Chapter 2

Behind the Crime

Solutions in this chapter:

- **Overview of Technologies**
- **Information Extraction**
- **Hidden Files**
- **Network Leakage**
- **Cryptography**
- **Steganography**
- **Malicious Acts**
- **The Human**

Introduction

In the previous chapter we laid the foundation for insider threat. We raised awareness and showed that it is a serious problem that requires serious attention. Insider threat has always been around, but with the widespread use of the Internet and computers, it has created a new arena that will allow insider threat to flourish like we have never seen it before. In an office full of paper there was always the potential for someone to copy a piece of paper that they were not supposed to copy, carry a paper copy of a document out of a facility, or tell someone what they read about. This problem of having a trusted insider use his or her access and knowledge to hurt the company always existed, but it was bound.

The methods that could be used before computers and the Internet to commit insider threat were limited:

- Extraction of a document in its original form
- Copy with potential miniaturization
- Faxing
- Verbal word of mouth
- Modification of information
- Destruction of information

To commit insider threat an attacker has to focus on compromising the confidentiality, integrity, or availability of critical IP across the organization. Even though there were methods to do this prior to the widespread use of computers and the Internet, they were controlled and bounded.

The easiest and most straightforward method, which still works today, is extraction of a document in its original form out of the organization. This could also entail physical relays, where someone gives the document to another department or group, who eventually transferred it out of the organization. One of the best and easiest ways I have seen attackers do this is via mail. Many organizations had guards that checked your bag when you left so it would be hard to walk out with a classified document in your briefcase. However, if I took that document from within the building and put it within an envelope with a stamp, no one would check it. There was a case where an organization would actually check mail and briefcases, but not FedEx packages. The reason is that there was a separate process and approval for being able to use FedEx, and when they set it up they were not thinking about insider threat. Yet an insider quickly found this as the best path for extracting infor-

mation from the company. Although we might not always think about it, attackers will always find the weakest link and exploit it.

This is one of the reasons why it is critical to have a document classification scheme in place. In this book we are talking about classification scheme in a general sense. It does not matter if you are a government intelligence organization or a Fortune 500 company; you need some way to be able to mark and protect information. The government might use the term "top secret," and a commercial company might use the term "company proprietary" or "trade secret"; however, the level of protection remains the same. The documents must be marked and done so in a clear and consistent fashion. If documents are not properly classified and marked, how would a guard or anybody be able to determine that you should not be leaving with a sensitive document? Some easy criteria have to be used so that spot checks can be performed.

Traditionally organizations would mark documents just at the bottom of the document. The problem with this method is if someone was quickly looking at the document, they might miss the marking at the bottom. Based on the limitations, organizations would then mark at the top and bottom of the document. Although this was easy to spot, it still had some limitations. The main limitation was that if someone made a copy of the document they could cover up or cut off the labels at the top and bottom and conceal the actual level of the document. This is why today companies use a watermark on the page that goes behind the text. This can be seen in the Figure 2.1.

Figure 2.1 A Watermark

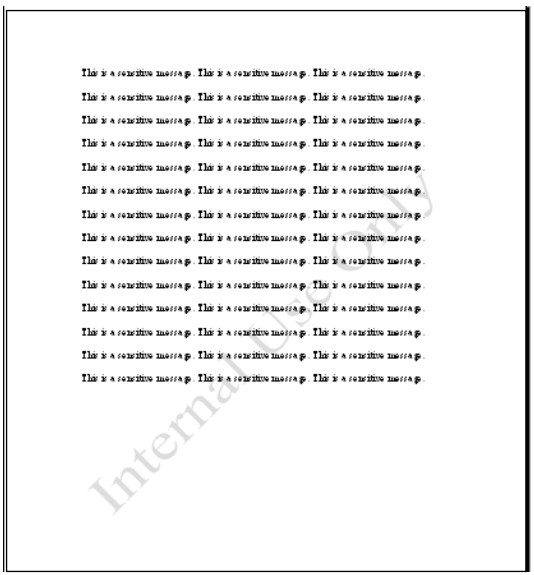

This makes it much harder for someone to try to remove and cover it up. I have seen cases where someone has tried to use white out to cover it up, but that does not work very well. Another method used is to mark each paragraph with a sensitivity label. Although this method works, it is harder for the person creating the document, plus a spot check could easily miss it because you would have to look much closer at the document.

A slight twist on this is to make a copy of the document and take the copy out with you. One advantage of making a copy is if the document has document control, where it has to be checked in or out if you left with the original, someone would know the document was missing. However, if I made a copy, I would be able to return the original to document control and take the copy home with me. The other advantage of copying is many copy machines have advanced features. These features could be used to make the document smaller in size and easier to conceal. If you have a single-sided document that is 100 pages, making a copy double-sided can cut the document in half.

Copy machines also having settings where you can reduce the contrast and make the copy either lighter or darker. I have seen cases where the watermark is so light that you could reduce the contrast by making a copy and removing the watermark. If you make a second copy with a darker contrast, the text will darken back up but the watermark will have been removed.

Many copiers also have size reduction capability. This is advantageous for many reasons. First it allows you to make a document much smaller in size. If I have a 100-page document and I reduce each page to 50%, now I can fit two pages per one printed page. If I use full duplex, I can reduce the document down to 25 pages. If I am a trying to remove a large number of documents, this can be the difference in slipping some documents in my briefcase and having to carry out a big box, which raises the suspicion level.

In addition to using copy machines to do this there are miniaturized cameras that have been developed just for this purpose. The cameras are very small and can easily fit within a pocket or embedded within a watch or a pair of glasses to make it very hard to detect. Now the attacker can take pictures of the document, leave with them on the tiny device, and then develop the pictures once they are at a safe facility. This technique and method has been around for a long time and is best illustrated in the movie, "Wall Street." It is actually not a bad movie to watch, but in the movie, the main character wants to commit insider trading. He works at the company but does not have access to his boss's office. He uses his internal knowledge of the organization to figure out where the sensitive documents are kept and when his boss leaves. Now he knows what he has to do and when he has to do it; he just needs to get access. To get access he interviews with the cleaning service that cleans

the offices and gets a job working in the evening. To perform his job with the cleaning service he needs keys to all the offices. Since the cleaning company times you for each floor, he actually cleans some of the offices very quickly so he can spend more time in his boss's office and not have it look suspicious. When he is in his boss's office, he pulls out a miniature camera that he bought at the local electronic store and starts taking pictures of all the sensitive documents. Although this method works, I have also seen cases where you could look through the garbage to extract data. An ideal case that would also fit well is if I have access and can make copies but I cannot leave with a large amount of documents. I could also get a job with the cleaning service and us the large trash can on wheels as a means for getting the document out of the building. I then come back later that night and extract the documents from the dumpster. There are lots of opportunities and unconventional methods people will use to accomplish their task of insider threat.

Although simple and straightforward faxing is a method that, despite all the technological advances today, still has its spot in the office. Prior to computers and even today, if you needed to send a document to someone very quickly, faxing was a method for doing this using traditional phone lines. Most organizations have or had fax machines on almost every floor, so it would be relatively easy for someone to be able to take a document and extract it out of the company using the fax. Once again, even companies that have guards checking bags when you leave usually do not have checks at the fax machines. One company I know of actually had checks at the fax machine. The fax machines were locked in a person's office and they had to perform all faxing. This allowed them to be able to review the fax prior to sending it out. Even in this case, the system could be bypassed if this person was on vacation or not in the office. However, there was even an easier way that an insider found. Fax machines could be purchased relatively cheaply. This person bought a fax machine, plugged it in at the office, shut the door, and began faxing, allowing him to bypass all corporate controls. In this case, the funny part about it is the insider expensed the fax machine back to the company and the boss approved it.

Although fax machines and miniature cameras present their own challenges, at least there is still physical IP leaving the organization so proper perimeters and checking can still potentially flag it. The harder problem is when there is nothing to catch or stop, and the information is stored in an intangible format—someone's brain. Unless technology advances there is no way to scan someone's brain and make sure that he or she is not leaving with anything critical, and even if we could, it would be much harder to determine that the person intends to do harm with that information. Verbal word of mouth is probably the easiest and still one of the most common methods for extracting data out of the organization. I read a sensitive document, remember key facts, and walk out of the organization with nothing tangible.

No matter what checks are in place, there is no way they can stop you. Then you would meet up with someone at a remote location and give them a data dump of the critical data.

There is no way someone will remember a 50-page document word by word, but most people can remember the key or critical facts that could still result in a compromise. You can perform this experiment with a book. After someone reads a book, ask them to recite the entire book, word by word—they will not be able to do it. However, if you ask them for the general plot, the critical facts, and even the key characters, I am sure they will be able to do it. This is a big problem and comes down to controlling access. If someone is not able to read a document, they will never be able to repeat it to someone else. This problem is where the saying during World War II came from: "Loose Lips, Sink Ships." People have to learn that what happens at work stays at work.

All the methods we talked about apply mainly to disclosure of information: someone who has special access revealing sensitive information to someone who should not have access. This is the main thrust and focus of insider threat, and it requires removal of information from the organization in some format, which could be hard and risky in some organizations. In cases where you cannot extract data from an organization, having access within the walls of the organization, an insider can still cause harm in two other areas: modification of information (integrity) and destruction of information (availability).

Ideally as an insider spy, my ultimate goal is to be able to get a copy of the data; however, I do not want to blow my cover and get caught. Therefore if I cannot get the information out of the company, if I can modify the data I still provide a better service for the ultimate company I am working for by causing harm to its competitor. If I modify sales projection, I can cause reputational damage. If I can modify a proposal with the wrong numbers, I can cause financial harm to the organization. If I modify critical formulas or controls for critical systems, I could cause lawsuits that would result in both financial and reputational harm. Though integrity is not always the first choice, it can still provide benefit.

Destruction of information is typically the main form of a disgruntled employee who is not looking for a direct benefit but trying just to cause harm to the organization or punish the organization for the way he or she has been treated. Destruction, depending on the level and technique that is applied, can be very easy and simple to do. It can also be used as a last resort if extracting the data out of the organization does not work. Ideally a competitor wants to be able to get a copy of the latest proposal a company is submitting so they can underbid them. However, if you cannot find out that data, but you can destroy all copies of the proposal two days before it is due, in essence you have accomplished the same net effect. By destroying the pro-

posal you will not be able to bid, and now since your company is not submitting a proposal I no longer have to worry about outbidding you. In some cases this could offer a big advantage because if you know you will be the only bid then price really is less important than before.

It is important to note that although these methods were used before there were computers, these methods will still work today. This is why organizations have to be so careful and make sure they look at the big picture. I have seen companies that are so focused on computers, they forget about the obvious and let people walk out the front door with documents in their brief cases. Their cyber security is so secure that you would not be able to do this electronically, but their physical security measures are so overlooked that an attacker could exploit that area with ease.

Although these traditional methods we just discussed could still cause harm to the company, it was something you could get your arms around and defend and detect. It was a known problem that was bounded in its methods and complexity. Therefore there were methods that could be used to counter this activity and the methods were known and understood, and if performed correctly, worked.

The following were some of the methods that were used to protect against this type of attack:

- Marked, color-coded covers
- Locks
- Guards
- Locked rooms with checks
- Spot checks
- Separation of duties

There is no such thing as 100% preventive measures, but the trick is to try to defend as much as possible and hope with proper awareness and training that most people will do the right thing.

One way to prevent information from leaving the organization is to carefully mark documents, both top and bottom, and with a watermark, as we discussed earlier. Something else that is helpful is putting a cover sheet on every document. A red cover sheet means the document should never leave the company, a yellow cover sheet means it can leave only if it is properly secure, and a white cover sheet means the document contains no sensitive information. Now you can argue that someone can just remove the cover sheet. However, a document without a cover sheet is an immediate flag because every document must have a cover sheet. If removing a

cover sheet is easy I can just as easily remove a red cover sheet and put a white cover sheet on top. Although that is correct, that is still a little harder.

There is no perfect solution—the trick is to keep modifying the solution to the point where is comes prohibitively harder for someone to bypass it. At one Fortune 500 company, they used the color-coded sheets, but each cover sheet had to have the document's name printed on it and a manager's signature verifying that the cover sheet and document named matched. Once again, people can always be tricked but this made it a lot harder for someone to bypass. Now you would have to involve other people to bypass the system, and even if you think your boss may not actually check, it is still a risky thing to do.

Fear of being caught is a big driving factor, an advantage of using this method. Most people will not chance having the boss sign a form in which the cover sheet does not match the title. With proper training, bosses should be made aware that checking documents on a regular basis is vital, plus periodically telling employees they are going to review documents to make sure they match the covers keeps people on their toes. If you are a boss that is known for carefully reviewing documents, once you catch one or two people and the word spreads, the chances of someone else trying to do this would be very slim.

