Biometric System Compromises

Theresa E. Elam
ForS 233: Security Management
November 19, 2002

In an information based world where identity theft and credit card fraud is a $4B+ market [1], corporate network break-ins are on the rise, and where information systems that hold and process highly sensitive data are connected to the Internet, the security industry is looking for better ways to prevent unauthorized access to systems. Current methods such as username/password, smart cards and bank card/PINs while functional, have well-established vulnerabilities and no longer hold the promise they once did. New, stronger methods of access control are being sought. In this regard, many have turned to biometrics, the use of unique human traits for identification purposes, as the answer. Although biometric systems have great potential, they also have weaknesses that must be understood before implementing them as a part of a security strategy.

Biometrics and Identification

Academically, the study of biometrics is the statistical analysis and measurement of human physiological and behavioral characteristics. More commonly (and certainly to the security community), the term "biometrics" refers to the automated use of those measurements to prove or disprove a person's identity. In the United States, we have been using biometrics for identification since 1902 when the New York Civil Service Commission, in order to prevent people from cheating on entrance exams, started fingerprinting applicants [2]. Today fingerprints are just one biometric method used. Others include facial recognition and thermography (temperature "signature"), hand geometry, retinal and iris scans, handwriting, and voice prints.

Biometrics is a method of identification, but what is identification? Simply, "personal identification is the process of associating a particular individual with an identity [3]." In the security realm, this process has two components -- recognition and authentication. Recognition asks the question "Who are you?" and seeks to discern which person, among many, you are. Authentication seeks to verify that the identity claimed is indeed the identity possessed. In day-to-day activities, recognition is commonly used as actual identification; when you meet your friend Alice on the street, you do not ask her to verify that she is, in fact, Alice. For security systems, however, this is not sufficient. Without an authentication phase we do not know if we are granting access to a friendly or malicious user.

Traditionally, automated (the type that is needed for security systems) authentication is based on a token or information. In other words, based on what a person has or something a person

knows. Tokens used for authentication include keys, driver's licenses, passports, or employee badges while passwords, personal identification numbers (PINs) and mother's maiden name are conventionally used pieces of verification knowledge. These methods have some obvious limitations. Tokens can be lost, forged, stolen or surrendered (someone is paid for their authentication token) with relative ease and knowledge can be forgotten or discovered (both with the help of the original owner and without).

With the use of biometric technology, a third method of automated authentication is added to the list -- what a person is. As stated above, a physiological or behavioral characteristic of a person is used to authenticate whom they say they are. Biometric authentication methods have the advantage over traditional authentication methods in that they are unique to a given person and they cannot be lost. It is also generally thought that biometric characteristics cannot be stolen or forged. This paper, however, will demonstrate that premise to be false. Further, once compromised, biometric characteristics cannot be replaced by destroying the characteristic and creating a new one the way you can with passwords, for example. Another unique disadvantage is that biometric characteristics are public in nature and it is difficult to increase their level of obscurity.

It should be noted that biometric characteristics, unlike tokens and knowledge, could be used for recognition *as well as* authentication. This is due to the fact that biometric characteristics are unique to a given person. This is not the case with token- and knowledge-based identity approaches. In these systems, it is the *combination* of whom the user says they are along with the token or knowledge presented that creates a unique identification. Biometrics offers a way to handle both recognition and authentication in one step.

Biometric Systems

A biometric system refers to the collective procedures and technologies by which a biometric characteristic is captured and utilized in the authentication process for security purposes. In general, biometric systems have four elements: enrollment, storage, access verification and the acceptance error rates. It is possible to find definitions of a biometric system that have more or fewer elements, however, this breakdown is best suited for discussing biometrics as they relate to securing sensitive systems because each of these elements offer opportunities to compromise the system (as discussed later in the paper).

Enrollment is the process by which a user and their profile (comprised of personal information along with access privileges) are entered into the system. Using a biometric reader, the chosen biometric of the user is digitally captured. The digital representation is then furthered processed to extract the unique characteristics of the biometric into what is called a "template". This is done for space and performance considerations -- using the complete digital representation for even a small-scale system would require prohibitively large amounts of storage space and processing power to be commercially feasible. The characteristics that are used in a given template are system dependent.

After a user has been enrolled in the system, the template must be stored in some manner to be used later in the access verification process. Storage methods include a database, smart cards, or files on the user's personal workstation.

Access verification is the final element of a biometric system and involves the matching of a template to the biometric characteristic presented to the system for access. In this process, the user presents the biometric (e.g., a fingerprint) to the system via a reader (usually one similar, if not identical to the one used in the enrollment process), the biometric is processed into a template and compared to those already into the system. If the template is valid the user is granted access to the system.

