CHAPTER 5

People Are the Perimeter

> *There's a difference between interest and commitment. When you're interested in doing something, you do it only when circumstances permit. When you're committed to something, you accept no excuses, only results.*
>
> —Art Turock

A few years ago, a senior manager began bringing his corporate laptop into the cafeteria at lunchtime. Typically, he'd find an empty table, set down the laptop, and then walk out of sight to get his lunch. As he perused the salads and main courses, made selections, and paid for his food, his laptop sat unattended in plain view of hundreds of people using the large cafeteria.

My security team noticed the neglected laptop and pointed it out to me. I discussed the issue with the manager a few times, but he continued leaving the laptop unattended. So eventually, I began taking the laptop and leaving my business card in its place.

Not surprisingly, the manager became somewhat annoyed. "Nobody's going to steal the laptop because there are all these people around," he said.

"Okay," I responded. "I'll never take your laptop or complain again on one condition. If you really trust everybody here, you'll take off your wedding ring and leave it on top of the laptop. If you do that, you'll never hear from me again."

He thought about this for a while. Then he said, "You made your point." And he never again left the laptop unattended.

The Shifting Perimeter

This incident helped crystallize in my mind a new perspective about how we should approach information security. It demonstrated how each person's daily decisions can affect the risk dynamics of the company overall.

The traditional enterprise security paradigm, often expressed in castle-and-drawbridge terms, described a wall of technology that isolated and completely protected the workers behind it. To protect our people and information assets, we focused our efforts on fortifying the network perimeter and the physical perimeter of our buildings.

M.W. Harkins, *Managing Risk and Information Security*,
DOI 10.1007/978-1-4842-1455-8_5

Today, however, a growing number of user interactions with the outside world bypass the physical and network perimeters and the security controls these perimeters offer. They take place on external web sites and social networks, on laptops in coffee shops and homes, and on personal devices such as smartphones. As the Internet of Things unfolds, those interactions will also take place on many more "things," such as wearables, cars, and even household appliances.

This changing environment doesn't mean the security perimeter has vanished. Instead, it has shifted to the user. The laptop left unattended in the cafeteria was clearly inside the physical perimeter, but the corporate information it contained was still potentially at risk due to the manager's actions. People have become part of the perimeter. Users' decisions can have as much impact on security as the technical controls we use.

Over the past few years, the idea of the people perimeter has won wider recognition and acceptance. Accordingly, organizations are placing more emphasis on employees' security awareness and behavior.

One reason for this is the rash of high-profile insider exploits, such as the leaks by National Security Agency contractor Edward Snowden. Another is that technical controls have not kept pace with the attackers. Many exploits are reaching users because technical controls, particularly those on endpoint devices, have failed to prevent them. We are therefore more reliant on the user's ability to detect suspicious activity. We also have been forced to deploy more back-end detection and response tools and staff to handle the flow of malware penetrating the corporate infrastructure. These ever-growing security operations teams, which become another layer of the people perimeter, typically are unable to keep up with the flood of malware and commit errors due to "alert fatigue."

There's a continuing emphasis on phishing attacks; the *2015 Data Breach Investigations Report* found that the percentage of users deceived by phishing actually increased from previous years, with 23% opening phishing messages and 11% clicking on attachments (Verizon 2015).

Older social-engineering techniques are also still effective, apparently. At hedge fund Fortelus Capital Management in London, the chief financial officer received an alarming phone call one Friday afternoon. The caller said he was from the company's bank, and warned of possible fraudulent activity on the account. The CFO reluctantly agreed to generate codes enabling the caller to cancel 15 suspicious payments. When he logged into the firm's bank account the following Monday, $1.2 million was gone. The CFO lost his job and was sued by his firm for failing to protect its assets (Chelel 2015).

As almost every company becomes a technology developer as well as a technology consumer, employee security awareness behavior will become an even bigger issue. Security lapses by the employees working on technology-based products can have far-reaching impacts, creating vulnerabilities in the digital services and physical products delivered to millions of customers.

Compliance or Commitment?

Each day, employees make decisions that can affect the company's information risk. Do I leave my computer unattended or not? Do I post this information on social media? Do I install this software on my device? Do I report this suspicious looking e-mail? When I'm in a coffee shop, do I connect to the corporate infrastructure via a secure virtual private network, or do I engage directly over the Internet?

We could view each of these decisions purely in terms of the potential for increased risk. However, there's also a positive side. If users become more aware of security and make better decisions, they can strengthen the organization's defenses by helping identify threats and prevent impact. Among CISOs surveyed recently by best-practices firm Corporate Executive Board, 50% said that insecure behaviors cause more than half of all breaches; but they also said employees are key to uncovering suspicious activity (CEB 2015).

Therefore, as information security professionals, we are in the behavior modification business. Our goals include creating a more security-conscious workforce so that users are more aware of threats and vulnerabilities, and make better security decisions. Furthermore, we need to influence employees' behavior both within the workplace and when they are home or traveling.