Locks, though fairly basic, can be very effective at controlling and preventing access. If someone cannot get into a room or cabinet it becomes harder for them to cause harm. If you cannot get access to something you cannot disclose it, you cannot modify it, and you cannot destroy it. The effectiveness of this control depends on the type of lock you are using and how the key to the lock is controlled. The lock is not what protects the data, it is the confidentiality and control of the key that keeps the data secure. Key accountability and key control are the driving success factors behind the effective use of locks. Therefore the more carefully you can control the key and track accountability, the higher the effectiveness of the lock.

Standard locks such as key-based locks controls access, but there is not a high grade of accountability of the key. People can give the key to someone else, they can make copies of the keys, and they can loose the keys. If 20 people all have a key to a given room and something is missing from the room, it will be very difficult to tell which of the 20 people did it. The other problem with standard locks that is often misused is the master key. A master key can get into any office. Even if I am the only one with a key to my office, but three people have a master key, we still have not achieved accountability.

The next general category of locks is something that can be tied to an individual, such as an access card. Each person has a different access card and is given permission to the areas needed to access in order to work. Even though 20 people have access to a room, since they each have a different access card, we can still have

accountability and know at any instance which specific individual gained access. The problem with access cards is people can still lend them to others and lose them.

Biometric-based locks are the final general category of locks. Here, the key of the lock is tied to a personal attribute such as a hand scan, a retina scan, or voice recognition, just to name a few. The advantage of this method is that it is tied to the person and cannot be borrowed or lost. It is always with you and depending on the method is usually good for 15 or more years. The disadvantage is that these devices can be fairly expensive and some people are still distrustful of giving away their personal attributes so the company can always track them.

Guards also serve as an effective measure for preventing insider threat. Guards can provide effective choke points and have detailed analytical capability if they are trained correctly. Guards can be trained to be located at all critical perimeters and perform spot checks, looking for suspicious activity and stopping individuals. They can also ask people what they are leaving with and why. Some people do not see the value in questioning because they say that it will catch only honest people because attackers will lie, but combined with other techniques, it can show intent. If I ask you as you get off the elevator if you have any laptops or data storage devices that you are leaving with and you say no, then at the exit another guard checks your bag and finds a laptop, now we know that you have intentionally lied and tried to bypass the system. Without the first step of asking you a question, during the spot check you could have claimed that you did not know you were not supposed to have a laptop. However, now because we used a two-step process, you were caught red-handed.

In addition to being placed at perimeters, guards also have value walking around the facility looking for suspicious activity and performing random spot checks. If I know that no one else is around, I can make copies and fax sensitive data without anyone knowing. However, if I now know that a guard could walk by at any moment looking for unusual activity, and the activity I am currently performing is highly suspicious, I might be less tempted to continue my activities. In addition to looking for suspicious activity the guards can also do what is termed a traditional spot check just to make sure people are following procedures and doing what they are supposed to be doing. If all sensitive documents are supposed to be locked up if they are not in your possession and a guard finds a sensitive document on a desk with no one around, this represents a security violation. Just because a person left a document on a desk does not mean they are out to cause harm. However, the fact that anyone else could walk by and see the document represents an avenue where an untrusted entity could take the document or read it and gain access to information they should not have access to.

If a document is very sensitive, it should be locked in a room and not allowed to be removed. Now people can come in read the document and leave but the docu-

ment can never leave with them. If the room containing the document does not contain any phone lines, copy machines, or other electronic equipment the chances of someone being able to extract the document is very slim. In addition, if the room requires two different people to open and you have to sign in and out and you are being monitored as you are examining the document, this also adds additional measures of protection.

We alluded to it when we talked about locked rooms, but separation of duties is another effective measure to defend against traditional insider threat. If only a single person is responsible for performing a task it is easy for that person to abuse privileges and cover up any malicious activity he or she is doing. However, if two people are involved with any task it would be much harder for any one individual to cause harm on his or her own.

Although these security methods seem basic and straightforward, they work and they solved the problem. Before computers, insider threat was still a problem, but companies knew how to control it and if they were concerned enough about it they could take action. It was not a problem that was out of control, where the attacker clearly had the upper hand.

Now that you have added computers, the Internet, and the mobile work force, just to name a few, we have a whole different ball game. Now, the potential for loss is huge and the methods they can utilize are numerous.

However, do not discount these measures. Even though they are not the main focus of this chapter, they still have value and use today, and sometimes the simple solutions work better than the more complex solutions.

Overview of Technologies

Technologies typically are going to serve as the basis for insider threat attacks. Even though simple, nontechnical attacks still work and are effective, they are not the focus of this chapter. They were covered in the introduction to lay the foundation. The main focus is to give an overview of key technologies and how they can be used to cause harm and damage across your organization.

It is important to remember that this is not a comprehensive list, and technologies are always changing. However, there are certain technologies that must be understood so you can try to stay one step ahead of the malicious insider.

The key technology areas that will be addressed are:

- Information extraction
- Network leakage

- Encryption
- Steganography
- Malicious attacks
- Beyond computers
- Humans

As you read through each section and start to understand each technology, how it works, and how it would be used, remember to think about the entire scenario in which this could be used to cause harm by a trusted insider. Understanding the core technologies is critical, but thinking about the big picture and how it would be used in an attack scenario will get you the best value in understanding how the malicious insider operates and how to defend against them.

Information Extraction

As we covered in the introduction, insider threat was occurring long before computers and long before the Internet was created. As long as organizations had sensitive data that they did not want anyone to obtain, insider threat existed. Even Julius Caesar had to worry about insiders and how they could hurt his organization. Probably the most famous example of insider threat is that of Jesus and Judas. Jesus' trusted disciples knew sensitive and critical data that no one else knew, and though Judas was a trusted insider, he was paid off to cause harm to Jesus.

In most cases computers do not invent new areas of crime, they just bring it to a whole new level. Computers and the Internet and companies having LANs did not create insider threat, it always existed. This new technological infrastructure just made the area of insider threat a more fertile ground for attackers because now their options and means and methods for causing harm are almost endless.

Before computers the main method of committing insider threat was information extraction; taking a sensitive document, taking it outside of the company, and giving it to someone else. Since technology did not create the problem of insider threat as we apply technology to the problem, the problem does not change. The main issue with insider threat in the current digital age is extraction of information, now using networks as opposed to putting a hard copy in a briefcase and walking out.

Therefore we will start looking at technologies in the area of information extraction, and then at additional technologies. The core technologies in the insider threat tool chest that we will examine are the following:

- Hidden files
- Removable media
- Wireless exfiltration
- Laptops
- PDAs/Blackberrys

NOTE

For most of the examples we will be using Windows. Everything that can be done with Windows can be done with UNIX. We are doing this for many reasons: First, many attackers are not that technical and therefore most likely will be using Windows as their primary OS. Second, most insiders are going to be using their desktop to commit insider threat and once again, Windows is the primary OS for desktops.

Hidden Files

Data exists in a digital form on your computer, on a server and in packets that are flying across your network. A simple but effective way to extract data out of your organization is to hide those files on your computer so if someone does a basic directory search they will not be able to see or find the documents on your system. In this section we are not talking about data hiding or steganography; that is such a broad topic in and of itself that it deserves its own section later. We are talking about basic file hiding on your system.

Similar Directory

The first method of hiding files on a system is to find a directory that contains a lot of files with similar names and create a new file or directory with a similar sounding name. The chances that someone knows or will check every single directory is slim. From a directory standpoint, c:\windows is a great spot for creating new directories that will never be spotted. Figure 2.2 is a screen shot from a portion of the c:\windows directory.

Figure 2.2 A Portion of the C:\Windows Directory

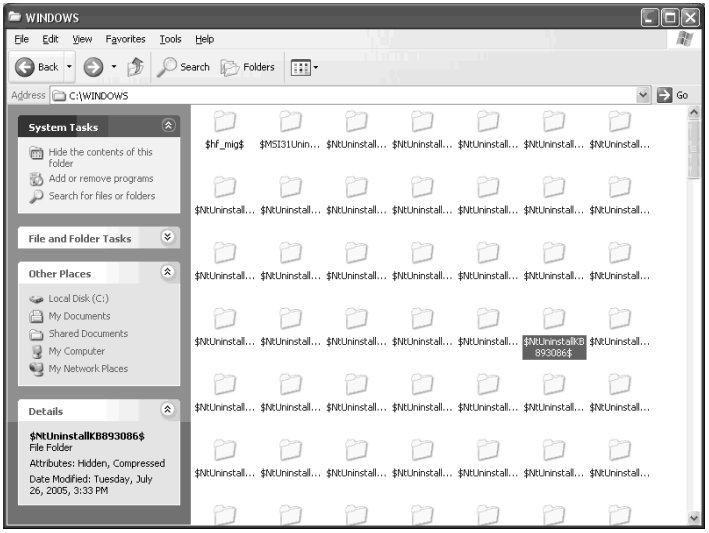

Can you tell me which directory is supposed to belong and which one is not supposed to? If I go in and create a new directory that starts with $NTUninstall, it will blend in and be virtually invisible. Adding to the fact that most people are very leery of deleting directories in the Windows main folder, and that there is minimal information on the correct folders that should be there, means most people when examining their system will ignore it. Add that with the fact that the folders are constantly changing based on the running of the system, no one is going to take the chance of trying to delete files that could cause their system to stop operating.

Similar File

From a file hiding perspective a great area to hide files is in c:\windows\system32, especially if you pick a file name beginning with ms (see Figure 2.3).

Figure 2.3 The C:\Windows\System32 Directory

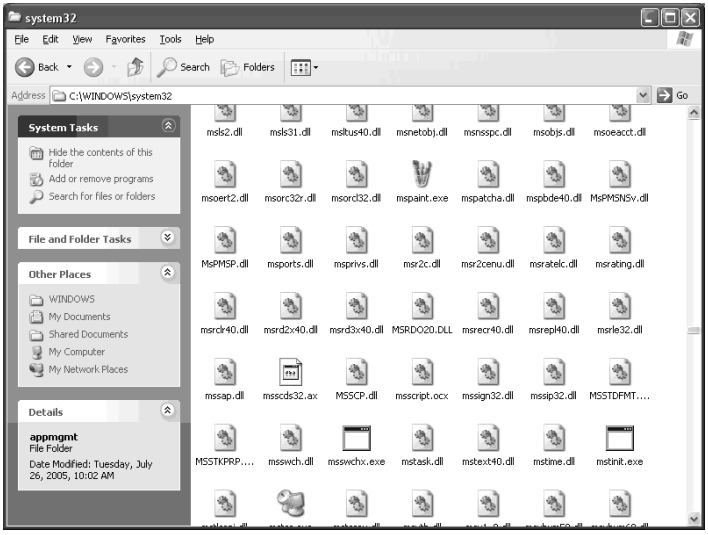

In this directory you have many files beginning with ms and of a variety of file types, once again making the attacker's job very easy to perform.

File Extension

Another variant to hidden files that works very well is to change the file extension. If I have an .exe and I change it to a file type of .dll or .doc, it will make it harder for someone to find. If attackers think I have a malicious .exe on my system, they are going to go into my system and look for it; if they do not find it, some people would stop looking. Now when you want to run the program you would have to rename the file type back to an .exe and run the program. Anyone who has worked at a company and has had to use e-mail is probably very familiar with this tactic. This is probably one of the most common tactics to get past e-mail filters. If your e-mail filter blocks Word documents or .exe files from coming into the mail server because they can contain viruses, you would just change the file extension on the file, attach it to your e-mail, and have the person on the other side just rename it back. To do this you would either right-click on the file and rename it or use the rename command from a DOS prompt to change the file type.

Hidden Attribute

Getting slightly more advanced than what we have been doing, you can use the hidden file attribute to hide the files. Files have attributes that can be set; one is the

ability to hide the file. By right-clicking on a file and clicking **Properties**, you can see the two main properties for a Windows file: Read-only and Hidden (see Figure 2.4).

Figure 2.4 Properties for a Windows File

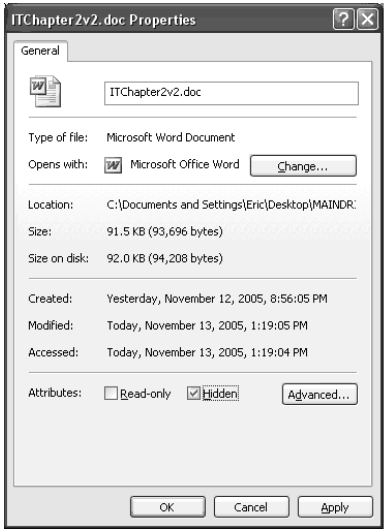

With the hidden attribute selected, you will not be able to see the file. You can also use the **attrib** command from a cmd prompt to change the attribute of files (see Figure 2.5).