There are two acceptance error rates used to determine the accuracy of a biometric system. They are the rate at which a system will reject an authorized user -- known as the false rejection rate or FRR -- and the rate at which a system will accept an unauthorized user -- known as the false acceptance rate or FAR. In a perfect system, these two rates would be zero. In practice however, not only is this impossible to achieve, but the reduction in one rate will often increase the other. For example, if you increase the sensitivity of your fingerprint system so that it accepts very few unauthorized users, chances are you will also increase the number of rejected authorized users (with the system at such a high sensitivity level someone with exceptionally dry hands may be falsely rejected). Because the two rates are interdependent, it is customary to plot them against each other to find the point at which they cross for given sensitivity levels of the system. This crossing point is called the equal error rate or EER (also called the crossover error rate or CER).

Methods of Compromise

Methods of biometric system compromises (a.k.a., attacks or hacks) can be categorized either generally by what part of the system they attack or specifically by type of biometric. Both categories will be discussed here.

System Compromises

As stated above, a biometric system is comprised of four elements; all containing methods by which a malicious user could gain unauthorized access. They are:

Enrollment - The compromise that is of most concern to the enrollment process of a biometric system is that of fraud. This is where the user who is enrolling for the system is not who they say they are but present the required credentials to enroll in the system. This is especially dangerous for systems that wish to use biometric technologies as a method of identification (and not necessarily access to a protected system). For example, the United States Immigration and Naturalization Service (INS) developed the INS Passenger Accelerated Service System (INSPASS) to reduce immigration inspection processing time at airports [4]. Travelers who wish to use the system fill out a form and go through an enrollment process at an INSPASS enrollment center that includes a hand scan for the creation of a hand geometry template. Upon successful enrollment the traveler is issued a card containing the template to be used at an INSPASS kiosk in designated airports. At the kiosk, the traveler inserts his/her card, answers a few questions via a touch-screen display and then places their hand on the biometric reader when prompted. If all goes well, the traveler is whisked through customs in a fraction of the time and with no human intervention or interview.

However, what if the person who enrolled is not who they claimed to be? This system relies solely on the fact that the person is who they say they are at the time of enrollment, which, with the proper forged documents, does not need to be the case. Interestingly enough, according to the Department of Justice and the International Biometric Industry Association (IBIA), the INSPASS system is being phased out. Richard Norton, the executive director of the IBIA is quoted as saying, "INSPASS taught us a lot of very good lessons about how to implement a system of border controls for what are now called registered travelers [4]."

Storage - Where the templates of a biometric system are stored is another vulnerability in the system. The most common storage mediums for templates are files in a database or on the user's personal workstation. In addition, More often than not these databases and workstations are connected to a network. A skilled cracker (the term for a malicious hacker) or disgruntled employee could gain access to the templates and thus have the "keys to the kingdom". Storing templates on smart cards does offer more protection for your system; however, smart cards are tokens and are thus susceptible to misplacement, theft and being sold to the highest bidder.

Access verification - It is at the access verification stage where most of the "sexy" compromises take place. This is where the criminal mastermind in a James Bond film places a forged contact in his eye to gain access to a storage facility of nuclear missiles. Faked biometrics are not as hard or uncommon as most people think and will be discussed in more detail in the "Type Specific Compromises" section of this paper. Also at the access verification stage are compromises known as "replay attacks". These attacks involve either causing the system to rescan a legitimate latent biometric (usually a fingerprint) or accept a recording of a biometric template.

A major factor that allows replay attacks to work so well is that most biometric readers are connected to a workstation via a Universal Serial Bus (USB) cable. Ease of use and not security was foremost in the design of the USB protocol, and as it turns out, the feature that most distinguishes USB is also the most dangerous from a security perspective. (This fact begs the question why so many manufacturers of biometric security devices would use USB as a cabling type, but the author is positive that being able to say "user friendly" in the marketing literature had something to do with it.) This feature is that USB devices are "hot-swappable" meaning that the device may be plugged into or unplugged from a computer while the computer is still running. This gives an attacker a perfect opportunity to plug-in a USB device containing a recording of a biometric template and play it for the computer. As long as the malicious device behaves as the computer expects (which is not hard to find out from vendor specifications) it will accept the recorded template as if the legitimate user were present. Obtaining the recording is not hard either. "Sniffing" is eavesdropping on electronic communications and can be done via software listening on the USB port or at the hardware level with devices that eavesdrop on the actual USB cable itself. Even though the information obtained by sniffing is just the template (i.e., the distilled, unique characteristics of the biometric reading and not the digital image itself), if one knows how the particular system creates its templates (again available from vendor

documentation) it is possible to recreate a copy of the digital image that the template was made from [4]. In the case of fingerprints, not only could the attacker gain access to your system, they can also reproduce your fingerprints!

