Sensible Authen

According to the author of *Beyond Fear*, it's not enough to know who you are; you've got to prove it.

The problem with securing assets and their functionality is that, by definition, you don't want to protect them from everybody. It makes no sense to protect assets from their owners, or from other authorized individuals (including the trusted personnel who maintain the security system). In effect, then, all security systems need to allow people in, even as they keep people out. Designing a security system that accurately *identifies*, *authenticates*, and *authorizes* trusted individuals is highly complex and filled with nuance, but critical to security.

Identification, authentication, and authorization.
Here's the shorthand guide:

- Identification: Who are you?
- Authentication: Prove it.
- Authorization: Here is what you are allowed to do.

BRUCE SCHNEIER,
COUNTERPANE INTERNET
SECURITY

tication

The background of the cover is a complex, abstract composition of glowing yellow and orange lines and patterns. These elements resemble a combination of a circuit board, with its intricate traces and components, and a stylized American flag, with its horizontal stripes. The lines are layered and overlapping, creating a sense of depth and movement. Bright light sources at the bottom left and bottom right create a lens flare effect, illuminating the central elements and adding to the high-tech, digital aesthetic.

Sensible Authentication

The three concepts are closely related, but in a security system it's critical that we tell them apart. Conflating the three—running them together, failing to distinguish each from the others—can lead to serious security problems. In this discussion we'll focus mainly on authentication.

AUTHENTICATION TECHNIQUES

Basically, there are three ways to authenticate an individual: by something the person knows, by something the person has, and by something the person is. All these ways have been used from prehistory until the present day, and they all have different security properties and trade-offs.

The first method is “something the person knows.” Think of passwords, secret handshakes, PIN codes, and combinations to locks. During World War II, American soldiers in Europe would ask strangers cultural questions like “Who won the 1940 World Series?” on the assumption that German soldiers wouldn't know the answer, but every American would.

One of the vulnerabilities of this kind of system is that the verifier learns the secret. Imagine two soldiers meeting in the dark. One asks the other, “What's the password?” The other replies with the password. Now the first soldier knows the password, even if he isn't supposed to know it. The same problem could arise with bank tellers, who could learn the PIN codes of bank customers and then use those codes for nefarious purposes. It is for this reason that many kinds of security systems have computers doing the authentication, not people. Not that computers are infallible, but they're easier to control and less easy to subvert.

The second method to authenticate someone is “something the person has.” The something might be a physical key, a membership card, or a cellphone SIM card. Like the “something the person knows” method, anyone can give this to anyone else. In fact, with either of these methods, all you're really identifying is that the person is of a particular group, not that the individual is a particular person. Knowing the secret handshake authenticates you as a member of the secret society. Having a copy of

a house key authenticates you as one of a group that has access to a given house. I might give you enough information for you to call my bank and withdraw money from my account. When you do this, the bank thinks it is authenticating the account owner, when it is really just making sure that the person on the other end of the phone knows enough information about the account and account owner to be an authorized user of the account.

In ancient Hebrew dialects, the word “shibboleth” means “ear of grain” (or maybe “stream”). According to Judges 12:1–6, the Gileadites defeated the Ephraimites in battle and set up a blockade to slaughter any fleeing Ephraimites. The sentries asked each person to say the word “shibboleth.” Any Gileadites stopped were able to pronounce the word with the *sh* sound. The Ephraimites, who had no *sh* sound in their language, were trapped when they pronounced the word with an *s* sound. Depending on your beliefs about accent and language skills, this story is either an example of “something the person knows” or the third way of authenticating a person: “something the person is.”

