CHAPTER 6

Emerging Threats and Vulnerabilities

Reality and Rhetoric

Curiosity is lying in wait for every secret.

—Ralph Waldo Emerson

One day, it's hard to read an online news source, pick up a newspaper, or watch TV without seeing reports of new threats: cybercrimes, data breaches, industrial espionage, and potential destruction of national infrastructure. These reports inevitably leave the impression that we are drowning in an inexorable tide of new and terrifying threats.

One has to question how much of this is rhetoric, and how much is reality. There are political and profit-driven motives for making threats seem bigger and more imminent than they really are. US government officials have warned that cyber attacks potentially can be "devastating, approaching weapons of mass destruction in their effects" (Levin 2010). Such warnings have been used to justify requests for increased national cybersecurity funding, as well as proposed restrictions on private networks. It's not surprising, therefore, that some experts have expressed skepticism about the real extent of the threat. In fact, academics at the George Mason University Mercatus Center have warned, "the United States may be witnessing a bout of threat inflation similar to that seen in the run-up to the Iraq War" (Brito and Watkins 2012).

On the other hand, common sense tells us new cyber threats really are emerging and growing. Malware production has matured into a sizable industry. More data is online and vulnerable to attack, and millions of new Internet-connected devices are inevitably introducing new risks.

Given the flood of often-conflicting information, how can we get an accurate picture of the threat landscape so that we can develop an appropriate security strategy? How do we determine which threats directly affect our organizations, and distinguish them from those that are irrelevant? How do we decide which threats require immediate defensive measures, as opposed to those that attract attention but don't yet present significant risks?

In this chapter, I'll describe methods for identifying the real threat and vulnerability trends among the rhetoric. I'll also discuss some key areas of threat activity that have been analyzed using these methods. My goal is to help information security groups stay ahead of the attackers and focus their limited resources on mitigating the most important threats.

Structured Methods for Identifying Threat Trends

To identify the real trends in emerging threats among the mass of news and speculation, we need to carefully examine the available information using a structured, analytical approach. Unfortunately, many security groups absorb information about emerging threats using methods that are unstructured and sometimes almost haphazard.

A typical process looks something like this. The security team relies on external sources, such as news feeds and alerts, as well as informal anecdotes, to gather information about emerging threats. Based on this information, the team holds brainstorming sessions to review the threat landscape. The output from these sessions is a list of "top risks." Security resources are then focused on mitigating the items on the list.

There are several problems with this approach. Information comes from a narrow, limited range of sources, resulting in a blinkered security perspective that tends to stifle creative thinking. Also, the information is usually fragmented, making it difficult for the team to identify trends and gaps in the data. These deficiencies continue through security planning and implementation. Because the team lacks a full view of the threat landscape, it's hard to determine which threats require immediate attention and how much of the limited security budget they deserve. As a result, risks are incorporated into plans on an ad hoc basis, and not all risks are adequately mitigated. Finally, security teams often don't have a structured process for communicating threat information to other people within their organizations. Because of this, people outside the security group remain unaware of emerging risks and don't know how to respond when they experience an attack.

At Intel, we realized the limitations of this approach several years ago and began trying to inject more rigor into our risk-sensing strategy. Over time, we've progressively developed a more structured risk-sensing process that helps us identify threats, prioritize them, plan our response, and deliver actionable information to other groups across the company. Through continued use, risk sensing has become a systemic process within Intel.

Our process for analyzing emerging threats includes several valuable techniques that may be unfamiliar to security groups at most other organizations. We use a product life cycle analogy to track threats as they mature from theoretical risks into full-blown exploits. We also use nontraditional analysis techniques, such as war games and threat agent profiles, to encourage creative thinking and identify threats we might otherwise miss. I'll discuss these methods in more detail later in this chapter.

The process is managed by a small core team, supplemented by a broad set of experts across Intel. This arrangement ensures continuity while enabling the team to mine a diverse variety of sources to get a more complete picture of immediate and future threats.

Security team members research a wide range of individual security topics in depth. Besides using typical sources, such as external feeds and analysis, they mine academic research and hacker discussion forums, and they network with other security professionals. Other team members scan the regulatory horizon to identify upcoming laws and regulations that may impact us. We also analyze internal investigations and other near-miss incident data. Team members communicate with each other frequently to identify areas of potential overlap.

We then hold regular meetings to analyze the threat landscape. Each security domain expert explains their findings to other members of Intel's security community. For each security topic, we review recent events and then look ahead to the future. Reflecting on what has happened helps us identify the key trends and the factors driving those trends, and it provides context that we use to analyze the current state. We then look ahead to predict the likely evolution of each threat based on the trends we've identified.