Acceptance error rates - Compromises at this level of the system are akin to default or weak passwords. These attacks are allowed to happen because the system is tuned to such a lax degree of sensitivity that malicious users are often allowed in because their biometric happens to fall within the acceptable error range of that of a legitimate user.

Type Specific Compromises

In addition to the general compromises that can be used against a biometric security system there are attacks that can be targeted to the specific type of biometric characteristic the system employs. A few of these compromises are discussed below.

Voice recognition - Voice recognition is the process by which a particular speech pattern is associated with an identity. This area of biometrics has significant potential for many reasons. First of all, it utilizes hardware that already exists in most computers today -- the microphone. Secondly, due to the popularity of personal dictation software that allows users to speak into their computer and have what they say typed into a document, the software to run voice recognition systems is fairly mature. Finally, voice recognition does not require anything "unusual" from the user. We all are used to speaking into a microphone and thus the access verification phase of a voice recognition system is considered very user friendly.

The enrollment phase of a voice biometric system, on the other hand, can be quite tedious. Voice recognition works by transforming the spoken word into text (the template for this system). In order to do this, the system must be trained to identify how a particular user says a passphrase. The process can take several days depending on the size of the vocabulary you wish to teach the system. There are systems that are text-independent (meaning that the system recognizes a user's voice irregardless of what they say) which are quite sophisticated and difficult to train but offer more protection against attacks for the passphrase that is used to gain access to the system can be changed on a regular basis.

Compromising a voice verification system can be accomplished by playing a recording of the legitimate users voice for the system. Because voice recognition systems must take into

account ambient noise and distorted samples (from say, a person with a cold or different emotional states), they are often tuned with a fairly high recognition threshold. This makes the attacker's job much easier for the recording used to gain access does not need to be of necessarily high quality in order to work. There is obvious room for improvement in this area of voice recognition systems if they are to be considered as a viable part of system security.

Facial recognition - Facial features are perhaps the most widespread biometric characteristic used for the purposes of personal identification. The idea behind such systems are based on what we humans do countless times a day without thinking. When we encounter a person we know we do not stop to verify who they are because we *recognize* them. It is this process of recognition that facial feature biometric systems attempt to recreate. This is done typically be analyzing not only the shape of facial features -- eyes, nose, eyebrows, chin, etc -- but also their location and spatial relationships within the face. Facial recognition systems are designed in one of two ways -- to recognize a controlled, mug shot style photo or to recognize a person in a dynamic, uncontrolled setting. An example of the latter would be the Pentagon's Image Understanding for Force Protection (IUFP) which aims to be able to recognize people at a distance and in public situations [6]. Interest in the development of this type of technology has increased dramatically since the attacks of September 11 with the idea of installing such systems at airports and around high-profile buildings in order to identify persons that are suspected terrorists.

Regardless of the type of system environment -- controlled or uncontrolled -- facial recognition technologies can be fooled. For mug-shot type systems, a digital image of the legitimate user and a laptop is all that is needed [5]. The attacker takes several digital photos of the legitimate user and saves them to his/her laptop. To access the system, the attacker brings up the image on the laptop and presents it to the system camera. The "tricky" part of this attack is usually in getting the distance from the laptop to the camera just right. Some controlled facial recognition systems offer a higher level of security in that they require a "live" person to be presented for authentication. However, "live" simply means moving. This is quickly circumvented using the same laptop with a short digital movie clip of the legitimate user.

Uncontrolled environment facial recognition systems such as IUFP can also be fooled with the use of disguises that distort facial characteristics. Proponents of the technology say that the systems are being improved by adding movement characteristics such as a person's gait to the

list of identifiers. Possibly the belief is that such movement characteristics are so ingrained in a person that they cannot be altered and thus would reveal an identity. It is hard to believe that a determined person could not change his/her movements in order to fool such a system, but only research and time will tell.

Fingerprints - As a society, we have been using fingerprints for a systematic, authenticated form of identification for centuries making it the oldest biometric identification method. It has long been known that the pattern of ridges and valleys on the skin of a person's finger is unique to that person. To extract a meaningful measurement from this pattern, systems use what are known as "minutiae". Minutiae is a term used to describe the special characteristics of a fingerprint that are formed when ridges stop, fork or double-back on themselves to form tiny islands. In general, a fingerprint contains about a hundred minutiae though information regarding ten to twelve is all that is needed to form a positive identification.