More specifically, it's an example of “something the person has that's a physical part of their body.” This is what we normally think of as identification. When we recognize people, we recognize their physical features. On the telephone, we recognize someone's voice. Our ancestors used this type of authentication mechanism even before they evolved into humans. In the animal world, cats spray to mark their territory, dogs sniff each other's butts, and whales have individual songs. More modern versions of this mechanism, called “biometrics,” include fingerprinting, voice printing, hand geometry, iris and retina scans, and handwritten signatures. Ear shape is a *facial* characteristic that's both reasonably distinctive and hard to alter, although it's not necessarily visible on U.S. passport photos. U.S. green cards and German passports require an oblique headshot, showing an ear. People are generally good at recognizing people by biometrics; machines, less so.

Biometrics has an advantage over passwords and tokens in that they can't be forgotten, although they can be lost. (People can lose fingers in an accident, or temporarily lose their voices due to illness.) And biometrics can't be changed. If someone loses a key or an access code, it's easy to change the lock or combination and regain security. But if someone steals your biometric—perhaps by surreptitiously recording your voice or copying the database with your electronic iris scan—you're stuck. Your iris is your iris, period. The problem is, while a biometric might be a unique identifier, it is not a secret. You

leave a fingerprint on everything you touch, and someone can easily photograph your eye.

MULTIPLE TECHNIQUES

Relying on a single authentication technique can be brittle. In the *Odyssey*, Polyphemus the Cyclops captured Odysseus and his men and sealed them in a cave with his sheep. Odysseus poked Polyphemus's single eye out, so when Polyphemus had to let the sheep leave the cave to graze, he could authenticate them only by feel. After watching this process, Odysseus and his men escaped by clinging to the undersides of sheep. Better authentication systems use two or more methods. An ATM, for example, uses "something the person has"—an ATM card—and "something the person knows"—a PIN. (Then it takes the person's picture, for audit purposes.) A passport is a physical card that is hard to counterfeit and contains a photograph. The door-locking device in my company's office uses both a PIN and a hand-geometry scanner.

Credit cards have two forms of authentication—the physical card and a signature—when used in person, but only one when used over the phone: the information on the card. Credit card companies have tried to improve security by requiring merchants to collect the cardholder's address for card-not-present transactions, but telephone and online credit card fraud is still much greater than in-person fraud (15 to 20 cents per \$100, versus 6 cents). Several French banks have recently introduced credit card numbers that are valid only once and are useless if stolen during a transaction, an excellent countermeasure to address the threat. And, for additional authentication, credit cards now have additional digits on the back that are not embossed on the front of the card or on the magnetic stripe.

Many systems perform identification and authentication at the same time. When you recognize a person, you're both identifying and authenticating that individual. When you look at someone's ID, you are both identifying and authenticating that individual. Other systems authenticate and authorize at the same time. A door key is an authentication token, and it also opens the door—in effect authorizing entry.

Systems that confuse identification with authentication can have

significant insecurities. Again and again I'm asked for the last four digits of my Social Security number as an authentication code, even though my Social Security number is a public identification number. I can't change it. I can't prevent others from having it. It's a unique identifier, but it's hardly a secret: a good number to identify me by, but a terrible one to authenticate me by. Your mother's maiden name is a similarly lousy authentication code.

I've described biometrics as an authentication tool, but sometimes they are misused as an identification tool. As authentication systems, biometrics answer a simple question: Does this biometric belong to that person? As a biometric identification system, they must answer the much harder question: Does this biometric belong to anyone in this large database of much-less-reliable biometrics of people? This confusion leads to active failures, and eventually to passive ones.

The reasoning is subtle, so let's work through an example. Automatic face-scanning systems have been proposed for airports and other public gathering places like sports stadiums. The idea is to put cameras at security checkpoints and have automatic face-recognition software continuously scan the crowd for suspected terrorists. When the software identifies a suspect, it alerts the authorities, who swoop down and arrest the miscreant. At the 2001 Super Bowl in Tampa, Florida, cameras were installed, and face-scanning software tried to match the faces of people walking into the stadium with a photo database of people the police wanted to apprehend.