If the manager was comfortable leaving his laptop unattended in our cafeteria, would he also leave it unattended at the local coffee shop? At the airport? Or somewhere else where the risk of loss was even greater? My belief is he probably would. When trying to influence this person's behavior, I wanted to achieve more than a level of compliance. I wanted to initiate a feeling of commitment.

The term *compliant behavior* implies making the minimum effort necessary to achieve good performance to a predefined standard. It's like checking boxes on a list of security compliance items. Ultimately, employees feel they are being compelled to follow someone else's list of instructions. Because of this, compliance requires supervision and policing, and employees may sometimes engage in lengthy recreational complaining. If employees are simply following a checklist, what happens when they encounter a situation that's not on the list? They stop and await further instructions, or perhaps they are even unaware of the threat or ignore it.

In contrast, *committed behavior* is intrinsically motivated and self-directed. Being committed implies that people are emotionally impelled to invest in security; they take responsibility and ownership. When people feel committed, they tend to deliver above and beyond the bare minimum. Rather than simply following a predefined list of instructions, they are empowered to make decisions and judgment calls in real time, with a focus on how their actions affect others as well as themselves.

If we can create this sense of commitment in our users, we can implement security not as a wall but as a collective security force that permeates the entire organization. Individually and as a group, every person in the corporation uses their skills in security to protect the organization, handling known attacks today as well as quickly adapting to new threats tomorrow.

When I needed to influence the manager's behavior, I looked for a way to establish this level of commitment. I sought to change the way he felt about the laptop, and to do this I tapped into his emotional connection to his wedding ring.

Creating a culture of self-motivated commitment rather than compliance can make a big difference, as shown in studies by management guru Dov Seidman. His group looked at behavioral differences between businesses with a culture of self-governance, in which an organization's purpose and values inform employee decision-making and behavior, and those with a culture of blind obedience based on command-and-control and coercion. Organizations based on self-governance experienced three times more employee loyalty and half as many incidents of misconduct, compared with organizations based on blind obedience (Seidman 2011).

The implications for enterprise security are clear. As the boundaries between personal and corporate computing dissolve, employees may be accessing information from any location, on any device. If users behave in an insecure way while they are in the office, it's likely they will also exhibit insecure behavior when they're elsewhere. Conversely, if we can create a feeling of commitment that causes them to own responsibility for security, there's a better chance they will behave more securely both within the workplace and when they are outside our physical perimeter. This change in behavior improves the security of the device they are using, the information they are accessing, their personal lives, and the enterprise.

Examining the Risks

Before discussing ways that we can modify user behavior, I'd like to briefly mention some examples of what can happen if we don't influence the ways that users think and act.

As an experiment, the US Department of Homeland Security secretly dropped disks and thumb drives in the parking lots of government and private contractors' buildings. Their goal was to see whether people would pick them up and plug them into their computers. As reported by Bloomberg News (Edwards et al. 2011), up to 60 percent of the people who picked up the items inserted them into their office computers. That number rose to 90 percent if the item included an official-looking logo. Clearly, the security behavior of employees at these facilities left quite a bit to be desired.

Insider threats unfortunately continue to make the news. A former JPMorgan Chase & Co. employee was arrested by the FBI on charges of stealing customer data and trying to sell it to an undercover informant. As noted by *CSO*, similar incidents have occurred multiple times at the bank over the past few years, illustrating the company's inability to account for insider threats despite its substantial annual spending on security technology (Lambert 2015).

Think about what can happen with newer, more sophisticated exploits. A sophisticated attack targeted government departments using fake voice-mails to distract users while malware downloaded in the background. Using social engineering and targeted e-mails, the attackers tricked users into visiting web sites harboring self-extracting archives. The archives contained a recording media file purporting to be a voice-mail from a female journalist seeking information for a news story, alongside other files that downloaded malicious content (CNET 2015).

As in the example above, today's threats may arrive in the form of carefully personalized spearphishing communications designed to win the trust of targeted users. These users then unwittingly provide access to the information the attackers want. In essence, trust—in this case, the organization's trust in the user—has become the attack surface.

Let's say a company is looking to hire a credit analyst with a very specific set of skills. Attackers notice this and apply online, using a résumé that lists the exact skills required for the job and contains the terms the company's résumé-scanning software is likely to be looking for. Suitably impressed, the company's human-resources specialists forward the application to the company's credit-department manager, who has access to all the systems storing customer financial data. The manager trusts this communication because it has been sent from another department within the same company. So she clicks on the link to the résumé. Unfortunately, that action triggers the execution of malicious code. The human-resources team effectively acted as an infection agent, ensuring the attack reached its real target.