This structured evaluation uncovers emerging risks we wouldn't otherwise see. We also look back at our previous predictions to see which ones were accurate, and to analyze the reasons why threats may not have materialized in the way we expected.

We communicate our findings to stakeholders across Intel in regular reports and briefings, including a wide-ranging annual assessment of the threat landscape. This communication provides further opportunities to get feedback from across Intel's business, which we can use to refine our risk-sensing analysis.

The Product Life Cycle Model

We have found that a product life cycle model is a useful way to track and prioritize emerging threats as they evolve and begin to present real risks to the enterprise. Like all security groups, we have a limited budget, and we need to direct our resources to mitigate the highest-priority threats.

This model, shown in Figure 6-1, recognizes that many threats initially emerge as theoretical risks, but are on a path to exploitation, and we need to evaluate and monitor them.

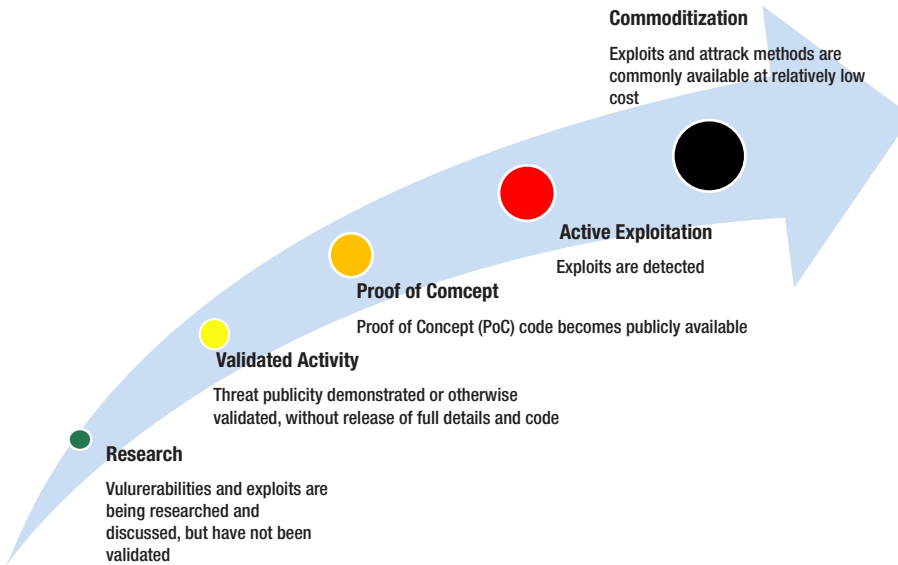

Figure 6-1. *The product life cycle model for tracking the evolution of threats. Source: Intel Corporation, 2012*

Often, researchers or hackers first reveal a possible attack or vulnerability at a security conference or publish information about it online. Next, attackers begin testing the use of this technique and making this information publicly available. Once the method has been proven, the threat enters the production phase as attackers start exploiting it in earnest. Ultimately, the threat becomes a mature commodity—source code is often freely available, many variants exist, and organizations treat the threat as part of the everyday landscape and build defenses accordingly.

This life cycle model enables us to systematically track the evolution of threats. It helps us determine when we need to allocate resources to fighting each threat. As each threat approaches maturity, we can examine how it is likely to affect us and plan appropriate mitigation.

In addition, at a product manufacturing company like Intel, this model provides a great way to communicate actionable information to business groups using terminology they understand—the product life cycle. When we provide our regular threat landscape assessments to stakeholders, each security topic includes a description of activity at each life cycle phase, thus providing a context that helps business groups across Intel determine how they should act on each of these emerging risks.

Let's examine some examples showing how we use this model in real life. Figure 6-2 illustrates the evolution of threats targeting smartphones and other handheld devices. Researchers and hackers began to take notice of handheld devices almost a decade ago, demonstrating weaknesses and theoretical avenues of exploitation. Initially, they focused on what were then known as personal digital assistants. As smartphones took off, attackers shifted their attention to this bigger market, which rapidly became a major area of threat activity. Monitoring this trend enabled us to prepare internally and inform Intel product development groups. As the threats matured and employees began using smartphones more widely at work, we then developed risk mitigation measures including technical controls and incident response plans.