I'll start by creating a wildly optimistic example of the system. Assume that some hypothetical face-scanning software is magically effective (much better than is possible today)—99.9 percent accurate. That is, if someone is a terrorist, there is a 1-in-1,000 chance that the software fails to indicate "terrorist," and if someone is not a terrorist, there is a 1-in-1,000 chance that the software falsely indicates "terrorist." In other words, the defensive-failure rate and the usage-failure rate are both 0.1 percent. Assume additionally that 1 in 10 million stadium attendees, on average, is a known terrorist. (This system won't catch any *unknown* terrorists who are not in the photo database.) Despite the high (99.9 percent) level of accuracy, because of the very small percentage

Sensible Authentication

of terrorists in the general population of stadium attendees, the hypothetical system will generate 10,000 false alarms for every one real terrorist. This would translate to 75 false alarms per Tampa Bay football game and one real terrorist every 133 or so games.

That kind of usage-failure rate renders such a system almost worthless. The face-scanning system needs to interact with another system—a security apparatus that must go into high alert with all its attendant cost, inconvenience, disruption, fear, and panic—and will still come up empty-handed in the end. The guards who use this system will rapidly learn that it's always wrong, and that every alarm from the face-scanning system is a false alarm. Eventually they'll just ignore it. When a real terrorist is flagged by the system, they'll be likely to treat it as just another false alarm. This concept, called the “base rate fallacy” in statistics, applies to medical tests, too. Even very accurate tests can be useless as diagnostic tools if the disease is sufficiently rare among the general population. A 90-percent accurate system, assuming a 1-in-10-million terrorist density, will sound a million false alarms for every real terrorist. And current systems are much less accurate than that; in March 2003, an Australian system was defeated by two Japanese men who simply swapped passports. It's “The Boy Who Cried Wolf” taken to extremes.

It's not just the face recognition software. The system presumes a photo database of terrorists. It seems unlikely that terrorists will pose for crisp, clear photographs. More likely, the photos in the database are grainy ones taken from 1,000 yards five years ago when the individuals looked different. We have to assume that terrorists will disguise themselves with beards, hats, glasses, and plastic surgery to make recognition harder. Automatic face-recognition systems fail miserably under these conditions. And remember, the system I postulated for this example presumes a face-scanning system orders of magnitude more accurate than the ones being sold today. A recent test of an airport system indicated it was less than 50 percent accurate, making it completely useless as an identification system.

Biometric authentication is different. Here the system compares a biometric on file, called the “reference biometric,” with the biometric of the person at the time it is being compared. This reference biometric is not a blurry photograph taken by an undercover spy; it's a known clear picture taken under the best lighting conditions. The person using the biometric system wants to be authenticated by the system and is not likely to make faces, wear dark glasses, turn sideways, or otherwise try to fool the system. And most important, the problem to be solved is different. Instead of answering the question “Who is this random person?” the system has to answer the much easier question: “Is X the person who X claims to be?”

We're far from the day where computers can reliably and independently identify people, but authentication is another matter. By the way, one of the things the U.S. government didn't tell us about the National Guard soldiers staffing the security checkpoints at airports after 9/11 was that they all memorized a small list of faces that they were watching for. That solution is considerably more effective than having computers do it, but it's hardly a long-term solution. ☐

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

This essay is excerpted from Bruce Schneier's latest book, *Beyond Fear: Thinking Sensibly about Security in an Uncertain World* (Copernicus Books, 2003). For more information see <http://www.schneier.com>.

BRUCE SCHNEIER, founder and chief technology officer of Counterpane Internet Security, is an internationally renowned security technologist and author of several well-received books: *Applied Cryptography* (Wiley, 1996), *Secrets and Lies* (Wiley, 2000), and *Beyond Fear: Thinking Sensibly about Security in an Uncertain World* (Copernicus Books, 2003). He also publishes *Crypto-Gram*, a free monthly newsletter about security stories that make the news. Schneier has written op-ed pieces for several major newspapers and has testified on security before the United States Congress.

Reprinted by permission of Copernicus Books.