Social media accounts can become sources of risk even when they haven't been compromised. There have been several examples in which senior executives accidentally revealed information that was confidential or problematic in other ways. In November 2014, Twitter's CFO accidentally publicly tweeted a plan to buy another company, including the fact that he wanted help to make the deal happen at a meeting the following month. The CFO was apparently trying to send the message privately (Frier 2014).

At Houston-based fashion retailer Francesca's Holdings, a former CFO frequently shared his thoughts via a personal blog, Facebook page, and Twitter feed (Silverman 2012). Unfortunately, he also shared information that caused problems for his employer. The company fired him because he "improperly communicated company information through social media."

Users frequently post information on external social-media sites that attracts the attention of competitors or the media. To boost their job prospects, interns mention product features they helped develop during their summer job at a well-known company; sales representatives reveal the names of major clients; even senior executives have been known to unintentionally disclose key corporate strategies. In fact, services exist that specialize in aggregating apparently minor snippets of information from social-media and other web sites to build an accurate view of a company's size, geographical distribution, and business strategy, including hiring patterns that indicate whether the company is expanding and which new areas it is moving into.

Adjusting Behavior

To counter these risks, we need to make employees aware and empowered so they act as an effective part of the security perimeter. Increasing recognition of this need has led to the development of a small ecosystem focused on increasing security awareness and changing behavior, ranging from companies offering best practices to groups focused on internet safety for children. This includes companies that train users to avoid phishing exploits, using simulated phishing scenarios and other tools. Security awareness professionals have come together to share best practices (see sidebar).

While I was Intel's CISO and then Chief Security and Privacy Officer, we focused on building security and privacy protection into the corporate culture, getting employees to own responsibility for protecting enterprise and personal information. Achieving this required a lot of effort, and we realized that it took just as much work to maintain a culture of security and privacy as to build it.

Training is a key part of security efforts at most companies, and Intel is no exception. We supplemented general training, which fulfills most legal requirements, with specialized training for employees who have specific roles or access to sensitive information. Another effective technique was to embed security and privacy training into business processes. When an employee requested access to an application that handles sensitive information, they were automatically prompted to take training that focused on the related security and privacy concerns. We also used online training including video and other visually stimulating material as well as entertaining, interactive tools to help engage users (see Figure 5-1).

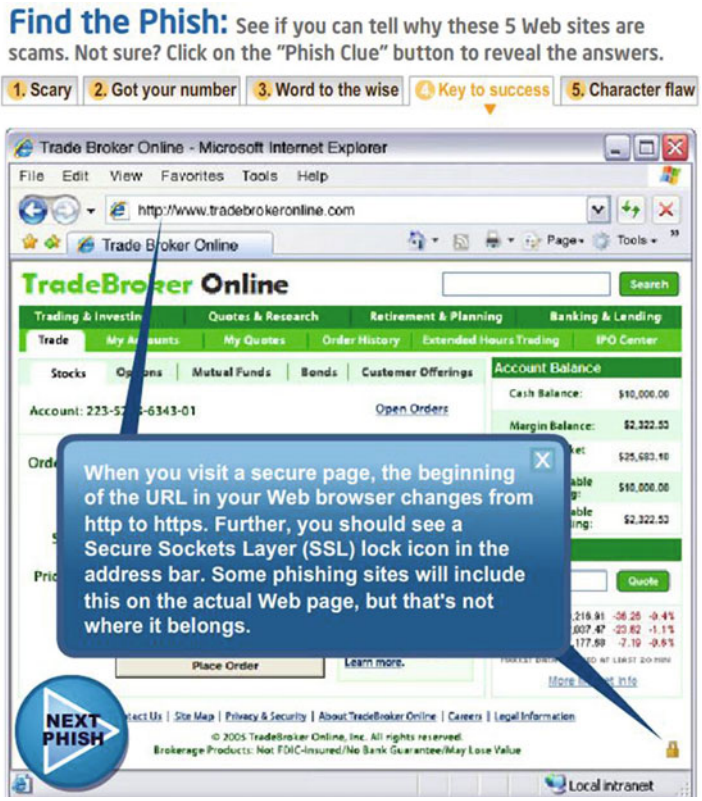

Figure 5-1. *Intel's internal "Find the Phish" interactive training tool helps employees spot web scams. Source: Intel Corporation, 2012*

However, it is not enough to create good training. If nobody takes the training, the effort is wasted. We found that incentives such as public recognition helped ensure employees underwent training and absorbed the lessons. Ultimately, if people continued to avoid security training, we escalated compliance efforts by directly contacting them and their managers.

We also found we could help maintain and increase awareness by publishing security-related articles on Intel's primary employee portal. Many of these articles included a personal aspect, such as preventing identity theft, keeping children safe online, and home wireless security tips. The focus on personal concerns recognized that the way employees behave outside the office is as important to enterprise security as their behavior in the office.

