covered in detail. For now, divide attacks into four categories:

1. Network attacks: Attacks not directed toward a server, such as DDoS attacks.
2. Worms: E-mail or web-based programs that travel from computer to computer on your network.

3. Attacks on peripheral servers: Attacks on servers that do not contain core business data.
4. Attacks on core servers: Attacks against servers that contain data that is essential to a business.

More categories can be added, or unnecessary categories can be deleted, depending on the needs of a business. (For instance, some organizations may want to add a category that specifically deals with an e-commerce site.) After categories have been created, the next step is to develop a basic cost structure for each category.

DoS attacks are a good example. If you have a firewall, or routing policy, that will block DoS attacks, then your costs would be limited to productivity losses from not being able to connect to the Internet while the attack was ongoing. If a routing policy that will lessen the impact of a DoS attack is not in place, productivity loss incurred while the network is unavailable may have to be factored. If a company generates revenue from the website, and it is located in a data center within the facility, then a DoS attack will cause loss of revenue from the website.

For each category created the goal is to develop as many fixed costs as possible. If it is known that it costs the company $100,000 an hour for every hour the website is down, that is a number that can be repeatedly factored into loss equations. If the company loses $90,000 an hour in productivity when the mail server is unavailable, that is also a fixed cost. Often, these numbers will be readily available from the appropriate departments.

1.4 Where Is the Network Vulnerable?

Before delving further into the book, it would be a good idea to assess network vulnerabilities. Being aware of some of the more common security problems found in a networking environment makes it easier to spot them on another network. A quick audit based on some common mistakes is a good start. As topics are covered in more detail, it should be easy to pick up other ideas to tighten security even further.

The most common mistake an administrator makes is using clear text passwords. Many administrators will disable telnet access to servers, but leave File

Transfer Protocol (FTP) access open, or they will use telnet to login into routers or switches, instead of creating a TACACS+[4] server. If possible, even e-mail login sessions should be done using encrypted usernames and passwords. Of course, encrypted logins have to be combined with a good password policy.

Domain Name System (DNS) servers are another commonly exploited vulnerability. The most popular program installed on DNS servers is the Berkeley Internet Name Domain (BIND). While recent versions of BIND have done a great job of increasing security controls, the vast majority of companies are still running older, less secure, versions of BIND.

Another common mistake network administrators make is to leave network passwords set to their default; this is especially true for the Simple Network Management Protocol (SNMP). The default passwords for reading data and writing to SNMP devices are generally public and private, respectively. Often administrators activate SNMP without thinking about the consequences of an attacker having full control of their routers.

Firewalls can also lead to poor security practices. Many administrators assume because they have a firewall in place their networks are secure. Firewalls do not solve all security problems. In fact, a firewall with poorly implemented rule sets offers little or no protection for a network. A firewall with good rule sets is important, but it is only a small part of a security policy.

Whenever possible, use managed switches instead of hubs. A managed switch offers security features such as VLAN control and MAC address control. These additional security features enable you to control what machines have access to your network, and can even allow you to control traffic within your network.

A wireless LAN (WLAN) is an incredible technology: It frees employees from their offices or cubicles and allows them to connect into the network from anywhere in your building. There is also a host of security concerns that need to be addressed before implementing a WLAN. Some of the security issues inherent in WLAN technology include the ability to easily port sniff other users connected to an access point, easy entry to your network for just about anyone, and of course, the use of an insecure default password.

4. TACACS is the Terminal Access Controller Access Control System, is documented in RFC 1492, and is an authentication and logging system.

1.5 The Network

The best way to learn is by example; to that end this section presents a typical corporate network for a 100-person company. This network is fairly insecure. Forging ahead, various chapters in the book will capitalize on the vulnerabilities in the network and demonstrate ways to correct them. Of course, there is no one correct security model. Security needs vary from company to company, but showing how to spot and correct weaknesses in corporate security helps administrators find holes in their own networks, and helps create better methods for dealing with security issues.

NOTE

In this example, the netblock 10.10.0.0 255.255.255.0 is used. This is one of the netblocks that has been reserved by RFC 1918 for private use. Think of it like using the 555 prefix for phone numbers in movies. The address block will function like a normal netblock, but the addresses are not routable across the wide area network (WAN).