Figure 2.5 Changing the Attribute of Files

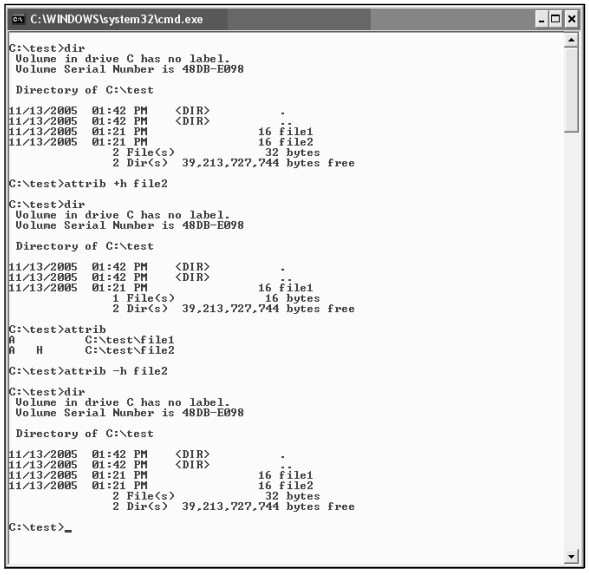

In this example we started off by showing a directory with two files in it. We used the **attrib** command to add the hidden attribute (+h) to the file. Then we did another directory listing and you can see that the file is no longer displayed. However, if you use the **attrib** command with no arguments you can quickly see that the file is still listed, with the H or hidden flag. Finally if you use the **attrib** command with the –h, you can remove the hidden attribute and the file will once again appear.

It is very easy for someone to hide files on a system, and if you are not aware of it, it is easy for you to miss them. I recommend always configuring Windows Explorer to always display all files, including hidden files. This way there are no surprises and you see everything that is on the system. If you are ever performing any analysis across a system it is recommended to always configure Windows Explorer to show all files.

From Windows Explorer, select the **Tools** menu and click **Folder Options**, then select the **View** tab. The penultimate item under hidden files and folders is set by default to Do not show hidden files and folder. My recommendation is to always change the setting to Show hidden files and folders as it is set in Figure 2.6.

Figure 2.6 Windows Folder Options

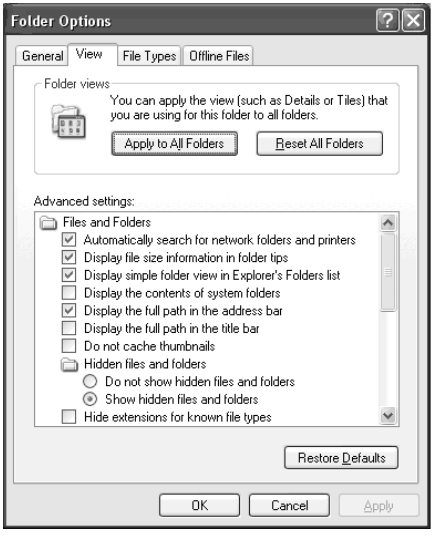

Even though attackers can still find ways around this technique, it still provides a basic level of protection.

Alternative Data Streams

The next technique we are going to look at, which is fairly advanced and very powerful, is alternative data streams. This is a feature that is available with NTFS file partitions on Windows. Since it is recommended that you always use NTFS as the file system, this technique will work on most systems. By right-clicking on a disk and selecting **Properties** you can quickly see what the file partition is (see Figure 2.7).

Figure 2.7 Local Disk Properties

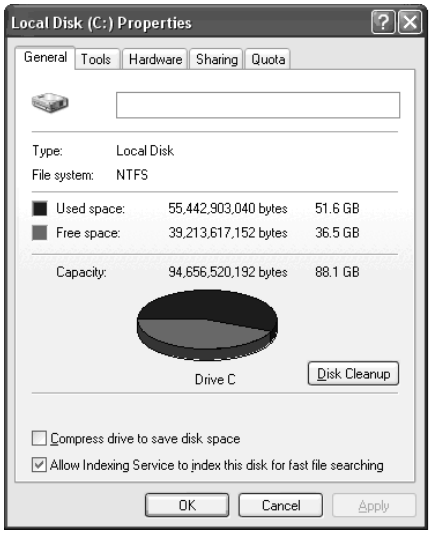

In Figure 2.7, you can look under file system and tell that it is NTFS, the default on most systems.

With NTFS there is a feature called Alternative Data Streams that turns every file into a file cabinet, and each file cabinet can contain many files. When I open the draw of my file cabinet at home I have my files, one for each of my clients, one for my bills, and so on. Here you can do the same thing. The file cabinet is the name of the file, and then I can attach additional files to the main file and each of the additional files are hidden and cannot be seen. When a secondary file is hidden within a main file using alternative data streams, the secondary file cannot be executed, but whenever the main file is copied or moved, the secondary file automatically moves along with it. Then if you want to use the secondary file, you have to unattach it from the main file and use it.

This is a technique used by inside attackers to hide sensitive data on a system and unless you know it is there it is almost impossible to find. Now I can take my

sensitive files, hide them in an alternative data stream, and either leave it on my system or copy the file somewhere else. Anyone looking at the file would have no idea that there are other secondary files, filed away in the main file.

There are two ways to create alternative data streams:

- **Attaching to a file**. Use the cp program from the NT resource kit. The format for issuing the command is:

```
cp hiddenstuff.exe boringfile.exe:stream1.exe
```

- **Attaching to a directory**. Use Notepad to open a file stream connected to a directory using the following command:

```
c:\ notepad <directory_name>:<stream_name>
```

Attaching to a File

To create an alternative data stream you have to use the cp command from the resource kit. Figure 2.8 shows the sequence for creating an alternative data stream.

Figure 2.8 Creating an Alternative Data Stream

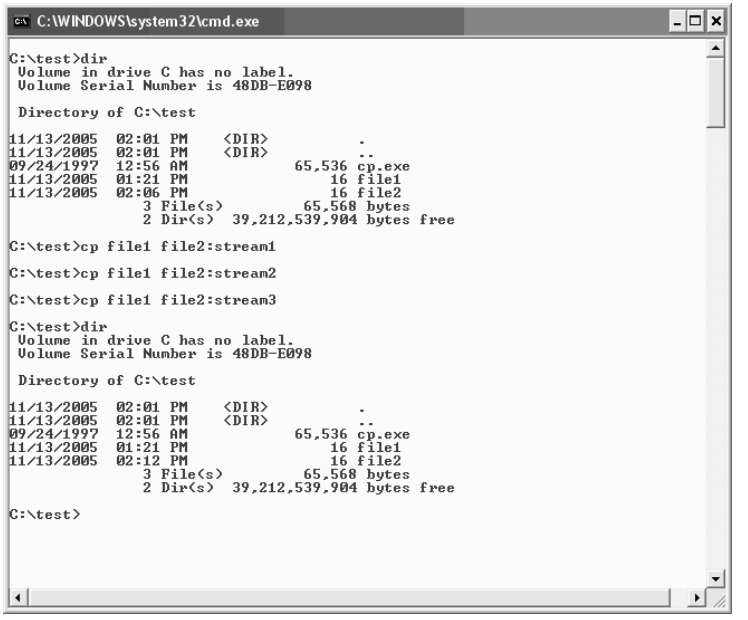

Typically this method is used if the file you want to hide already exists so you can attach it to a primary file as an alternative data stream. Notice there is nothing

unusual about the file that has changed. It still looks the same in terms of size. Even if you check Properties from within Windows Explorer, there is nothing unusual that stands out (see Figure 2.9).

Figure 2.9 File2 Properties

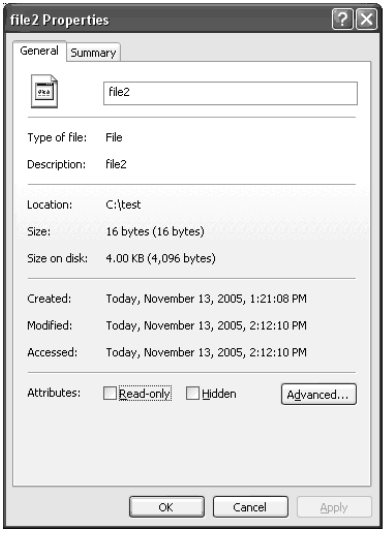

Alternative data streams are fairly nasty, since if you know the file is there you can extract it, but if you do not know it is there is could be very difficult to find.

Attaching to a Directory

Attaching a stream to a directory works best if you are creating a new file from scratch; however, it can also be used if the file already exists. If you start Notepad with a filename of a directory and a filename, it will create an alternative data stream attached to the directory. If it is a new stream, Windows will prompt you on whether you want to create a new filename (see Figure 2.10).

Figure 2.10 The Notepad Prompt

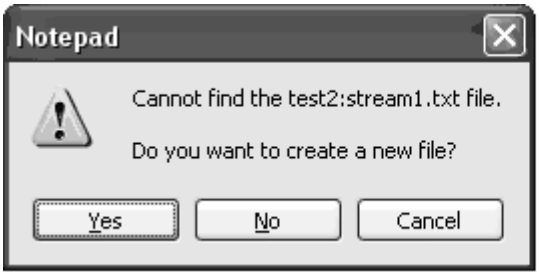

Select **Yes** and then type or paste the contents of the information you want in the stream. When you are done close the file and the stream will have been created. This can be seen in Figure 2.11.

Figure 2.11 Creating an Alternative Data Stream

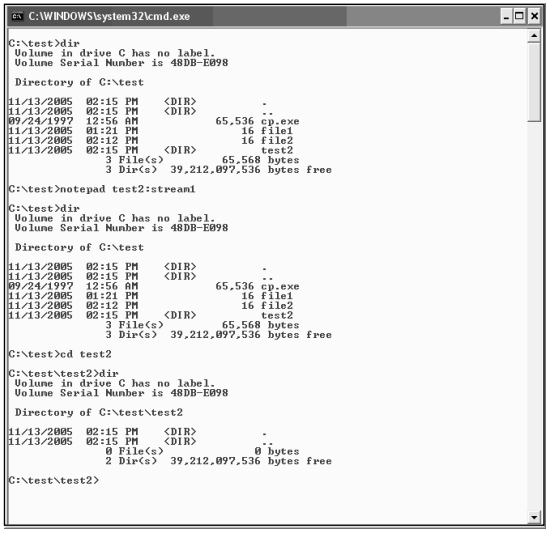

We open up Notepad and create a stream attached to directory test2. After we are done you can see that there are no visible signs of the stream in either directory.

Removable Media

Although alternative data streams and hidden files are interesting concepts, we still have to be able to walk out the front door with our information. Hidden files will make it harder for someone to find, but we still have to put the information on something so we can leave the company with it.

One of the best ways to remove information from an organization in a covert manner is with removable media. Technology is an amazing thing, especially in this regard. The longer you can wait the more storage you can get in a smaller form factor. Memory sticks that used to contain 128 MB of memory now, for the same size, can contain over 4 GB of information. In many cases you can find something half the size that contains twice as much data storage capability. The real question is how much do you want to spend; today you are not talking that much money.

As I write this chapter, I am sitting on a train with a USB drive the size of the key to my car that contains 4 GB of data. While this is very powerful, because as I travel I can back up my data on the fly, it also raises interesting questions with regards to secu-

rity. What if I lose the USB drive and someone else gets hold of it? During the train ride I can back up all my critical data in a few minutes on a device that easily can be concealed. What if someone else can get three minutes with my laptop when I am at a meeting or presenting and I step out to take a break? One of my themes that really drive home this point is that anything that can be used for good, can be used for evil. Organizations and IT people love these new portable devices, but we have to step back and think about the implications this technology poses.

Based on the ease with which someone can extract data out of the organization, the real question is whether allowing USB drives on work computers is worth the risk. People can back up their work, but they can also use it to cause significant harm to the company. Since most organizations have file servers, does it make sense to force all users to back up and store sensitive files on the file server and not run the risk of having USB ports accessible? You might bring up the point that for traveling users they would need this capability, but laptops present their own problems (covered in the following section).

So far with removable media we covered only small USB drives. There are many other storage types, however, that plug into serial and parallel connections and even use local wireless. Once again the decision has to be made on which is the lesser of the evils, allowing the access or not allowing the access. Since there are alternative ways for users to accomplish the same task in a more secure manner, my recommendation is to lock them down.

Removable media also comes in different form factors. There are watches and jewelry that contain USB plugs, and some of it is quite good. You could detect the earlier versions because they were unusual—geek-wear that most normal, self-respecting people would not wear. However, the new USB watches that contain over 8 GB of storage are no bigger than and look similar to the Tiger Woods self-charging Tag watch. There are also rings, earrings, and necklaces with storage capability very well concealed. Now if you see any people at your organization wearing an extra amount of bling, you might want to keep a closer eye on them.