How did we know our security efforts paid off? We accumulated a variety of evidence, including independent benchmark results from Corporate Executive Board (2011), which indicated that Intel employees consistently ranked in the top 10 percent of companies for secure behavior. We also observed that employees acted as part of the security perimeter by alerting us to suspicious text messages or e-mails they'd received.

A Model for Improving Security Awareness

Some of the industry's most valuable work to advance organizational understanding of security awareness comes from best-practices firm CEB. The company offers a security awareness service that includes surveying key security-related employee behaviors at member organizations, and benchmarking the results against other organizations. The program attempts to understand the psychological reasons for insecure behavior; it then focuses on those psychological drivers when suggesting tactics to change employees' behavior. To date, CEB has collected some 300,000 employee responses from over 400 organizations.

The program has found that despite the importance of secure employee behavior, most organizations deliver only moderate amounts of training: just over an hour per year, on average, and only three to four employee communications per month. Survey results suggest that organizations can increase training time to as much as six hours a year before experiencing diminishing returns.

CEB emphasizes that organizations need to use an understanding of psychology to tailor their awareness efforts; awareness programs must target the specific root causes of employees' risky behavior in order to be effective. The company initially identified five psychological factors influencing security awareness and behavior: lack of knowledge of policy, lack of self-interest in security, inadequate perception of the risk to the organization, a low emotional commitment to security, and a perception that secure behavior imposes a high burden. It recently added a sixth factor: the ability to display good judgment.

CEB's findings suggest that the perceived burden of secure activities affects employee behavior more than any other psychological driver. That's the bad news. The good news is organizations can fix the perception of the burden both by reducing the burden itself and by addressing the other drivers. For example, employees' emotional commitment to security increases if their managers engage with them directly to emphasize the risks. Clear enforcement of policy compliance increases employees' self-interest in secure behavior and heightens their perception of risks.

Based on survey data collected over the years, CEB has developed a model that organizations can use to help plan and assess their security awareness programs. The model presents a four–stage progression toward higher security awareness and commitment, from basic check-the-box compliance to active involvement in security. It recognizes that in a complex threat environment, we can no longer rely only on policies that prescribe specific employee behaviors: we also need to enable employees to actively support security activities including breach detection.

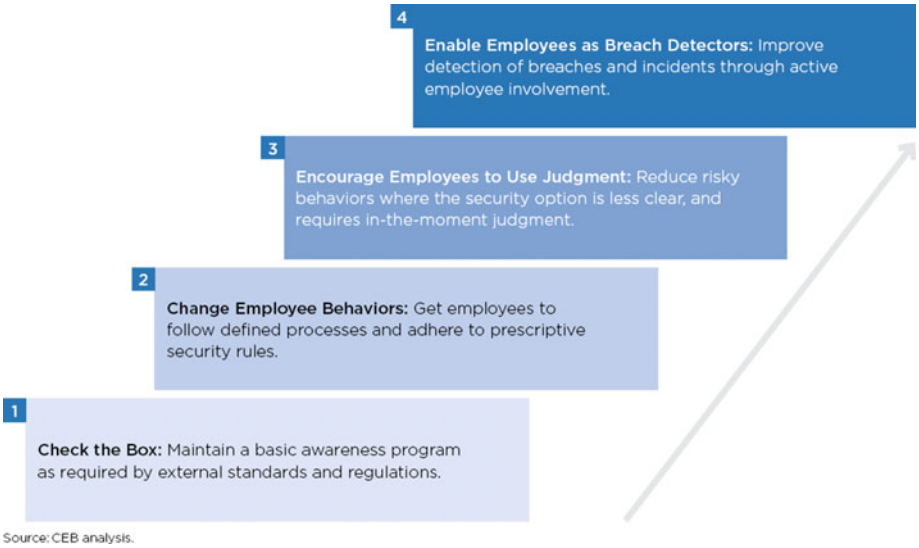

Figure 5-2. *A four-stage model for programs seeking to improve security awareness and behavior. Source: CEB Inc., 2015*

Employee awareness programs at Level 1 (Check the Box) simply respond to external regulatory requirements and don't explicitly aim to change specific employee behaviors. At Level 2, programs try to encourage users to adopt specific, simple behaviors, such as avoiding sharing passwords and sending sensitive information to their own personal e-mail addresses.

The third and fourth levels display greater levels of judgment and commitment. At Level 3, employees are able to make good judgment calls in the moment, especially in situations where the right answer is not immediately obvious. For example, they remember to pause before clicking to check whether an e-mail comes from a legitimate source or contains a phishing link. At Level 4, employees become an extension of the information security organization; they not only avoid security risks, but also notify information security when they see something suspicious. A key behavior here is an increase in reporting events such as spearphishing attempts.

The encouraging news is that that CEB surveys show a gradual improvement; the percentage of employees avoiding insecure behaviors such as password sharing has slowly increased over the past six years. Resistance to phishing has improved faster, though from a lower base. "I think at progressive companies the aspiration is changing," says CEB practice leader Jeremy Bergsman. "Most companies have been moving from Level 1 to Level 2 over several years, and are starting to think about Level 3. But progressive companies are moving beyond employee behavior as a risk to be reduced, and working on ways to make employees a control—an early warning system (Level 4)."