1.5.1 The Network Infrastructure

Figure 1.1 shows the network infrastructure for this company. It is fairly simple: a router connected to a firewall that has three interfaces: public—to the router, and two private interfaces—one to the employee network and one to the server farm.

The firewall rule set is also fairly simple for this network. No traffic is allowed in to the employee network, all traffic is allowed in to the server network. The rules for the server network were tighter, but as new software was added to the servers in the server farm, it became difficult to keep track of which ports needed to be opened so all ports were opened.

The company uses a TCP/IP network infrastructure, but no auditing has been done to see what other network protocols are running on the machines. The netblock 10.10.10.0 255.255.255.0 (a Class C block of addresses) is assigned to the company. The IP addresses have been distributed throughout the network without subnetting them.

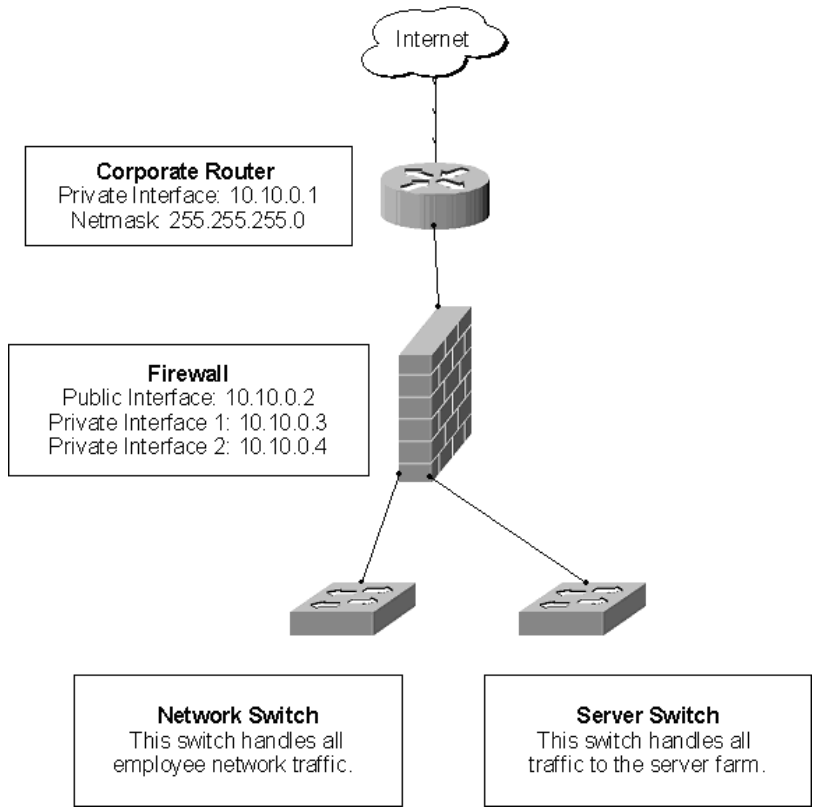

Figure 1.1 The network infrastructure

Finally, even though they are using managed switches, the network administrators have not assigned different VLANS to the ports on their switches; all machines connected to the switches are using the default VLAN.