Probably the biggest culprit of them all I saved for last, because it always creates quite a stir when I bring it up at conferences: iPods. Yes, I am actually saying that organizations should have policies about bringing iPods to work, and should restrict and control the usage. I have no problem with people listening to music at work, but I do have a problem with people bringing a 4+GB hard drive to work in which they can copy and store anything they want conveniently disguised as a music device. I am not against iPods (I actually own several); it is an awesome invention, but you just have to realize the risk it presents.

Put aside all your personal preferences and biases on the matter, and ask yourself if your organization should take the risk of allowing people to bring in personal

devices that can store large amounts of data, remembering that it is their personal device and you have no control over it. In essence if you allow them to come into work with it and you have no policy stating they cannot, then since it is their personal property you have to allow them to leave with it.

Laptops

Laptops are probably one of the greatest, most powerful weapons for insider threat available to the attacker today. At this point I know some people must think I have lost my mind. First I insult iPods and now I am going after laptops. You are probably thinking that in the next section I am going to say Blackberrys are evil too (well, just wait until you read the next section). The point is control of IP is what companies should be focusing their energy and attention on, and any device that can provide a convenient avenue for circumventing those controls should be examined very carefully. I agree that laptops play a critical role in the IT arsenal of any organization, but we have to be aware of the threats that they pose and put measures in place to carefully control them.

Just to prove my point on how dangerous a laptop can be if it is not carefully controlled, I want you to do this simple exercise. Randomly walk around and borrow someone's laptop at your organization. Before you look at it ask yourself how much the laptop is worth. You are probably thinking $1K to $2K, focusing on the price it would cost to replace the physical equipment. Then I want you to do a quick analysis of the laptop looking at what files, data, e-mail, and such is on the laptop, and what damage it would cause if that information was exposed to the media or a competitor. Do not spend a lot of time analyzing it; 45 minutes max will be able to give you a good idea of what is on the system. *Now* ask yourself how much the laptop is worth (really thinking how much monetary loss and damage the company would suffer if this laptop was put into the wrong hands).

Many executives or managers before they travel dump all the critical data from the file server onto their laptop so they have all the information they need while they are on the road. In many cases the laptop is no longer used just for traveling, it is also their desktop system. The term portable desktop has emerged, showing that laptops are no longer computers that people use just when they travel; they are now systems that are used all the time, and for smaller companies that do not have file servers, the laptop is the file server and contains *all* the critical IP belonging to a company.

Laptops are a necessary evil, and we have to make sure that we understand the inherent risks and put measures in place like encryption, and minimize wireless to make sure the information stays protected as possible. Purging is also a good practice, where any unnecessary information or data that is not needed is removed from the

system. Now if for some reason the system is compromised, the amount of data loss is minimal.

We have to make sure we avoid being digital pack rats. We would jokingly tease my grandparents that everything they ever bought they still had, because they would never throw anything out. Their garage and closets are filled with things that they can't even identify, and could not find even if they wanted to. We say we do not understand why people do that, but then I look at my office and realize that anything I ever created digitally I still have. Half the stuff I probably could not find if I wanted to, but I have it somewhere. Especially because laptops are so easy for someone to take, we have to make sure we do not become digital pack rats. Store the minimal amount of information needed and minimize your overall risk with what you are carrying around with you.

PDAs/Blackberrys

I definitely could not live without my Blackberry, and I am not sure if it is a quality-of-life addiction, but nonetheless, a Blackberry and PDA enhance our lives by bringing information to our fingertips anywhere we might be. You can be in a meeting and not be able to take a cell phone call but can quickly look down at a PDA or Blackberry and pick up emergency messages or even send quick replies to messages that might be very urgent. Why make the person wait six hours to get an answer to a question because you are tied up in a meeting, when you can just shoot them a quick reply back on your mobile device in less than 30 seconds?

With PDAs, however, take all the problems and issues associated with a laptop and multiply them. Essentially, PDAs have the same information that a laptop might contain, except with even less protection and in a much smaller form factor. At least laptops have a password that is required to log in before you can access the data, and even though it can be bypassed it still provides some level of protection. With most PDAs, though, there is no security, or if there is, most people turn it off. If I lose my Blackberry, someone can get access to e-mails, files, contacts, and credit card numbers, social security numbers, and other very sensitive information that many people store (not that I would) on their PDAs.

I had a client that started having some strange problems where their servers were being accessed and data was being copied, but there was no sign of attack, no visible break-in, no back doors or Trojans that you normally would see with a typical account. To make a long story short, one of the admins lost his Blackberry, purchased a new one, and never made a big deal about it. However, within all his e-mails and files were all the IP addresses, passwords, configuration information, and so on, and evidently someone got hold of it and started using that information to log in

remotely. Therefore we quickly had to go in and change all the passwords on critical systems and control external access into the network. Even though the attacker came from the outside, I would clearly list this as an insider threat because a trusted insider with special access, in essence, took all the passwords and gave them to an untrusted entity so they could come in and do whatever they wanted to do. It wasn't on purpose, but unfortunately in this game that does not matter. Whether it was done intentionally or accidentally, the fact remains that an insider created an open door for an outsider to gain access.

The really scary part about this case was when I questioned why they did not have procedures and formal notification when critical equipment with sensitive data is lost, so immediate action can be taken. The response was that it happens all the time, and it would be too much work to change passwords every time it happened. If that is the case, you have a few options:

- Limit the information that is on the PDAs.
- Limit who actually has PDAs.
- Put together better awareness methods to minimize the chance of equipment getting lost.
- Do not use PDAs, because if you are not willing to accept the responsibility that goes along with them, then the organization should not have to accept the risk.

The other scary thing about PDAs is that people leave them lying around. During authorized insider threat penetration tests that I performed, people left PDAs on their desk when they went to meetings or lunch. Since it was authorized and I had permission, it was easy hypothetically to take the PDA, download the data to my computer (which takes about five minutes), and then put the PDA back. In most cases the person does not even know it is gone. In one case, when I went to return it, the person was sitting there, so I had to think quickly. I said I found the PDA in the break room, and the person admitted to often forgetting where he put it, and said he must have left it behind. As I stated earlier, the information that is found on PDAs is scary and organizations must do a much better job in protecting and controlling the information.

Wireless Exfiltration

The last technology that we are going to cover under information extraction that requires no introduction at all is wireless. Wireless is probably one of the easiest ways for an outsider to become an insider. There are many maps showing all the wireless

access points in a given city; Figure 2.12 shows a map of commercial wireless antennas in the United States (from http://www.cybergeography.org/atlas/wireless.html).

Figure 2.12 Commercial Wireless Antennas in the United States

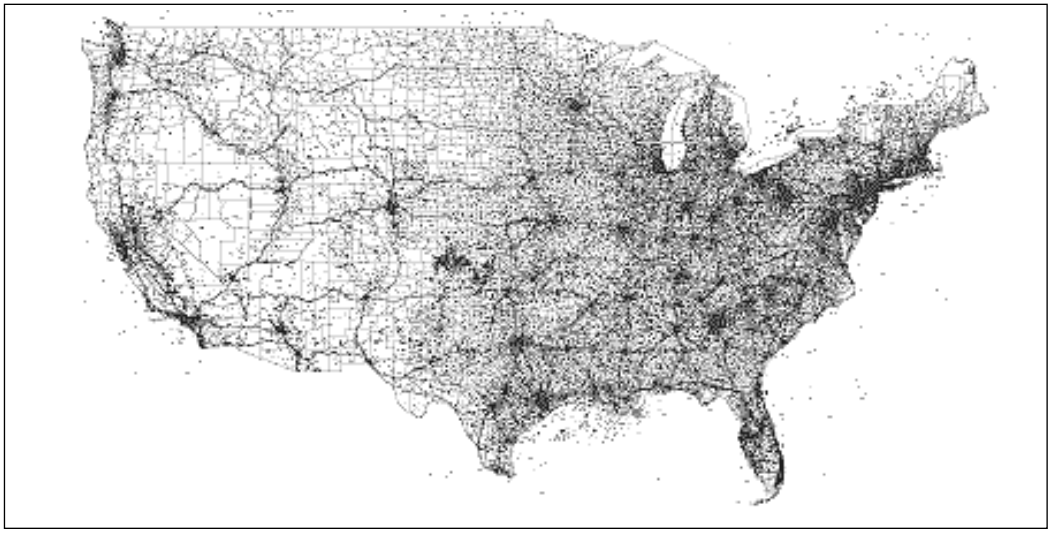

If you do not think wireless is a problem you might want to look closely at the map in Figure 2.12 and at other war driving maps that have been produced for almost every city (available on the Internet). War driving is the technique of driving around an area with a wireless antenna and GPS locator and a preloaded map, and plotting out the location for all wireless access points, both protected and unprotected. If you think you can put up wireless and not be found you are wrong. Not only will other people be able to find you relatively quickly, but with tools like Network Stumbler anyone can quickly find wireless access points in a given area (see Figure 2.13).

Figure 2.13 Using Network Stumbler

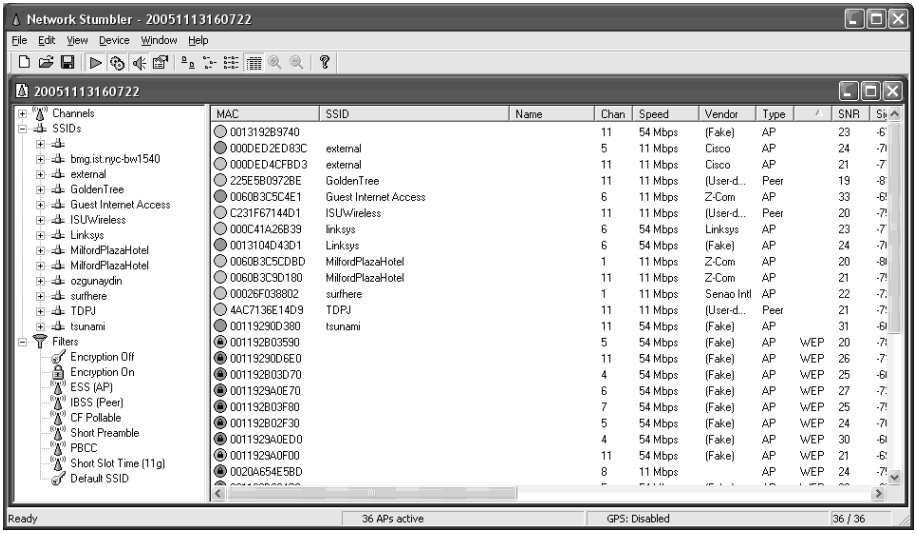

As I was writing this section I just opened up Network Stumbler to see what I would find and a whole list of APs popped up including at least a dozen that are not protected. Even the ones that are protected are not necessarily protected with strong security. Encryption on just means some from of encryption is being used, but many of them have weaknesses and ways to be broken if they are not configured correctly.

There are three general categories of wireless that we are concerned with:

- Authorized wireless
- Rogue wireless
- Ad-hoc wireless

Authorized Wireless

Even authorized wireless connections can create problems for an organization. They can be misconfigured and cause harm just like authorized users can make mistakes and cause harm. The first problem with authorized wireless is not properly restricting who can gain access to the device. There is still a large percent of authorized wireless access points that anyone could connect to.

I still remember talking with a friend, who is a financial advisor, about wireless. He said that his very large company does not allow wireless use, so some people are setting up rogue access points. He said that he usually just connects to Company X, which is a large government contractor, and uses their wireless because it is not properly secured

and seems to work fine for him. I quickly told him never to send me an e-mail again, because I do not want all my sensitive data to go through that other company's network. He sheepishly said that he had a feeling that was a bad idea.

The second problem with authorized wireless is that they are connected directly to the private network. Even if proper security is put in place, wireless should be used as a means of access and still terminate at a perimeter, preferably the firewall. Then if the firewall determines the access is OK, then and only then should it be allowed onto the network. A wireless access point is not a firewall and it is not an IDS, therefore do not let it bypass your firewall and IDS. Any connection needs to be validated by your security devices, no exceptions.

Even if proper authentication is used, your data still might be going across the air, unencrypted. Even though you have slightly reduced the problem it is not much better. Now someone sitting in the parking lot or building across the street cannot directly connect to your network and access your data, but they can sniff all the traffic and if any sensitive data is being sent over the wireless network, they can still intercept it and analyze it. There is a lot of confusion on what companies actually are getting and what a given protocol will provide, but the best way to know for sure is to test it. When was the last time you set up a sniffer and examined your traffic to make sure it is really encrypted? People tell me all the time, we are fine because the product provides encryption. Great, but are you using the encryption and did you set it up correctly? The only way you will know for sure is if you test it.