For example, a large telecommunications firm that participated in the CEB security awareness program wanted to empower employees to act as controls supporting information security. It provided each employee with a weekly report tracking his or her behavior, including the documents they accessed, and the devices and external locations used to connect to corporate systems (Figure 5-3). Employees were responsible for reading the reports, thus sharing responsibility with information security for detecting

breaches. The firm found that employees detected suspicious activities faster than would have otherwise been the case; users also proactively improved security by suggesting other activities that should be tracked and added to future reports.

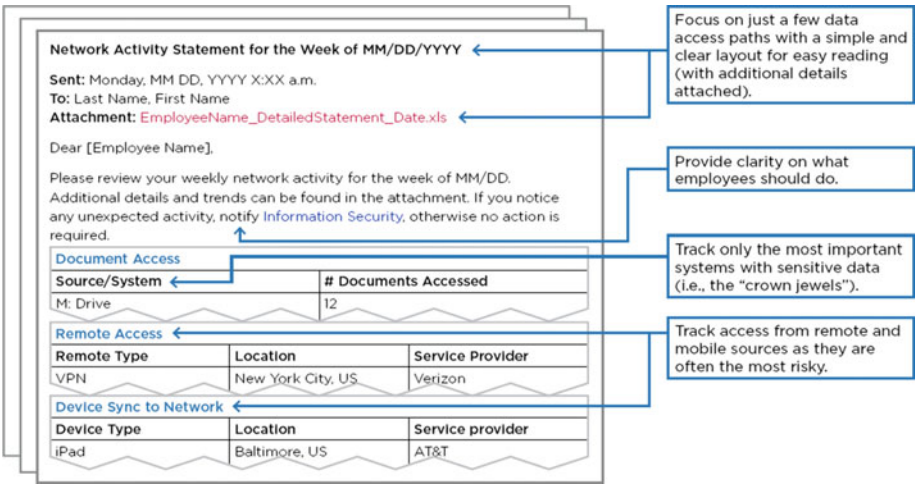

Figure 5-3. *Example of weekly tracking report showing employees their activity. Source: CEB Inc., 2015*

INTERNATIONAL ASSOCIATION OF SECURITY AWARENESS PROFESSIONALS

Exchanging ideas with other security professionals can help improve security awareness programs. The International Association of Security Awareness Professionals (www.iasapgroup.org) is an independent association of corporate security specialists who are seeking to do just that. IASAP is a non-profit, fee-based association dedicated solely to security awareness programs. Its goal is to serve as a trusted forum of security awareness professionals collaborating to improve employee security behavior. "Clear guidance has not been as available for employee security behaviors as it has been for technology solutions," says IASAP board member Michael Diamond. "Several awareness professionals noticed this gap, and that ultimately led to the formation of IASAP."

Some members have built programs from scratch; others inherited established programs in need of fresh ideas and new energy. Members meet in person two to three times per year to learn about other members' security programs and present their own. There's also a members-only sharing platform supporting Q&A, feedback, benchmarking surveys, member polls, guest speaker webinars, and teleconferences. Some members feel comfortable posting program resources that are available to other members for re-branding within their own programs.

Broadening the Awareness Model

I think that the CEB four-stage model shown in Figure 5-2 is a very useful tool. One limitation, in my opinion, is that the model is based on the traditional organizational view of information security. I believe that we need to expand the model to capture a more complete view of information risk. You might think of the following two additions as Levels 3a and 4a, respectively, of the model.

- Privacy awareness, which has become a critical concern. Just consider the number of breaches that have targeted personal information at retail, healthcare, and government organizations.
- A specific focus on engineers and other technology professionals, including those creating technology-based products and services. If engineers don't have a foundational understanding of privacy and security, they cannot design privacy and security into the technology they produce. As a result, a company's products may contain vulnerabilities that introduce significant risks for the business and its customers.

The Security Benefits of Personal Use

Employees use an ever-growing variety of personal devices every day, both inside and outside the physical workplace. This trend started with smartphones and laptops; it also includes wearable devices such as smartwatches and fitness monitors. Information security specialists naturally tend to focus on the security risks of using these devices for business purposes. As I discussed earlier in the book, I've found that the productivity benefits of personal devices often outweigh the risks. But even the security implications are not as one-sided as they might seem at first glance. I believe that, in some respects, allowing personal use may actually encourage better security.

In general, people are likely to take better care of their own possessions than someone else's. They feel a stronger connection to their own car than to one provided by their employer. If people are using their own computing device, they may take better precautions against theft or loss. And they may feel the same way if they are storing personal information on a corporate device. At Intel, we allowed reasonable personal use of corporate laptops, and therefore many employees stored personal as well as corporate information on their laptops. Because of this, they had a personal stake in ensuring the devices didn't get lost or stolen. I believe this sense of ownership contributed to our lower-than-average laptop loss rates.