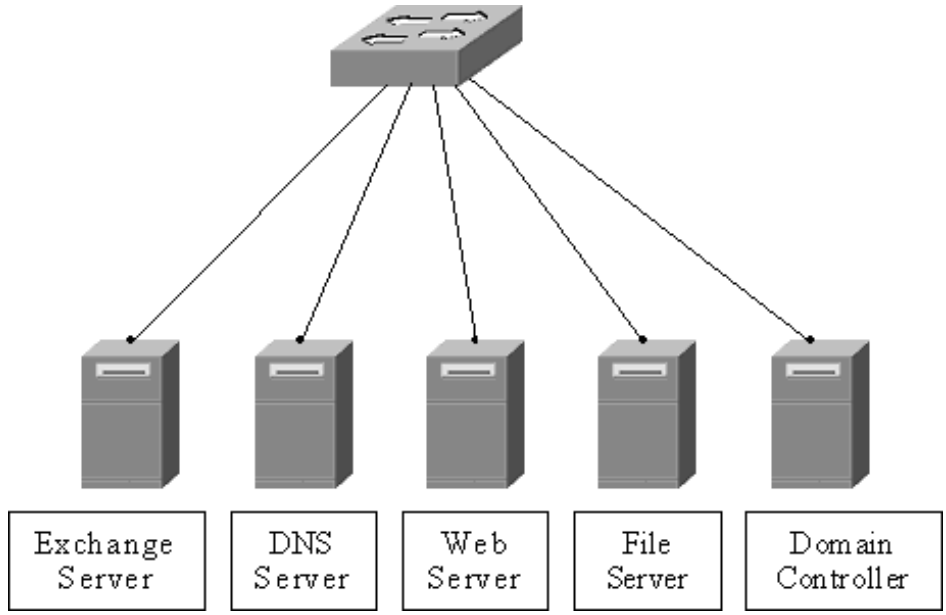

Figure 1.2 The server farm

1.5.2 The Server Farm

The server farm (Figure 1.2) consists of five servers; all but two perform unique functions. The File server also doubles as a Remote Access Service (RAS) server that allows employees to dial in to the network from home, while the domain controller doubles as a monitoring server.

The file and exchange servers and the domain controller are all running Windows NT, with service pack 4 installed. The web and DNS servers are both running Red Hat Linux 6.2.

New accounts are created on an as-needed basis, and there has been no auditing of account information to date.

1.5.3 The Employee Network

Various employee groups, such as human resources and accounting, are connected via hubs to the network switch (Figure 1.3). The employees use a mix of Windows 98, Windows NT Workstation, and Windows 2000 Professional workstations. Again, there has been no workstation auditing to date, and no one has set a policy to limit the type of workstations that can be added to the network. There is also no password auditing or policing system in place.

All workstations on the network are assigned IP addresses by the domain controller when they log onto the network.

The company is also experimenting with WLAN technology. The two conference rooms have been outfitted with access points that allow anyone with an 802.11b-enabled card or computer to connect into the network.

There are many gaping security flaws within this network. As each area is delved into more deeply, they should become more apparent.

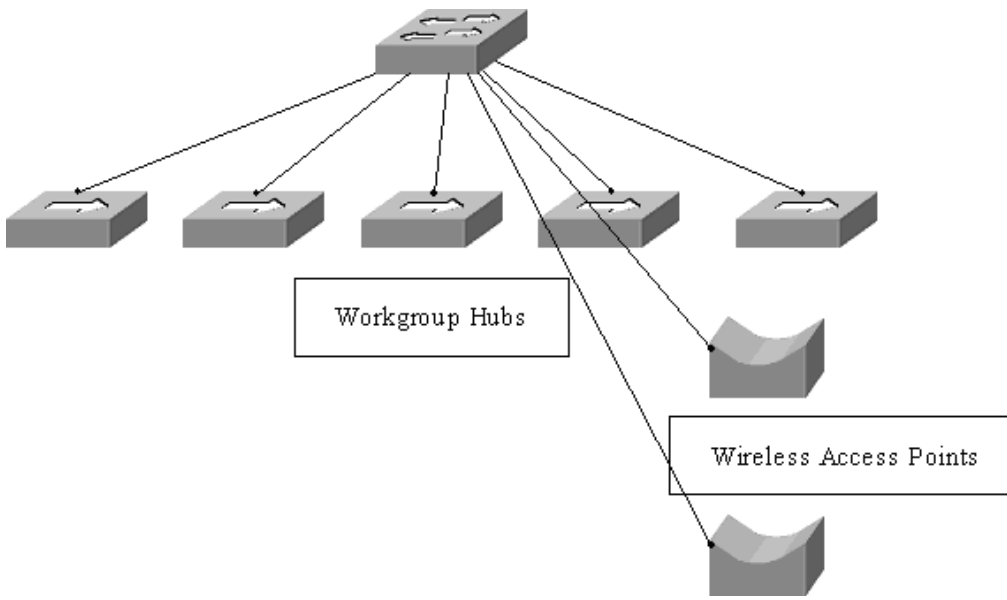

Figure 1.3 The employee network

1.6 Summary

Because networks are so commonplace within organizations, network security is important for all administrators. Maintaining good network security is a full-time task that has to involve the cooperation of all employees within an organization.

Because network security is so important, and involves all aspects of day-to-day operations, it is important that security policies be communicated from the top down, and that all managers are involved in the planning of network security policies.

Many organizations, especially small ones, don't feel that they need to worry about network security. The truth is, any organization that is publicly connected to the Internet has to make an effort to secure its border.