6.915 Computer and Network Security **Fall Term, 1995**

Lecture 10 : October 12, 1995

Lecturer: Ron Rivest *Scribe: Matthew Condell*

1 Introduction

In this lecture we will discuss the following topics:

- NFS attack in the New York Times.
- Biometrics.
- Mutual Authentication.
- Other Attacks.
- The Future.
- Auditing and Intrusion Detection.

2 NFS Attack

In yesterday's New York Times an overblown article by John Markoff (handout 16) publicized an often overlooked security flaw. The flaw, pointed out earlier in a posting to a cypherpunks mailing list, is that people usually trust untrustworthy paths to executables. Spoofing clients can be fairly easy for an enemy who is sitting between a client and a server. Upon seeing a request for a program, the enemy can send back its own version of the program faster than the file server. The client will accept this version and assume the version from the server is a duplicate so it will discard it accordingly. An effective way to prevent this kind of attack is to use digital signatures on the programs.

3 Biometrics

Biometrics is the use of a person's physical characteristics for authentication. Biometrics are generally not as good as passwords since error rates under 1% are hard to achieve. Passwords, while they can be compromized, do not have any false rejects. Among the biometrics we will look at are:

- Voice identification
- Signature dynamics
- Fingerprints
- Hand geometry
- Retinal scan
- Facial recognition
- Typing characterstics

When evaluating biometrics, we are concerned with the time it takes to measure each metric, the equipment expense, the user acceptability of the method, and false-acceptance and false-rejection rates. When looking at false-acceptance and false-rejection rates, we often look at the equal-error error rate. This is the point on a graph of percent error vs. threshold set for a match, where the false-acceptance and false-rejection rates are equal.

3.1 Voice Identification

The first problem when designing a voice recognition system is how to recognize the voice. Pitch doesn't work too well so the method often used is modelling the characteristics of the resonant system of the vocal tract using linear predictive coding (LPC). Voice recognition has many similarities to work in speech recognition.

There are many problems with speech recognition that make it unpopular with people who use it. It requires a tedious enrollment process for the computer to learn users' voices. It's effectiveness depends upon ambient noise levels. For example, if a train goes by while you are trying to authenticate yourself, the authentication will probably fail. Also, voices tend to change both over time and when users are sick, increasing false-rejection rates. Finally, many people just do not like talking to a computer.

A couple of technological problems also make voice identification difficult. It requires a lot of complicated processing including procedures such as Fourier transforms. Additionally, replays are fairly easy because all one has to do is record a person saying the pass phrase and replaying it to the computer. One way to protect a system against such replay attacks is the computer can display a different pass phrase that the user must read for each authentication. This makes replays more difficult.

3.2 Signature Dynamics

As with voice identification, the simple solution for signature recognition doesn't work very well. Just comparing signatures for exact copies has problems because they are easy to forge and most people have lots of variation in their signatures. To solve this, more complicated dynamics are used such as pressure, time, velocity and acceleration in addition to the x and y coordinates of the signature. The processing is similar to voice identification where you compare the characteristics of each segment in the sequence that constitutes the signature.

Signature dynamics is not too expensive to implement, but does have some problems. While it allows for some variation in individuals' signatures, it won't accept all variations. It is also prone to forgery. One way to help reduce the forgery problem is to use a similar method of changing pass phrases as described above.

3.3 Fingerprints

Fingerprint analysis is a fairly straight forward comparison of the characteristics of a user's fingerprints such as arches, loops, whorls, etc. (see page 59 of the *Computer Security Handbook*). While it has a fairly simple interface, consisting of placing a hand on a scanner, it also has several problems. The transducer needs to be well done to be accurate, though it is not perfect and getting dirt on it can lower the acceptance rate. It's accuracy is also dependent upon the finger that is being identified. The technology required to implement the scanner is relatively expensive. Additionally, there is a social stigma against using fingerprints since they are associated with their use for criminal identification.

3.4 Hand Geometry

Hand geometry involves measuring the physical characteristics of the user's hand to determine identity. Measurements such as the lengths of the fingers, width of the hand, and finger thickness are used. This is one of the best biometric methods in terms of accuracy and user acceptance. It is questionable about how easy it is to fake out this method.

3.5 Retinal Scans

A 1935 study by Simon and Goldstein showed that the pattern of blood vessels in the retina, like fingerprints, are unique to every individual. This means that by taking an infra-red scan of a person's eye one can match the pattern to that of a known template using a technique called Fourier transform cross-correlation. This technique is not popular with users since it is not comfortable and reequires a light to be shined into one's eye. Another problem with this method is the alignment of the head in relation to the scanner is critical to make a positive identification.