One of the big reasons why authorized wireless is still causing problems is because many companies are being forced into wireless sooner than they would prefer and therefore are not taking the time to set it up correctly. With many organizations, rogue wireless has gotten so out of hand, that many companies feel the only way to stop people from setting up rogue points is if they provide wireless to their employees. However, this solves the rogue problem only if it is set up correctly; otherwise you are spending a lot of money to have the same problems you did before you set up your own sanctioned company wireless network.

Rogue Wireless

Wireless is truly a plug-and-play commodity. From your favorite electronic store you can buy a wireless access point for around $50; an insider goes into work, unplugs the cat5 cable from his or her computer, and plugs it into the wireless access point. Then they take a second cat5 cable, plug one end into the wireless access point, one end into the computer, and you now have a rogue access point at your facility. Five minutes and $50 is all it takes. Many users know the dangers of rogue wireless and set it up anyway because the convenience is more important. Others do not under-

stand the dangers and think that if you buy something it must be secure, because stores would not be selling and setting up a rogue AP unintentionally.

Either way a rogue AP presents an easy way for an attacker to gain access to a network. The problem is, your organization needs to be prepared to detect and defend against rogue APs. Since they have become such a problem, there are actually commercially available wireless solutions that will access, track, identify, and disable rogue APs. These solutions work very well—the only thing you have to be careful about is that you do not take what you believe to be a rogue AP offline, when in reality, it is not on your network and does not belong to you, it belongs to another company that has office space in your building. As with any solution, measure twice before you cut.

Ad Hoc Wireless

By default with most configurations of Windows Operating Systems, when you turn on your wireless card, ad-hoc wireless also is turned on by default. With wireless ad-hoc or host-to-host wireless your computer is advertising itself as an open connection that someone can connect to. This vulnerability is shown in Figure 2.14 with the two bottom wireless connections.

Figure 2.14 Vulnerabilities in a Wireless Network

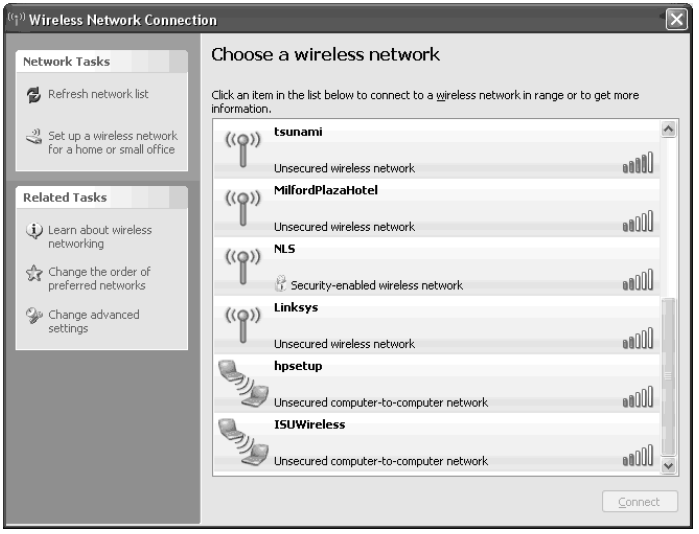

Now if someone connects to these connections they most likely will not have Internet access but they potentially could access any files or programs that are on that computer. Based on our earlier discussion of laptops, you can now see how bad

the problem can get. Now you have an executive who has all the critical IP of the company on his laptop. He leaves his laptop and wireless on and goes to dinner, and as he is sipping a martini, someone is accessing and copying all the sensitive IP off his system. When he comes back from dinner, he has no idea what just happened and how bad it really is.

Regardless of which method of wireless is used, this provides an easy gateway for someone to get access to critical resources, and unfortunately in most cases, bypassing the perimeter. The typical rule of strong perimeter design is all connections must go through the firewall. However, for some reason, with wireless, many organizations decide to put it behind the firewall so if it is not configured correctly (which in most cases it is not), it allows an untrusted outsider to become an insider, putting your entire network at risk.

Wireless is so easy to perform today that people tend to forget what a huge risk it could present. Therefore it is critical that measures be put in place to control and secure wireless, otherwise the base of trusted users on your network could be a lot larger. In one case a security administrator said I was given an unfair advantage based on how the wireless was configured. I have 8,000 users, which I thought was what I was up against in combating the insider threat problem, but to do misconfigured wireless, the number of trusted users I had to deal with was more than 20 million. Although wireless has value, make sure the reward is worth the risk you are taking.

Network Leakage

Acquiring access to information does not accomplish anything if you cannot extract it from the organization. In prior sections we talked about walking out the front door with the IP either in paper form or digital form. However, that can still be cumbersome and risky. Ideally from an insider's desk they would like to be able to extract data from the organization via the network. The network provides an avenue and connection to the rest of the world. From your desktop computer you have full access to the entire world via the network connection that is connected to the Internet: full access to the Internet, anytime that you need it. Therefore it should come as no surprise that network leakage is a major problem for organizations. If people have the access not only are they going to use it, they are also going to abuse it.

We have talked a great deal about controlling access and limiting access, and one of the key principles that we have talked about is principle of least privilege—giving someone the least amount of access they need to do their jobs. Although we understand and adhere to this principle in some parts of our company, why is it that the one area that has the most power and the most potential for abuse is the one area in which we ignore the principle of least privilege: the network, or more specifically, network access?

Being able to access any systems or resources on the Internet from within a trusted private network is clearly a problem. The potential for people to take sensitive documents and tunnel them through the opening in the network to any resource on the Internet is an insider threat's dream come true. Changing this will require a major paradigm shift. My first question in most organizations is to ask why people need access to the Internet to do their job. Give me a reason why every employee needs to access the Web, e-mail, instant messaging, and so on, with little or no protection. Some organizations are slowly understanding that the risk is too great, and are implementing solutions like proxy servers. With a proxy server the employees have to authenticate to a server that will then analyze and service all requests. The employee never directly talks with resources on the Internet; the proxy will perform the actions on his or her behalf. Although this is a good start, it is still not as properly or finely grained controlled as it needs to be.

Most perimeters at organizations are set up to block or limit inbound connections, and perform little filtering outbound. Even organizations that are starting to filter outbound traffic still allow certain traffic out. Two types of traffic that are always allowed out of an organization are Web and e-mail traffic. Those two protocols/applications are the life-blood of any organization, and are needed to flow so employees can do their jobs. Therefore in this section we are going to look at both Web and e-mail in additional detail.

Web Access

The Web is no longer just a source of information; it has turned into a front end for almost any application or protocol that you would want to run. Initially there were separate protocols for performing e-mail and file transfers and sending messages back and forth. Although they still exist, the Web has transformed itself into a multipurpose protocol in which almost anything can be done by utilizing the simple power of a Web browser.

E-mail has always represented a potential risk (it is covered in the next section), but at least there was a central server that could limit and control e-mail leaving the company. Now with the power of the Web, users via a Web browser can send and receive e-mail, bypassing most of the filters that are in place. Even proxies that are starting to filter out Web-based e-mail has had difficulty via the use of encryption. Now Web-based e-mail is setting up SSL connections that look similar to e-commerce sites, which are allowed, but now all the content is encrypted. Not only is it almost impossible to block, but now you cannot even monitor or see what is going on. Figure 2.15 shows one of the more popular Web-based e-mail programs, Hotmail.

Figure 2.15 The Hotmail Login Screen

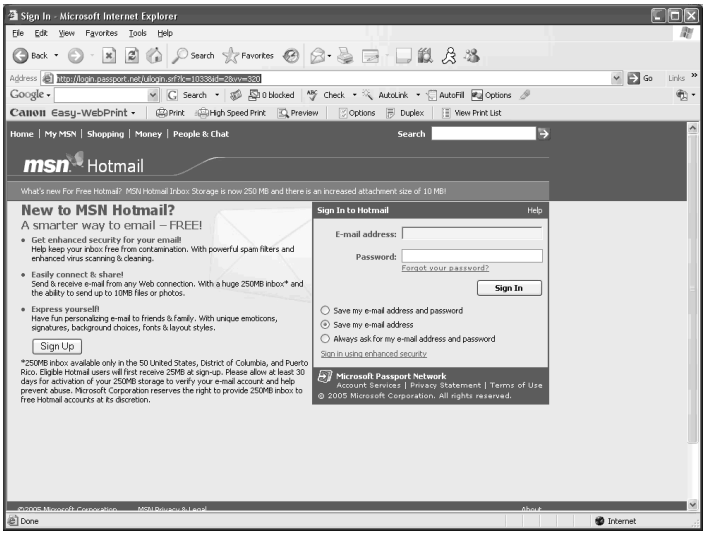

No messy clients, no special protocols—all through the power of the Web you can send and receive e-mail and even attach sensitive documents out of an organization.

Although this is a specialized application, most corporate e-mail servers also have Web front-ends set up so people can check e-mail anywhere, anytime. The problem with these front-end applications is that many people use them so if they do no have their computers with them, they can access e-mail from a Web kiosk at a hotel or airport, to log in, check e-mail, and disconnect. The problem with many of these Web-based applications is that information could be cached or locally stored and if not done correctly, other people could quickly retrieve the information and log into your mail server. Once again, what type of information usually is stored in your mail box? Almost everything, including critical IP for your organization.

If surfing the Web was not bad enough, now there is a whole suite of products called *anonymizers*, with the goal of making your Web surfing experience anonymous to anyone watching. Now if your company is trying to watch or carefully control Web surfing, via an anonymizer they now no longer know where you are going or what you are actually doing. One of the more popular sites for doing this is http://www.anonymizer.com (see Figure 2.16).

Figure 2.16 The Anonymizer Home Page

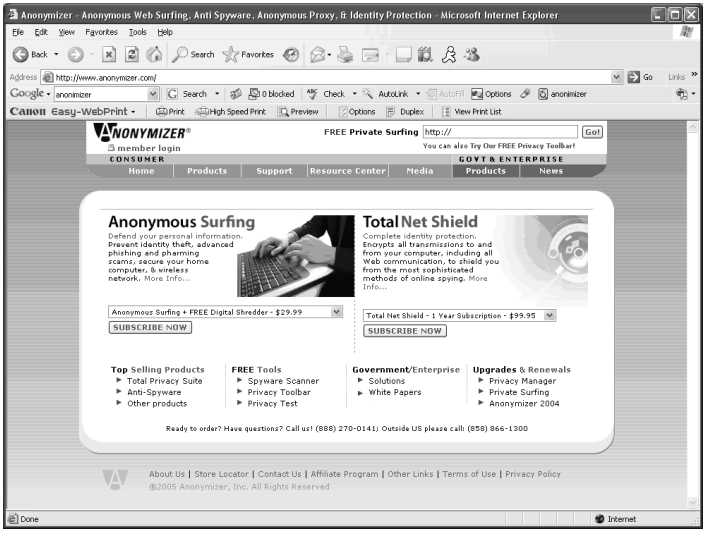

This site currently charges, but there are many sites that are still available for free. More and more corporate sites are blocking or not allowing access to these sites from within the corporation. Although this strategy is effective, so all user surfing can be watched and controlled, the well-versed employee usually will find a way around the controls.

E-mail

E-mail is another application that has huge potential based on the power of being able to send information, including attachments, to anyone in the world with the click of a button. If you are looking for an easy and simple way to perform an insider threat and extract data from the organization, e-mail is your answer. As you will see later, even if it is being monitored, the use of encryption or steganography can make it harder to track and block. Even in the ideal case you could just use a Web-based client that we talked about earlier to bypass the corporate mail server.

Just to show you how simple, yet powerful e-mail can be, let's look at e-mail basics. In order to send e-mail out to another system, your e-mail client has to connect to a mail server. The mail server then connects to the destination's domain mail server and transfers the message. There is a little more that goes on behind the scenes, but the key point to remember is in order to send out e-mail you have to connect to an e-mail server. Most organizations tighten down and properly secure their e-mail server, but there is an obvious solution for the insider that most people miss. Why not run an e-mail server on your local desktop? Now you are controlling

everything. You can configure your local e-mail server to do whatever you want and it will make outbound connections. There are several of these programs available for download. A quick search of Tucows gives a long listing (see Figure 2.17).