There are two major categories of fingerprint system attacks -- latent replay and the use of artificial fingerprint prosthetics. The type of sensor used to read the fingerprint for comparison can determine the type of attack that is best used on a system. The two most common types of sensors used are capacitive and optical with a digital camera. Other sensor types include ultrasonic, electric field and temperature [7].

Capacitive sensors read the capacity of a material for storing an electric charge. Because the capacitive values of air and skin are different the sensor can tell when a finger is applied to the sensor and thus needs to take a reading. However, it is not the *specific* capacitive value of skin that causes the sensor to scan, it is merely the presence of something that has a different capacitive value than air. It is in this feature where the weakness lies. If an attacker can cause the sensor to trigger a re-scan of a latent print (left by the legitimate user the last time they used the scanner), more often than not he/she can gain access to the system. To trigger such a re-scan an attacker can cup their hands over the sensor and breathe -- the moisture in the breath condenses on the latent print causing a change in the capacitive value on the surface of the sensor. This change tricks the device into activating a scan of the latent print [5]. Other trigger techniques include placing a bag filled with water on the sensor's surface as well as dusting the latent print with graphite powder then covering the sensor surface with tape and applying pressure. The graphite powder and tape technique also worked for transferring latent prints that

were not actually on the sensor but rather found on items that the legitimate user had come into contact with [5].

Optical sensors are not as easy to trick as capacitive sensors, though they are by no means difficult. What makes optical sensors more challenging is the fact that an actual object must be placed on the sensor for a scan to be triggered. Optical sensors work by capturing the image of a fingerprint that is caused by the illumination the surface of the finger with a light source (usually an LED). The image is captured by a digital camera and processed for comparison. Because an object is needed for a scan to take place, latent replay techniques will not work with optical sensors and the use of prosthetic fingerprints must be employed. The creation of a prosthetic fingerprint can be done both with and without the help of a legitimate user and can be made using widely available materials such as gelatin, liquid silicon rubber, a fingerprinting kit (basically a brush and commercial grade graphite powder), photo equipment and a laser printer. In general, to create an artificial fingerprint with the help of a legitimate user, a mold of the subject's finger is created using molding wax. Liquid gelatin or liquid silicon rubber is then poured into the mold and allow to dry. The resulting form is the prosthetic finger that can be easily worn and disposed of (with gelatin fingers you can eat the evidence!). For fingers created without the help of the user, a latent fingerprint is gathered. A picture of the print is taken and the negative is applied to photosensitive printed circuit board (PCB) where a mold is developed using an etching bath. This mold is then used to create the fake finger. Detailed descriptions of these techniques can be found in works by Ton van der Putte and Jeroen Keuning [7] or Tsutomu Matsumoto, Hiroyuki Matsumoto, Koji Yamada and Satoshi Hoshino [8].

Acceptance rates for prosthetic fingers are in the range of 67%-100% for both capacitive and optical sensors [5,7,8]. Additional security features added by manufacturers such as temperature, blood pressure and heart beat sensors were also tricked at similar or better rates [5,7,8]. The reason for this is twofold. First, these sensors are looking for merely the presence of a temperature difference, blood pressure or heartbeat, not a *specific* measurement. With the use of very thin artificial prints created using silicon, these measurements are allowed to pass through the attacker's live, but unauthorized, finger. Secondly, the sensors must allow for natural variations in temperature, blood pressure and heartbeat due to stress, emotional states, fitness level, etc. Because of this, the range of measurement that is allowed is often more than generous enough to accept an artificial finger made of gelatin or silicon.

Retinal pattern/iris scans – The two types of biometric characteristics concerning the human eye are iris and retinal pattern scans. The iris is the colored, round membrane that surrounds the pupil. Retinal pattern refers to the pattern of blood vessels formed beneath the light-sensitive lining at the back of the inner eyeball. Iris scans are high-quality pictures taken with a digital camera as the iris is illuminated by two fairly weak infrared sources. For retinal scans the process is similar however, the equipment involved is more specialized due to the optics needed to capture a picture of the back of the eye. In addition, the participation of the user to sit for a retinal scan is more involved and precise than that of an iris scan. For these reasons the retinal scan is not as widely used as the iris scan but it is also for these reasons that retinal scans are very difficult to obtain without the help of the legitimate user and no documentation could be found that demonstrated a retinal scan compromise that used a simulation of a retinal scan.