Another company's experience provided some empirical evidence supporting this idea. The company conducted a tablet pilot deployment in which, for the first time, it allowed personal use of corporate devices. The company found that breakage and loss rates were dramatically reduced compared to its past experience with mobile devices. The CIO's conclusion was that employees simply take better care of devices when they use them for personal purposes.

Perhaps we should be similarly open-minded when considering the security implications of wearable devices. I met with managers at a large company who were pondering the security implications of smartwatches and fitness monitors, which

employees were already bringing into the workplace. Understandably, some people at the company wanted to make sure the devices could not interact with the corporate network. I observed that in the future, wearables could be harnessed to help identify users in ways that are less cumbersome for users than traditional controls such as passwords. Fitness devices, including some smartwatches, count the user's steps and monitor heart rate, and could therefore be used as biometric security devices in the future. As the devices evolve and accumulate more user data over time, they may become more adept at identifying each user's physiological and behavioral "signature." In addition, some smartphones include fingerprint recognition, which in itself can be a powerful authentication mechanism if the technology has been properly designed and implemented.

As security professionals, shouldn't we think about taking advantage of the benefits these technologies offer? We should seek to integrate into security strategies the broader variety of existing devices, which have useful features such as cameras and voice recognition and also contain data about our use patterns. Many of these devices already communicate with each other; why not take the next step and use the technology to eliminate the pain of using passwords? Why not find a way to reduce risk and cost, while providing a much better user experience, by using these devices to authenticate us automatically?

It may also be worthwhile to reexamine other assumptions about the security implications of personal devices. Some companies have policies forbidding the use of cameras in their offices. However, a smartphone includes a camera that employees can use to capture the off-the-cuff design sketches often scrawled on whiteboards during brainstorming sessions. This intellectual property can then be stored and encrypted on a hard drive within the enterprise. Is it safer to allow employees to photograph the image, or to copy it onto a piece of paper, or to leave it on the whiteboard where anyone might see it? Companies may come to different conclusions, depending on their culture and appetite for risk. But this is another illustration of the importance of considering all the possible business benefits as well as the risks when making technology decisions.

Roundabouts and Stop Signs

To try to reduce driving accidents at a dangerous curve in Chicago, the city painted a series of white lines across the road. As drivers approached the sharpest point of the curve, the spacing between the lines progressively decreased, giving the drivers the illusion they were speeding up, and nudging them to tap their brakes. The result was a 36 percent drop in crashes, as described by Richard Thaler and Cass Sunstein in their book *Nudge* (Yale University Press, 2008).

This traffic-control method succeeded in making drivers more aware and improving safety while keeping the traffic flowing with minimum disruption. I think this example provides a useful metaphor for information security. Some security controls are like stop signs or barriers: we simply block access to technology or data. But if we can shape the behavior of employees rather than blocking them altogether, we'll allow employees, and therefore the company, to move faster.

To use another traffic metaphor, a roundabout at an intersection typically results in more efficient traffic flow than an intersection with stop signs, because drivers don't have to come to a complete halt. The roundabout increases drivers' awareness, but they can proceed without stopping if the way is clear. Statistics have shown roundabouts are often safer than intersections.

Of course, we need to block access in some situations, such as with illegal web sites. But there are cases where it's more efficient and productive to make users aware of the risks, yet leave them empowered to make the decisions themselves.

Consider the case of a large multinational company whose business relied heavily on its significant intellectual property. To protect that proprietary IP, the company implemented data-loss protection software, including an application on employees' laptops. But instead of simply blocking transmission of information flagged as sensitive, the company configured the software to warn employees whenever it detected potentially insecure behavior. If an employee tried to transmit a confidential document, the software displayed a message that explained the potential risks and suggested ways to protect the information, such as encryption. After all, users may have good reasons for sending confidential documents, and preventing transmission could be detrimental rather than beneficial to the business. The company found that this warning caused 70% of users to change their behavior, representing a major reduction in risk. Yet because of the way the software was configured, users didn't complain about the security burden. The roundabout approach reduced risk without interfering with users' productivity.

Here's another hypothetical example. It may make sense to warn users visiting certain countries that they may be accessing material that is considered unacceptable. A US employee traveling on business might be working in a local office of a country with strict religious guidelines. The employee has a daughter who's in a beauty pageant, so it would be natural to check the pageant web site from time to time. But the images could be offensive in the country, so it makes sense to warn the employee to exercise caution. At Intel, we found that when we warn users in this way about potentially hazardous sites, the vast majority heed the warnings and don't access the web sites.