Figure 2.17 Searching for E-mail Server Programs on Tucows.com

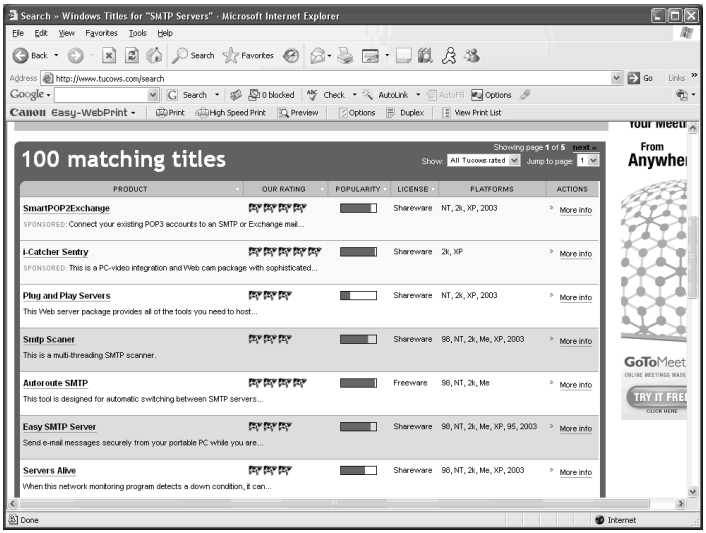

One of the programs that I have found very useful is SL Mail (see Figure 2.18).

Figure 2.18 SL Mail

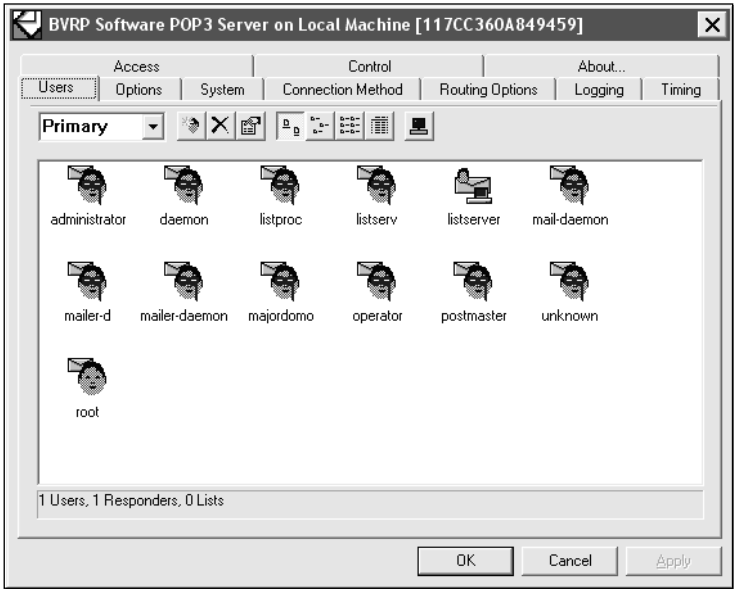

It is highly configurable, very powerful, and runs on a local desktop. Now if you want to send mail out or perform any tricks with e-mail aliases or spoofing addresses, you no longer have to worry about e-mail servers filtering out the message or not allowing them to run. You are running your own e-mail server and can configure it any way that you like (see Figure 2.19).

Figure 2.19 Configuring Your E-mail Server

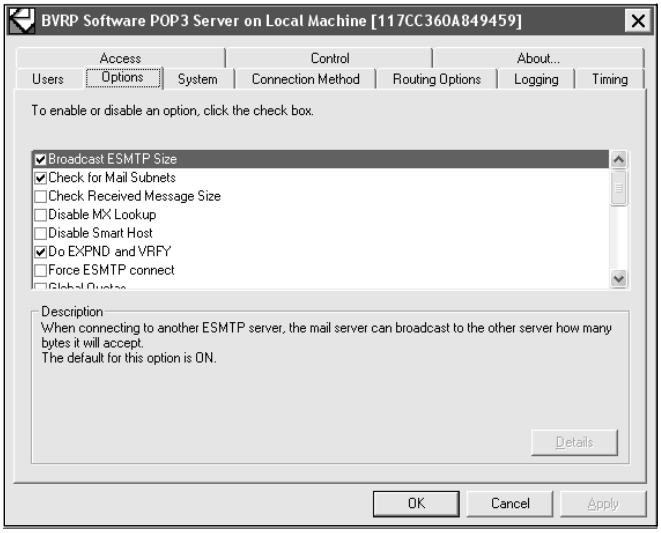

Figure 2.19 shows some of the many options that you can use to configure or set up the system to send e-mail from your local system. These programs have legitimate uses, an attacker can always use them to bypass the controls that are in place.

Spoofing e-mail is also just as easy. You just have to go into your e-mail settings and create a new account (see Figure 2.20).

Figure 2.20 Spoofing E-mail

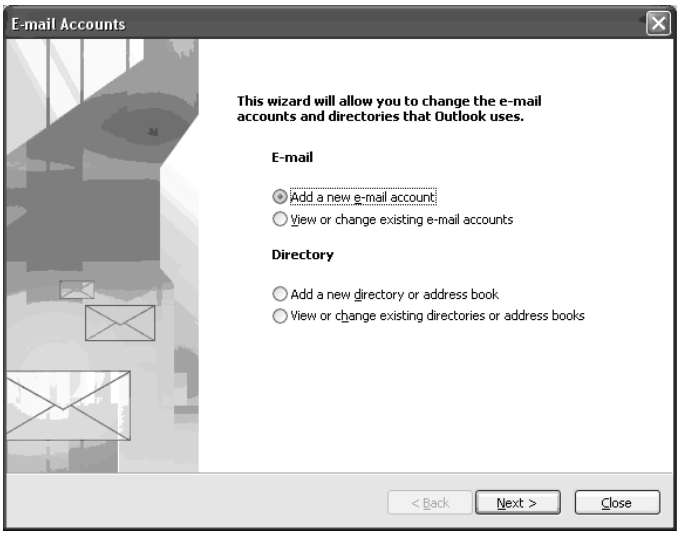

Once you are within the e-mail accounts, just select **Add a new account** and enter whatever information you would like (see Figure 2.21).

Figure 2.21 Adding Information to a New E-mail Account

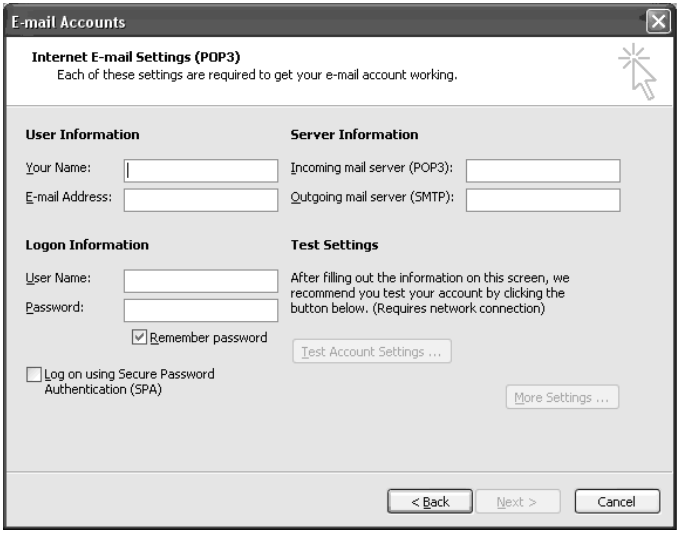

You can create a spoofed e-mail address for a competitor or just play games with your friends. The only trick is getting your e-mail server to send it out, but since you control your e-mail server, that is trivial. It is important to remember that if you

are spoofing a legitimate e-mail address, when the person receives the message and replies back it will go back to the real person's e-mail address. Therefore if they realize they are getting responses to e-mails they never sent out, they might figure out there is a problem. However, if you are an insider attacking the system and trying to trick someone into getting access, it might work. For example at many companies you need to get an e-mail from your boss approving access to a certain resource. In this case it is usually a one-way communication, so if I can send a spoofed approval, it could work to get the additional access.

Cryptography

One of the main measures of protection organizations are using to protect sensitive information from leaving the organization is by monitoring traffic. They either are looking for keywords, actual filenames, or the content of critical IP to track and prevent information from leaving the company. To keep from being caught, one of the technologies that insider attackers are using is encryption. Cryptography garbles a message in such a way that its meaning is concealed.

With cryptography you start off with a plaintext message, which is a message in its original form. You then use an encryption algorithm to garble a message, which creates ciphertext. You would then use a decryption algorithm to take the ciphertext and convert it back to a plaintext message. During the encryption and decryption process, what protects the ciphertext and stops someone from inadvertently decrypting it back to the plaintext message is the key. Therefore the secrecy of the ciphertext is based on the secrecy of the key, not the secrecy of the algorithm.

Therefore to use an encryption program you have to generate a key. The key usually is tied to a username and e-mail address (see Figure 2.22).

Figure 2.22 Using an Encryption Key

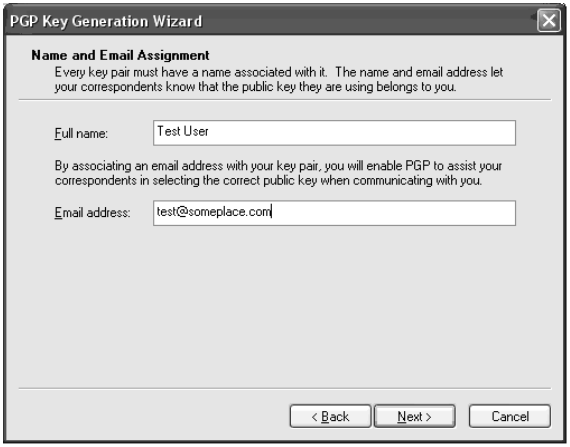

However, no validation is performed so you can put in bogus information that could be used later to launch a man-in-the-middle attack where you can trick someone into using a false key. If you know the public key for a user you can encrypt a message; but only if you know a private key can you decrypt a message. Therefore the public key can be distributed via a trusted channel, but your private key should never be given out. If someone can get access to your private key, then they can decrypt and read all your messages.

Detection

The benefit of encryption is that someone cannot read the content of the original message; however, encryption is detectable.

From a nonmathematical standpoint, if you are expecting to see ASCII characters and English text in a message and you see the following, clearly you can tell there is a problem.

```
dUTiR9wm+A0b3RFqJnpghU0QkyFpXliyt+c/lQkk1M5Q60azvBCi+XRJW/k7p15Y
d8fK/k5pSDiH+YwxDE62j8hs27sT7srttcEj+X4WcBimnUxIh2m22UARtGgXJOmp
zITSRgP9E6gab/iXiygH3IutNG2ovzgILaIeR7YTnYoGLVRgUjtuuRjSgX7aKLtz
7NcNLLtrNmdNZwCX0QlXdVi2NUOJPA96CSFgRfHw26GZ+3Jx4T4F7xRMDx2zIAsRBl
T7zqujZozh1hAqXpiVxk4bkkdeBBUuv8DXayLd96+ADnXlLBwU4DK+4i4J2T
rp8QB/9zVTo0sfigXAR8vI78EQrq/U225kXHLqPwPVgwuBOkZhlaTfEhTRZDmK84
rJZWwGUZAcEwTGPNzm/SYiJf2rS1ziP/hdyblfm3jgUkxi3pAoWAtyukLdwT1O+4
+9I3bnTnxjgE4olo/WRGJx/CV0ou2PQYjMOJsKWBZrR3pL8fhgO3kBT3Sx2OimCG
8HzZV5C0KZVnVyiFQKi1uBETtbUDXHjvdzmqFzfksght5+P+H+98SDu0MyvqrHaJ
hOHMTXwKdohLwTlrpBbe9ZT26oV+8nLlQlVca72mFACWyEfc4AH52BouOyIEQ/g8Yuthl
uxTvHjvY7JdtGJKLhEhYV5B/0ZjCtHa1fSBhY011ZrBZthVr9YjruWvDBo
/UukEcHOGyrw0UXEte3YmQceZOSVQCe1V3k675gvUi/HnAWG8dfneT5v/f1u2Vjp
WWdKMNcloeI4gQE8V91aA9AXPycG6wR/4KsInOB57WD5nsAC/FHE9sShLHsJ/2Cl
QyarhpGikhUNLGKFBNc5RYameGrp37pZn44ev4OgGLsAKzt99lST2cT8cXrglHgE
igbnp6xbHkLF2mkXm0nsA7B1cDm4dZ959IOp8mCFFbcsUzY3zPDMD69kPVPBa6q4
+pK89ioZ8fIqlf5Arip1IqSJhy2DGiZChzjij9ldWL5blppD0q5/0sGFAQNphq2h
hfWZyWRiY6Gz9LmwpyOmdI7YnCS5XQ2lbFpkKtLVy6flhVdyjxAscBkoqMcmUgc0
HKgcQlRXC/syLvpk3tEgiivVYBDq7VWGHYoh4swwqA4VOZDHKjoaO/FaIPayWo0f
RrWosMmKPtySgnD8WDZKyCXEshyx7jApapsxIeoTI5Z7soAHta5Yvl1Li8H8+KKB
```

Although this is not a mathematically sound way of doing the analysis, it works at a high level. However, mathematically, there is also a way to detect encryption. One of the qualities of strong ciphertext is that the output of encryption is random. This is important because given the ciphertext of a message, there should be no way to determine the plaintext message. If the ciphertext is not random, that means there are patterns that can be used to predict the plaintext message, which is not a good thing. However, this quality of strong encryption can also be used as a detective message.