3.6 Facial Recognition

Work on developing facial recognition techniques is currently being done by Tommy Poggio's group at the AI Lab. It is modelled after how we recognize each other–by sight. The computer takes a picture of the user's face and compares it to a picture it has stored on disk. The technology for this is still being developed, but probably won't be ready for use for another ten to twenty years.

3.7 Typing and Mouse Characteristics

A final technique is to recognize a user by the characteristics of his typing or use of the mouse. This was commonly used during World War I to determine which operator was typing the German's Morse code. This method can have problems if the user hurts his hand or his typing style changes.

3.8 Sandia Study of Biometrics

A study done by Sandia National Labs (Handout 11) compared the effectiveness of commercially available biometric equipment. The crossover error rates for the methods studyed are summarized in the table below:

Technique	**Error Rate**
Voice (Alpha)	3%
Voice (ECCO)	2%
Signature	2%
Retinal scan	0.4%
Hand geometry	0.1%
Fingerprint	9% false reject, no false accepts

They also compared other characteristics of the techniques:

Characteristic	**Best**	**Worst**
User acceptability	Hand	Voice
False reject	Hand	Fingerprint
False accept	Hand, Retina, Fingerprint	Voice
Throughput	Hand, Retina, Fingerprint	Voice, Signature
Template size	Retina	Voice
Difficulty of imitation	Retina	Voice, Signature
Cost	Voice	Retina

As mentioned above, these biometrics are not big improvements over passwords because of their high error rates. They can be used in conjunction passwords, though, so that the system checks both what the user is and what the user knows.

3.9 Mutual Authentication

Mutual authentication is a dificult problem since it is often harder for the user to recognize the computer than the computer to recognize the user. How can the user be sure he is actually talking to the correct system? Rebooting the system before logging in is effective, but has obvious drawbacks. Using trusted devices, such as smartcards, is a possibility. Another is a two-level password system where the user enters a password, the computer gives a response, then the user can enter a second password. A zero knowledge protocol can be used to make sure the system is authentic. Finally, some systems, such as Sunos, hardwire logins so that only the log-in program can have a completely white screen, while all other programs must have some kind of border around their screens.

3.10 Other Attacks and Ideas

There are many vulnerabilities in authentication systems that can be exploited. For example, a bug in TENEX allowed intruders to gain partial information about passwords. TENEX does password comparison one character at a time and would invoke a page fault when a bad character was entered. This allows an enemy to gain information about which characters of a password are correct and which are not. Another bug is in non-atomic password checking. An interrupt durning the password checking process can leave you logged in.

Some other vulnerabilities not discussed in lecture that Prof. Rivest had in his notes included: sniffing on ethernet and X-windows, keeping default passwords that are unchanged from the manufacturer, trojan horses and by-passing the need for passwords using `.rhosts`.

Coersion is another problem that needs to be dealt with. How do you prevent someone from gaining enterance to an area by forcing a legitimate user to make his way through the security precautions? One way is to only allow one person to pass through an entrance way at a time. A second is to have emergency passwords that will allow a user to gain entrance, but will alert someone that there is a problem.

3.11 The Future

Now that we have an idea where we currently are technologically, let's look where we might be headed in the future. Smart cards are a possiblility, along with computers that you can wear on you clothes (like on your belt) that can represent you in all transactions. It may also be possible just to continually monitor everyone's location using a system like GPS. Other biometrics may be used in the future, too, including DNA and brain waves.

4 Auditing and Intrusion Detection

Why would we want to audit the activities on a system? Why go beyond only using accesse control methods? The short answer is that there is no perfectly secure system so more fences are needed to make the system more secure. Auditing is a backup when access control fails.

Auditing is similar to a surveillance camera–it allows you to watch an intruder once he is inside. It can provide evidence that can be used to identify and possibly convict an intruder or an authorized user who is misbehaving. It can detect unusual behavior and provide warning that other defenses have failed. It can help a system to recover from damage by identifying what parts have been hurt and can help identify other systems that may have been damaged, too. These features may help auditing to act as a deterrent to intrusion. It is important to remember, though, that it only detects intrusion and doesn't prevent it.

Intrusion detection is very difficult because it is hard to define what is an unusual or suspicious activity as opposed to a normal one. Once you have caught an intruder, though, what do you do about it? You can shut down the system which will keep the intruder out, but also keep legitimate users out. The weakness that allowed the intruder to enter can be repaired. You can call CERT (the Computer Emergency Response Team) or FIRST to alert them of the problem. This could be a security problem since they post the flaw and/or a solution to a bulletin board to alert others of the problem. Some hackers use this bulletin board for ideas. Finally, you can allow the intruder to enter the system (either by the old security flaw or create a flaw to use as a trap) and monitor him.