Iris scan systems however, could be fooled using a high-quality image of a person's iris [5]. The trick, of course, is how to obtain a picture of an iris. The documentation reviewed did not go into detail about this procedure but from what was gathered it appears the best way is to obtain an image from the system itself either with or without the help of someone known to the system. Using a high-quality laser printer the image can then be printed out for use. However, presenting the iris to the system is not as simple as holding up the image to the camera. The reason for this is that iris scans are based on relative measurements to the center of the pupil; to gain this starting point the scanner must take an in-depth aperture of the pupil which cannot be done with a flat image [4]. To get around this a small hole is cut in the picture of iris and the attacker then places his/her own eye behind the picture. The system then uses the attacker's pupil as the reference point to continue the scan of the fake iris and grant access to the system.

Ways to Improve the Security of Biometric Systems

It is obvious that biometric systems, like all security systems, have weaknesses. Does this mean that biometric systems should not be used? No. It does mean that you need to treat the use of biometric systems as you would other security measures and use risk analysis, common sense and a layered approach when designing your system. Some thoughts on how to improve the security of a biometric system are:

Combine with other technologies – Do not use a biometric system as your total solution. The marketing and hype surrounding biometrics makes it very easy to believe that it will address all of your security problems. Biometric systems are a security tool to be used in conjunction with other technologies to form a diverse system.

Biometric tokens – For the storage of biometric templates, smart cards are the most secure method. An attacker must have both the biometric characteristic and the smart card for access to the system. This extra layer of security makes this solution more secure than storing templates in a database or on a workstation.

Encryption of data transmission and storage – Add a challenge/response public key infrastructure (PKI) layer to all transmissions and storage of biometric data. Not allowing any clear text representations of biometric data to reside or be transmitted across your system would significantly harden your biometric system. Again, this is another layer of security to add to your biometric system that will deter a good portion of your attackers.

Policy – Not very sexy or newsworthy, but policies work and are easy to customize to your specific needs. In terms of biometrics, examples could include making sure all fingerprint scanning surfaces are kept clean to avoid the replay of latent prints. This may sound silly, but it would have avoided a few of the attacks described above and forced the attacker to try a more advanced technique like a fake finger; something they may not have been willing to do.

Conclusion

This paper has described some very specific weaknesses in biometric systems. However, it should be noted that the attacks above assumed quite a bit of contextual information and were often carried out in a “friendly” environment (e.g., a test lab). The purpose of this paper was to demonstrate that while biometrics can offer greater security for your organization, they are by no means a panacea. Because biometrics offer a new direction in security does not mean that you should throw out more traditional security foundations and techniques.

References:

1) Stobbe, Antje. "Biometrics – hype and reality". *Deutsche Bank Research, Economics, Internet revolution and new economy*, no. 28. 22 May 2002. Available on-line at <http://www.dbresearch.com>.

2) "Fingerprint Identification." Available on-line from the Federal Bureau of Investigations <http://www.fbi.gov/hq/cjisd/ident.pdf>.

3) Jain, A., Hong, L., and Pankanti, S. "Biometric Identification." *Communications of the ACM* 43, no. 2 (February 2002). 91-98.

4) McMillian, Robert. "The Myth of Airport Biometrics." *Wired*, 9 August 2002. Available on-line from <http://www.wired.com/news/conflict/0,2100,54418,00.html>.

5) Thalheim, L., Krissler, J., and Zeigler, P. "Body Check." *c't*, November 2002. Available on-line from <http://www.heise.de/ct/english/02/11/114/>.

6) Dupont, Daniel G. "Seen Before." *Scientific American*, December 1999. Available on-line from <http://www.sciam.com/article.cfm?articleID=000CCD75-3886-1C74-9B81809EC588EF21&pageNumber=1&catID=2>.

7) van der Putte, T. and Keuning, J. "Biometrical Fingerprint Recognition: Don't Get Your Fingers Burned". In IFIP TC8/WG8.8 *Fourth Working Conference on Smart Card Research and Advanced Applications*, 289-303. Kluwer Academic Publishers, 2000. Available on-line from <http://cryptome.org/fake-prints.htm>.

8) Matsumoto, T., Matsumoto, H., Yamada, K., and Hoshino, S. "Impact of Artificial 'Gummy' Fingers on Fingerprint Systems". In *Proceedings of SPIE Vol. #4677, Optical Security and Counterfeit Deterrence Techniques IV, Thursday-Friday 24-25 January 2002*. Available on-line from <http://cryptome.org/gummy.htm>.