With a normal message, all characters do not appear with the same frequency; some appear more than others. This should be evident if you have ever watched the Wheel of Fortune—people will always try some letters before others. Samuel Morse, who invented the Morse code, wanted to make it as efficient as possible, so he would give the simplest codes to the letter that appear with the highest frequency and the longer codes to the letters that appear with less frequency. He determined frequency by counting the number of letters in sets of printers' type. The following are the values he came up with:

12,000	E	2,500	F
9,000	T	2,000	W, Y
8,000	A, I, N, O, S	1,700	G, P
6,400	H	1,600	B
6,200	R	1,200	V
4,400	D	800	K
4,000	L	500	Q
3,400	U	400	J, X
3,000	C, M	200	Z

Therefore E, followed by T, are the letters that usually appear with the highest frequency. Based on this analysis if you plot a histogram of the letters in a typical document, you would get a very uneven histogram that would look like the one shown in Figure 2.23.

Figure 2.23 A Histogram Plotting Letters in Normal Text

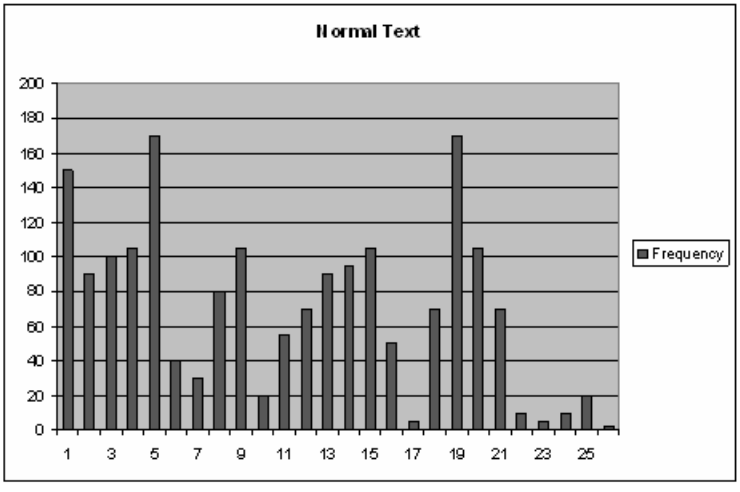

However, if you plot out the histogram for encrypted information, since ciphertext by nature is random, you will get a very flat histogram (see Figure 2.24).

Figure 2.24 A Flat Histogram Plotting Letters in Encrypted Text

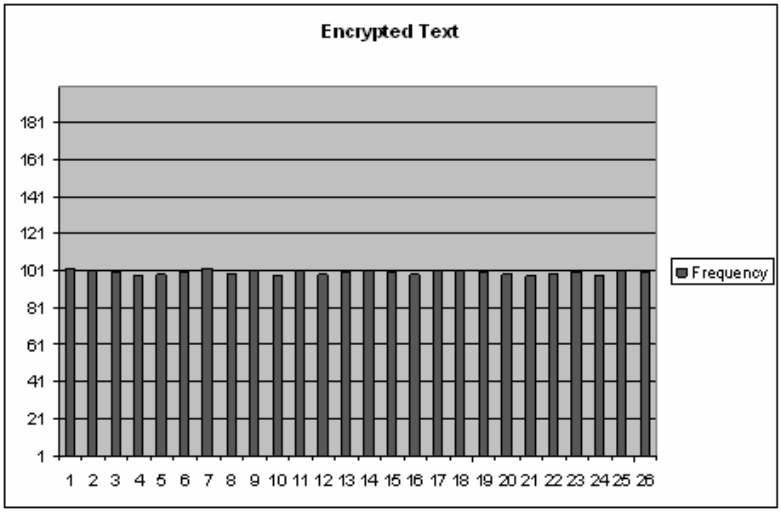

By looking at character frequency of the data you are examining you could detect whether the information is encrypted or not.

The good news is an organization can tell if information is encrypted; the bad news is if robust encryption is used you might never be able to decrypt the information and figure out the meaning of the message. However, in many cases, the mere fact that someone is encrypting information is enough to suspect them and watch them closely. If the organization does not have a sanctioned encryption program, in essence no one should be using encryption. The fact that someone is using encryption means they are using their own noncompany-sponsored program. This immediately raises a question. Second, the insider probably is only using encryption for certain communication, not all communication. Now the question becomes, who is the insider talking to and what are they saying that is so sensitive that he or she feels the need to encrypt the data during this transmission? Even though you might never know the actual content of the original message, you should be able to perform enough analysis to be able to warrant a more detailed investigation, and watch the person very closely.

Steganography

Steganography is data hiding, and is meant to conceal the true meaning of a message. With cryptography, you know someone is sending a sensitive message, you just do not know what it is. With steganography, you have no idea that someone is even sending a sensitive message because they are sending an overt message that completely conceals and hides the original covert message. Therefore cryptography often is referred to as secret communication and steganography is referred to as covert communication.

The analogy I like to use to differentiate between the two is the following. If I inherited a lot of gold I would want to keep it safe and protected. One option would be for me to purchase a big safe and keep it in my living room with all the gold locked inside. If you came to my house for a drink you would see the big safe in my living room. It would be locked so you could not open it. Even though the information inside was protected and you had no idea what was inside, you would know I must have something very valuable to lock it in such a big safe. That is cryptography. You do not know what I have but you know I am protecting something.

Another option would be for me to cut out the sheet rock on one of the walls, make a shelf, and stack all the gold on the shelf. Then I would put a picture on top of it on the wall. Now if you came to my house there was nothing unusual. I have some pictures on my walls, but it is quite normal for people to hang pictures on the walls. In this case nothing looks strange and you have no idea that I have anything of value. That is steganography. The sensitive gold is hidden and you have no idea it is there. The problem with straight steganography is if you noticed that the picture was crooked and went over to the wall to straighten it out, you would notice the wall is cut out behind the picture, and when you lifted up the picture you would see all my gold.

To maximize the efficiency of both technologies they often are used together. First I would install a wall safe and put the gold in the wall safe (cryptography). Then I would put a picture on top of the wall safe to conceal the safe (steganography). Together they provide a robust solution for inside attackers to protect their data and hide their information so no one at the company has any idea of what is going on.

To illustrate the power of steganography, we will use a program called S-tools. S-tools uses a simple drag-and-drop interface to perform steganography (see Figure 2.25).

Figure 2.25 The S-tools Interface

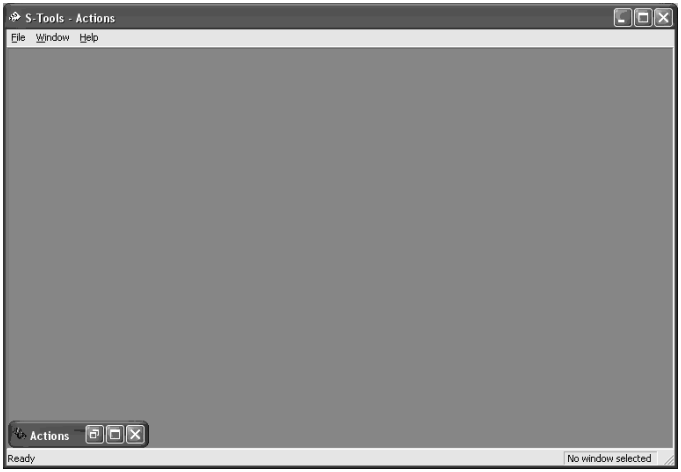

When S-tools hides information within a file it not only uses a pass-phrase to protect the information, it also uses encryption to keep the information secure (see Figure 2.26).

Figure 2.26 Using S-tools to Hide Information within a File

After you type the required information, S-tools automatically will hide the covert message within the overt message (see Figure 2.27).

Figure 2.27 Hiding a Covert Message within an Overt Message

The image on the left is the original image and the image on the right is the one with data hidden in it. Based on the how the algorithm works there is no visual difference between the two files. The binary composition of the two files are different, but visually, they are identical.

Steganography is a very popular technology for attackers because it allows them to completely conceal the true meaning of the communication and hide it in an overt image of any type.

There are two general ways to protect against steganography on your network: removal and detection. In some cases an organization does not care if someone is using steganography, they just want to make sure sensitive information does not leave the organization. Typically the easiest way to remove steganography is to convert the image to a different file type. For example the most common type of files on the Internet today are .jpeg files. The reason .jpeg files are so popular is because they are compressed files and are smaller than normal files. Therefore if you take a .jpeg file, convert it to a .bmp file, and then convert it back to a .jpeg file, although the image will look the same, the binary content will be different. During uncompression and decompression, the least significant bits are modified and this is where data is hidden with steganography. This conversion will remove any data that has been hidden within the file.

The second method is detection. For each steganography technique there are different ways to detect whether data has been hidden, but with S-tools it is fairly straightforward. S-tools works by making duplicate colors in the color table.

Therefore by writing a simple program to check for duplicate colors, you can tell whether data has been embedded within an image.

Normal File:

D:\DH\Data\BMP>bmpmap test.bmp

Filename: test.bmp

Actual size: 66146

Reported: 66146

Near duplicate colors: 2

Embedded File:

D:\DH\Data\BMP\STools>bmpmap test_h.bmp

Filename: test_h.bmp

Actual size: 66614

Reported: 66614

Near duplicate colors: 1046

You can see that in a file without data embedded within it the number of duplicate colors is less than 20. In a file that has data embedded within it, the number of duplicate colors is much higher than 20. Therefore a simple check of the duplicate colors would allow an organization to be able to detect steganography.

Malicious Acts

Malicious attacks typically are used by external attackers to gain access to a system. Since someone on the outside does not have access to critical systems or files they often need to use exploitation methods to attack a system and get access. Typically with an insider attack, the insider has access and the challenge becomes how to hide and extract the data from the organization. In addition, since most insiders are non-technical, exploitation methods are used only as a path of last resort.

Although the use of malicious acts as the primary measure for insider attack often does not happen, to gain additional access or elevated privileges malicious acts do come into play. In a situation where I have some access but not all the access I need to access the data, I might use a malicious attack to elevate my privileges or gain access to a system. These attacks usually involve running exploit code against a system. Some insiders may write their own exploits or even find their own zero day

exploits, but many will use tools that already have been written. The advantage of utilizing zero day code is because most IDS would not be able to detect it or find it because there is no known pattern or signature.

Malicious attacks can also take other forms including the following:

- Deceit and deception
- Perception management
- Information corruption
- Theft
- Destruction
- Unlocked system

Although not the main focus or technology normally used for insider threat, in some cases an insider can reach deep into his or her arsenal and use these techniques to gain access to the information needed or to cause harm to the organization.

The Human

The weakest link in any security problem is the human factor, and insider threat is no different. No matter how well thought out your security plan is, no matter how much technology you utilize, no matter how many resources you have, humans still will be one of the key areas of focus for the insider to do harm.

Even though the insider typically has some access needed to do the job, often additional access is needed in order to acquire the access he or she needs to fully compromise critical IP. There are many ways that can be done, and we looked at several already, from malicious attacks to using out-of-band systems like voice mail. One of the easiest methods is to compromise the human element and have a human tell you the information you need, or at least give you enough data points to be able to make your job as the attacker easier.

This concept of essentially tricking someone into giving you information they normally should not give you is called *social engineering*, and it has been around forever. When my wife accuses me of lying to her, I respond jokingly that I am not lying, I am social engineering her (which, by the way, I do not recommend if you are married!). In short, social engineering is lying, it just sounds better than saying you are a liar.

The formal definition of social engineering is pretending to be something that you are not, with the goal of tricking someone into giving you information they normally should not give you and that you should not have access to. Essentially, if

the person knew your true intent, he or she would not tell you the information, but since you are "in disguise" you are able to trick him or her. Some simple but effective examples of social engineering are:

- Calling the help desk requesting a new account to be set up
- Calling IT and acting like a vendor to find out how a piece of software is configured or where it is installed
- Impersonating a manager to get an employee to send you a proposal
- Impersonating a manager to get approval to access a sensitive directory

The interesting fact about social engineering is that many people overlook it because it seems so simple. But remember, just because something is simple does not mean that it is not extremely powerful.