In the case of information security, there's an additional benefit of making controls as streamlined as possible. We all know if controls are too cumbersome or unreasonable, users may simply find ways around them. We kept this concern in mind when developing a social media strategy at Intel IT (Buczek and Harkins 2009). We were well aware of the risks associated with social media, but attempting to stop the use of external social media web sites would have been counterproductive and, in any case, impossible. We realized that if we did not embrace social media and define ways to use it, we would lose the opportunity to shape employee behavior.

As part of our initial investigation, we conducted a social media risk assessment. We found social media does not create new risks, but can increase existing ones. For example, there's always been a risk that information can be sent to inappropriate people outside the organization. However, posting the same information on a blog or forum increases the risk by immediately exposing the information to a much wider audience.

So we developed a social media strategy that included several key elements. We determined that we could reduce risk by implementing social media tools within the organization, so we deployed internal capabilities such as wikis, forums, and blogs. Initially, employees used these tools mainly to connect socially rather than for core business functions; we later integrated the tools into line-of-business applications to achieve project and business goals. We also worked with Intel's human-resources groups to develop guidelines for employee participation in external social media sites, and developed an instructional video that was posted on a public video-sharing site. The video candidly explained that Intel wanted to use social media to open communications channels with customers, partners, and influencers, to encourage people to adopt the technology, and to close the feedback loop. The information also included guidance

about how to create successful content and general usage guidelines. We also used technology to help ensure that employees followed the guidelines. We monitored the Internet for posts containing information that could expose us to risks, and we also monitored internal social media sites to detect exposure of sensitive information and violations of workplace ethics or privacy.

The Technology Professional

So far, I've focused mainly on the security roles of end users. But think about the broadening roles that technology professionals play at many companies. Historically, technology professionals have performed back-office IT roles at most companies, such as managing infrastructure and internal applications. Many also work on web sites and online services. We're now moving into a future in which companies in all industries will become creators of technology embedded in physical as well as digital products, and they'll hire developers to create that technology. These technical professionals are also part of the people perimeter, and their actions can have major positive or negative effects.

We've already seen several well-publicized problems caused by vulnerabilities in products. Fiat Chrysler recalled Jeeps in 2015 after researchers showed they could hack into a 2014 model and hijack its steering, brakes, and transmission. The researchers used an unsecured communications port to execute the attack (Dark Reading 2015). Similar concerns prompted the FDA to order organizations to stop using older drug infusion pumps made by Hospira when it was found that an unauthorized user could hack into them and change the dosage the pump delivers.

In traditional IT roles, technical professionals manage almost every element of the technology spanning our networks, data centers, and users' computing devices. They develop and install software. They configure, administer, and monitor systems. Their actions or inaction can make the difference between a system that is vulnerable and one that is reasonably secure.

Those systems include servers, which are still the IT assets most commonly attacked and robbed of data. An attacker may initially gain access to your company by compromising a user's laptop, but the biggest prize—databases of corporate intellectual property and personal information—still reside on the enterprise servers. To steal that information, attackers now typically often use a compromised end-user device to search the network for servers with inadequately configured access controls. Surveys show many attacks continue to exploit security holes that organizations could easily have fixed. Among organizations surveyed for the 2015 *Data Breach Investigations Report*, more than 30 of the exploited vulnerabilities had been identified as long ago as 1999, yet presumably not addressed at the victim organization. As the report notes, "Apparently, hackers really do still party like it's 1999."

Similar trends can be seen in the incidence of software errors. Many of the most serious, frequently exploited vulnerabilities in software are due to well-known errors that are "often easy to find, and easy to exploit," as noted in the 2011 CWE/SANS Top 25 Most Dangerous Software Errors (CWE/SANS 2011). Furthermore, the situation does not seem to be improving. As David Rice, author of *Geekonomics* (Addison-Wesley Professional 2007), puts it, most software is not sufficiently engineered to fulfill its designated role as the foundation for our products, services, and infrastructure (Rice 2007). This is partly due to the fact that incentives to improve quality are "missing, ineffectual, or even distorted," he concluded. To compete, suppliers focus on bringing products to market

faster and adding new features, rather than on improving quality. Rice estimated, based on government data, that "bad" error-ridden software cost the United States a staggering USD 180 billion even back in 2007.

Not surprisingly, the typical recommendations for improving IT security often sound remarkably familiar. That's because they address problems already known to most organizations, but not fully addressed. As the Data Breach Investigations Report notes, the question is not which vulnerabilities should be patched (all of them should): "The real decision is whether a given vulnerability should be patched more quickly than your normal cycle or if it can just be pushed with the rest." Previous editions of the report have recommended basic precautions such as ensuring passwords are unique; regularly reviewing user accounts to ensure they are valid and properly configured; securing remote access; increasing employee awareness using methods such as training; and application testing and code review to prevent exploits such as SQL injection attacks and cross-site scripting, which take advantage of common software errors.