Most people have a default level of trust—if you do not do anything to convince them otherwise, they generally trust you. Some of the other reasons why social engineering is effective include the following:

- **Exploits human behavior**. As we stated, if you are nice and friendly most people will trust you. Most people's personalities assume it is better to have a default mode of trust unless you do something to change their minds, than to go through life paranoid and not trusting or talking to anyone.
- **Trust is good**. Some level of trust is not a bad thing. When I meet someone for the first time it is much better to start talking and assume some level of trust than to sit in the corner thinking everyone is evil. But always remember the saying from the movie, *Italian Job*, "I trust everyone, it is the devil inside I do not trust."
- **If you are nice you must be honest**. Despite all the effort that we spend teaching our children not to trust strangers, most people grow up still thinking that it is OK to talk and socialize with nice people. Although interacting with people you do not know can be intriguing and give you a new outlook on life, you always have to be careful that there is a distinct difference between friendliness and trust.
- **People love to talk**. This is probably one of the biggest reasons social engineering is so effective. If you just give someone an opportunity and ask questions that are nonthreatening, you can get someone to tell you almost anything.
- **If you listen, you must be trustworthy**. Many people are lonely, and one of the signs of loneliness is that they have no one to talk to. Actually,

the real sign is that they have no one who listens to them. Many people just want to hear themselves speak and when they talk with someone it is a competition of who can get more words in and dominate the session. Many people are not good listeners. Therefore if you are a good listener people are more prone to talking. It is critical to remember there is a difference between listening and not talking. I had one person tell me that when he talked with someone, he let the other person talk 90 percent of the time, therefore he was a good listener. You cannot make that correlation. Just because someone is not talking does not mean he is listening; he could be daydreaming or thinking about other things. Listening is a process of engaging with another person, understanding what he is saying, and being sincere. If you talk with someone who hardly talks, it could be because he could tell you were not interested and not really listening.

- **People love to tell secrets**. By having a secret that no one else knows people feel empowered. However, that empowerment is diminished if you do not tell anyone about the secret. Therefore the fact that someone knows something no one else knows, they have this eagerness to tell someone. In addition many people love to gossip and find out what other people are doing. Taking this and carefully planning a conversation can be a perfect avenue into social engineering.

The other item that makes social engineering so popular is that you need minimal, if any, information to perform social engineering. Although no information is required to perform social engineering, the more information you have, the higher the chance of success. If you are going to pretend to be a manager or from a different department, the more supporting data you can use or names you can throw out to put the person you are talking to at ease, the greater the chance of the attack being successful. Typically with an external social engineering attack, some general recon is appropriate so you are not going into it blind. However, if you are talking about internal social engineering, although you can perform additional recon, most insiders just by nature of their job have enough information to make the attack highly successful.

Social engineering can come in many flavors and types. I have found that just by being nice and helpful, usually you can get whatever information you need. However, in cases where that does not work, the following are some additional types of social engineering that can be used:

- **Just asking**. The most general type of social engineering and the one that often is overlooked because it is so obvious is just asking for the information. In many cases if you are confident about how you present yourself

and you ask for something with authority, many times people will give it to you. Many buildings have guards in the front of a building but if you walk in the building you do not need an escort. If you walk in and look around and act like a visitor you will be stopped and have to sign in. However, if you walk with confidence as if you belong and know where you are going, in many cases, you will not be stopped.

- **Impersonation**. Pretending to be someone else is probably the most popular method of social engineering. If I pretend to be from the help desk I can probably get an executive assistant to open a CEO's office so I could supposedly fix a problem he was having. Even if the assistant says I do not know about any problem, you can say he was complaining to my boss about a problem, but if you prefer I just leave it and wait until he gets back that is fine. Most people do not want to be the one to be blamed for not having a problem fixed, so in most cases you would be given access.
- **Misleading and redirection**. With both of these you are giving someone false information with the purpose of misleading them, confusing them, or causing them to make a decision they normally would not make. For example, I heard a rumor that we had a fire code violation, and sometime this weekend they are going to come in and inspect the office. If they cannot get into an office, they might fine the company, so make sure you leave your office open this weekend. Clearly you are presenting someone with false information so you can gain access after hours.
- **Anger**. Although I would recommend using this method as a last resort, it will work in some cases. Instead of being nice, kind, and trying to sweet-talk someone into doing something, you get very angry and raise you voice. Some people do not like conflict so if you start getting upset and angry they would rather just give you what you want than risk making a scene. The risk of anger social engineering is that you could annoy the other person and he or she could get angry; then nothing would get accomplished.

To defend against social engineering, the best defense measure is to raise user awareness on the dangers that social engineering can pose to a company. In addition, strict policies with clear guidelines that are tightly enforced will also limit the chances that people could be tricked into given access. Employees have to realize that a company is taking social engineering seriously, and if it is enforced and proper measures are taken, the amount of social engineering attacks will decrease. This is not a problem that will be solved instantly, but with time it will get better.

Summary

This chapter covered a wide range of technologies and methods that can be used by an insider to cause harm to a company. Just like with viruses and other attacks, attackers often are taking a base method and modifying or creating variants of it to make it more powerful and difficult to attack. The same thing will occur with the insider threat. This chapter covered the base technologies and some variants, but expect to see additional variants and new methods evolve. As companies become more and more savvy on prevention and detection techniques for the insider threat, attackers are going to be forced to enhance their means and methods of using these technologies, and even in some cases, developing new technologies.

This often is referred to as the leap-frog approach. Either the "good guys" (security professionals) or the "bad guys" (attackers) can start the game, but usually it is the attackers. Most security measures are put in place because of a need; if there is a not a need, why spend money on something? Remember from our analysis that threat drives the train. If there is not a threat, than whether there is a vulnerability or not is irrelevant. Threat is tied back to possible danger and created by attackers. Every time an attacker develops a new method or technique for exploiting a system, he or she is creating a new threat, and in most cases, there is a high chance the company has a resulting vulnerability. The reason is simple: attackers develop methods to break in and cause harm, which is manifested through a vulnerability. If there is not a vulnerability then the attacker will spend all this effort on a technique that will never amount to anything of value. Attackers are very smart and clever and usually develop threats out of a need because they saw an opportunity they want to take advantage of.

Based on these reasons, we assume that the attacker will start the game. The attacker will leap-frog over the security professionals. This usually is done by the attacker finding a way into a system or a way to compromise data from an insider perspective, that a company has not thought about and has an inherent vulnerability to the attack. Once this starts being exploited by the attacker, security professionals start figuring out a way to fix the problem and stop the attack. Once they are able to do this they leap over the attacker. At this point the attackers' chance of success is either stopped or greatly diminished, and they have to figure out a new way to accomplish their goals, which usually entails the compromise of sensitive IP. In most cases they will modify an existing attack, or in some cases, create a whole new attack that once again is successful. Now they have leap-frogged over the security professional. This game usually continues, and is what guarantees that security professionals will be employed for a long period of time.

The problem with any new area, and insider threat is no exception, is the time it takes for security professionals to leap over the attackers. This usually takes a long period of time; however, the time it takes for attackers to leap-frog over the security professionals is usually fairly quick. Therefore it is critical that we are properly prepared going into this battle. The more we can understand the technology that is behind the attack, and start thinking about both prevention and detection measures, the better chances we will have at protecting our organizations from attack.

We covered a lot of different technologies that are used by attackers. Some of the technologies you might have known about and we presented a different perspective; some were very cool, and ones James Bond might use, and others seemed very basic. You might have questioned why something this basic was in this chapter. Remember that this chapter is not titled, "Cool and Crazy Technologies," it was meant to give you some insight to what is used to commit the crimes. Attackers do not care if the method is cool or not, they care about getting the job done. If the technology will help them accomplish the goal, they will use it. In most cases, if you have two ways to accomplish something—one that is very complicated and high tech and one that is simple and straightforward—and they both have the same chance of working with the same success rate, an attacker would always pick the more straightforward one. The simpler something is, the less chance for mistake. Also, if a simple method will work as good as a more advanced technique, why waste the advanced technique? An attacker only wants to use just enough energy to get the job done. Anything else could raise there profile, have a higher risk of getting caught, and give away an advanced technique.

Therefore, you have to look at how effective any technology that we review is going to be. Even though it might seem basic, the question you have to ask yourself is whether it work in your organization. If it will get past your security defenses and cause an attack to be successful, then you have to worry about it. The good news you should say to yourself is, if it is a basic attack, it should be easier to defend against than a more complicated attack.

In this chapter we dissected each technology by looking at each one separately; however, in practice, many are used together. Traditionally some form of social engineering usually comes into play before someone actually extracts IP out of the company. Essentially the social engineering piece is a good way to test the waters to try and find out as much information as possible to minimize the chance of getting caught and maximize the chance of being successful.

Another example is that if I am going to use steganography, I am always going to use it with encryption. This way if someone actually determines that steganography is being used, they will not be able to read the content of the message. In addition, if my ultimate goal is to get sensitive IP to a competitor, even if steganography is going

to be used, I am not going to send it directly to someone at the competitor's e-mail address. Even though it might be pictures of my vacation, it could raise the suspicion level and create a link that I would prefer didn't exist. Therefore I would encrypt the IP, then hide it with steganography and send it to an e-mail alias outside of the company. Then from my personal home connection with no link to the company, I would download the message, extract the content, and even use relays before I send it to the competitor. Even though my company would have a hard time of finding out, I do not want anyone to put a link between me and the company; therefore relays provide an extra level of protection.

As you think back about the chapter, think about how an attacker would mix and match these different technologies to cause harm to your organization. The more you can think like an attacker and understand how they operate, the better you will be at defending the insider threat.

Technology is the foundation for how an attacker would compromise a company, but it is the raw materials of the attack that have to be fashioned and refined to come up with an actual attack. The refinement from taking a core technology and putting together an attack is the means and methods of how attackers operate. That is the level we need to think about. Understanding the technology is important, but what would be the means and methods an attacker would use to launch the attack is the end state. Being able to prevent and defend against a given technology is critical, but being able to protect against an entire attack is the ultimate goal. By knowing the means and ways an attack is going to use a technology against you will help you build the best defense possible. By understanding the full method, there might be easier ways to be able to defend against the attack ,and ultimately the more you know, the more in-depth defense you can apply to securing your organization.

The way this usually is done in practice is to create scenarios. Taking the information you learned in this chapter about the technologies and by understanding the real threats, you should build a series of scenarios of how an insider would compromise your organization. Try to put as many details into the scenarios as possible. Once you have the scenarios, figure out how you would defend against it and what changes you need to implement so the attack will not be successful. In order to do this it is critical that you put the scenarios in priority order, focusing heavily on likelihood of the attack. If the scenario has a high chance of occurring and could have a huge loss, you would spend more time defending against it than if it had a low likelihood of success.

Just to give you a sample, the following is a high-level scenario that can be used as a starting point:

- **Executive targeting**. After careful analysis, one of the employees working at your organization has identified that whenever the CEO travels, he has all the critical corporate IP on his laptop. By accessing help desk tickets they know that this executive has complained numerous times while on travel that he does not have the information he needs to do his job. To fix the problem they load all the data onto his system prior to his trip. The insider also knows that the executive has been locked out of his system many times and therefore his password is always set to the same password. Through social engineering they can acquire the password from the executive assistant. The insider checks corporate bulletins and knows the CEO is giving a presentation in California next Tuesday, so he will be traveling on Monday. By watching the pattern of the executive, the help desk loads his system with all the data prior to him leaving and then he comes in the morning of his trip, picks up his laptop, and goes on the trip. By talking with the help desk they are going to configure the laptop on Friday and leave it on his desk Friday evening so the CEO can pick up the laptop on Monday. This represents a window of opportunity over the weekend to go into the office, log in to the computer, copy the data to a USB drive, and leave.

  Likelihood: Very high.

  Impact: Very high.

  Problems: Once you put together a scenario and determine that it has a high likelihood of occurring and the impact could be high, you then want to figure out what the problems are, and possible solutions. The problems are people not being properly training on social engineering: the CEO not being made aware of security and presented with alternatives to accomplish what he wants with the same ease but more secure. Finally, sensitive corporate records and the laptops they are stored on are not properly protected.

For purposes of this book we listed a high-level scenario to serve as a starting point, but it should be developed in a lot more detail to make sure your company is ready and prepared to deal with the insider no matter what method of attack is used.

Scenarios are critical because the more planning you perform the better off you will be in the long run. Although you can never predict the future with 100 percent accuracy, the more you can do to figure out what is going to happen and plan for it, the better off you are. In the military they are constantly running scenarios and training against them because they know it is not 100 percent accurate, but it will be pretty close, and the more your organization can be prepared the better off you will

be. The more time you take to create the scenarios, the more valuable they will be in a preparedness and training tool. The less time and thought you put into them, the less value they will have.

Although this chapter talked about some cool technologies, the chapter was about awareness and allowing you to use that knowledge to build defenses before attackers strike.