The fact that these measures do not appear to be rigorously applied at many organizations takes us back to a key theme of this chapter: that the commitment of employees is as important as the policies and procedures you have in place. If administrators and developers are committed rather than just following directives, if they feel personally responsible for the security of the enterprise, and they will be more conscientious about ensuring the right technical controls are in place.

Insider Threats

High-profile national security breaches by insiders such as Edward Snowden and Chelsea Manning have made insider threats a considerably more prominent issue during the three years since the first edition of this book was published.

Among the 557 organizations participating in the 2014 Cybersecurity Watch Survey (CSO et al. 2014), 28 percent of cybercrime events were attributed to insiders. Furthermore, insiders accounted for the highest percentage of incidents in which sensitive or confidential information was stolen or unintentionally exposed.

Insider attacks also cause additional harm that can be hard to quantify and recoup, such as damage to an organization's reputation. Insiders have a significant advantage because they can bypass physical and technical security measures such as firewalls and intrusion detection systems that were designed to prevent unauthorized access. The organization's trust in the insider is used as the attack surface. In at least one case, the insider was the person one might least suspect: the head of information security at the Iowa state lottery, who hacked his employer's computer system, and rigged the lottery so he could buy a winning ticket in a subsequent draw. By installing a rootkit on a lottery system, he could secretly alter the lottery's random number generator, enabling him to calculate winning numbers in advance and buy a winning ticket in advance (Thomson 2015).

Unfortunately, even security firms are not immune to compromise; well-known cybersecurity company FireEye hired an intern who was later discovered to be a top Android malware developer. Unfortunately, his job at the security firm involved researching and analyzing Android malware, which raises the concern that he could have used his inside knowledge to develop malware capable of evading technical controls (Fox-Brewster 2015).

Yet surveys have also suggested that many insider attacks are opportunistic, rather than highly planned affairs. Many insiders take data after they've already accepted a job offer from a competitor or another company, and steal data to which they already have authorized access. In some cases, misguided employees may simply feel they're entitled to take information related to their job.

Clearly, all organizations need to be aware of the insider threat. It may not be possible to thwart all insider exploits, but we can take actions to reduce their likelihood and impact. Perhaps the biggest step we can take is to instill a culture of commitment. User behavior analytics technology can also help by detecting behaviors or access privileges that are outside the norm; perhaps technology could have prevented the case in which a former nursing assistant at an Orlando health network inappropriately accessed about 3,200 patient medical records, with no apparent motive. Besides disclosing the breach, the health network had to notify the affected patients and offer support, fire the employee, reeducate the workforce, and increase its efforts to audit and monitor access (Brinkmann 2015).

To help manage insider threats, consider a three-part approach: deter, detect, and discipline. Remember that successful implementation will require the involvement of the entire organization.

Deter

- Build security awareness and instill a culture of commitment, using the techniques discussed in this chapter.
- Make your company a great place to work. Employees are less likely to get disgruntled, and therefore less likely to seek ways to harm the company.
- Let people know you're watching. Technology can help monitor users' activity. Showing users their activity reports can help involve them in protecting the business. It also lets potentially malicious insiders know they're being watched.

Detect

- See something, say something. A committed workforce will tell you if they see something suspicious.
- User behavior analytics tools are becoming more and more effective at finding anomalies in access permissions and user activity, and identifying whether a user's actions are far enough outside the norm that they merit investigation.
- Form a team that focuses on insider threats and investigations. This should operate as a cross-functional team with involvement from human resources, legal, physical security, and information security groups.

Discipline

When an insider incident occurs:

- If it's an honest mistake without a big impact, immediate remedial training may be the best remedy.
- If the impact was low, and the incident seems more an error of judgment than a malicious act, a less heavy-handed approach may be appropriate—perhaps a written warning or a comment in the person's performance review.
- If the intent is clearly malicious, or the impact is significant, consider the options of termination and even engaging law enforcement.

Finding the Balance

One reason that organizations are focusing more attention on security awareness is that their technical controls have failed to prevent attacks from reaching employees and thus the core of the enterprise. Rapidly evolving new exploits, often involving social engineering as well as malware, have outstripped the capabilities of the security tools companies have relied on in the past.

Now, innovative security technology is becoming available that uses machine learning and artificial intelligence techniques to prevent malware much more effectively, on every type of device. This is great news for all consumers of technology. The adoption of this technology should result in a substantial reduction in risk, due to a precipitous drop in malware. The danger is that some will see this as an opportunity to dial back their security awareness efforts. I think this could be a mistake. We will always need to maintain a level of diligence and discipline in security and privacy awareness. However, we may be able to shift the emphasis of training toward prevention and future risks, and focus on how we should design, develop, and deploy technology that better protects privacy and resists attacks.

No matter how good our technical controls are, we will still need people to act as part of the perimeter. We need to create a sense of personal commitment and security as well as privacy ownership among our employees. If we succeed in this goal, we will empower employees to help protect the enterprise by making better security decisions both within and outside the workplace.