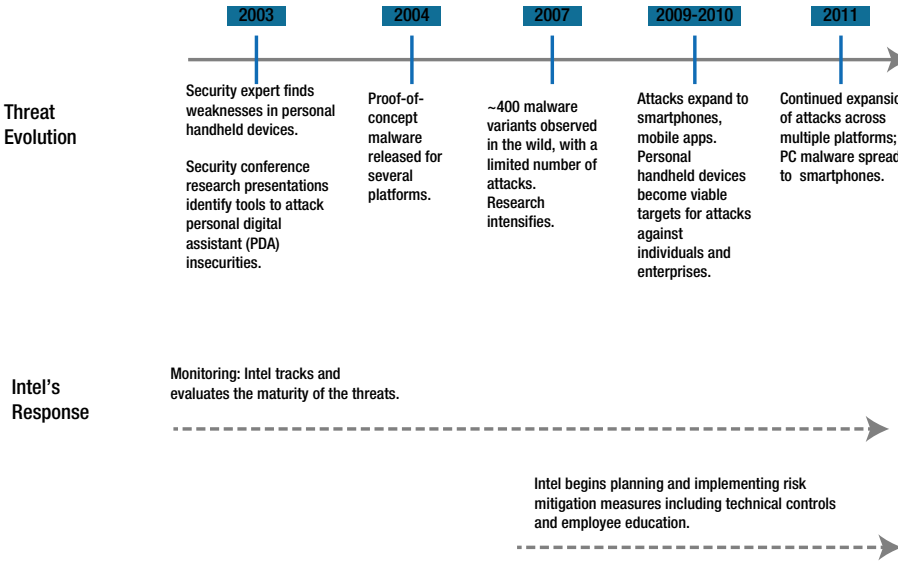

Figure 6-2. *How Intel used the product life cycle model to track and respond to smartphone security threats. Source: Intel Corporation, 2012*

By visually comparing activity across multiple threat areas, we can quickly identify major areas of activity and see the likely timing and extent of their impact. This chart shows the areas experiencing the most exploits today. It also shows us areas in which there are numerous proof-of-concept tests and other activities that suggest major problems in the near future. And it indicates areas of focused research that may ripen into active exploitation over the long term. Figure 6-3 shows how the activity in the areas of social computing and smartphones has shifted heavily to active exploitation, as previously predicted. It also shows an increase in research into threats to applications, which is likely to metamorphose into full-blown attacks in the future.

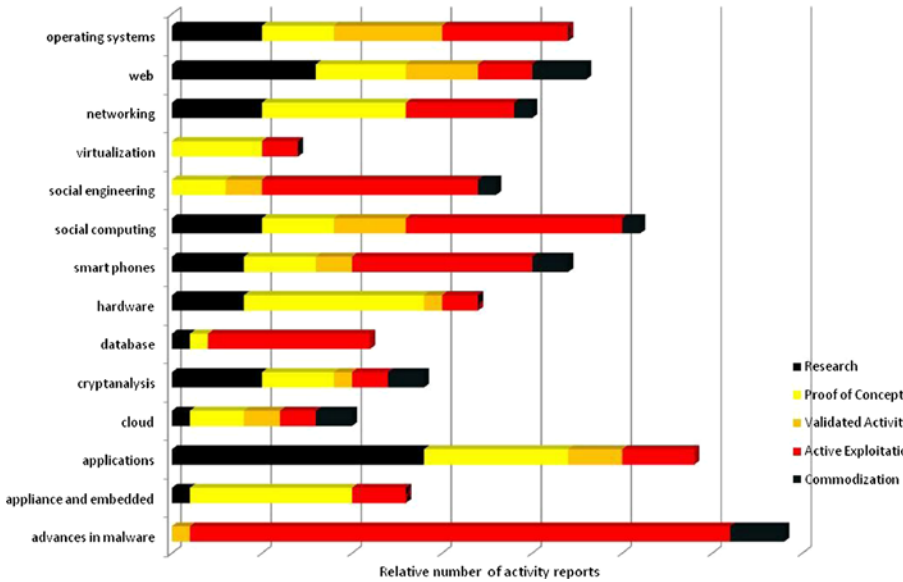

Figure 6-3. *A visual comparison of security-related activity across different technology areas. Data are for illustration purposes only. Source: Intel Corporation, 2012*

Though the depth of detail in Figure 6-3 is valuable to our security team, we have found a simpler, consolidated view can help communicate the essential trends to a broader audience. We have recently begun supplementing our threat analysis materials with charts like the one shown in Figure 6-4. These are based on the activity identified using the product life cycle model, but we add further trend analysis and group the activity areas into four main clusters, depending on their level of activity and maturity potential and on their potential impact to the company.

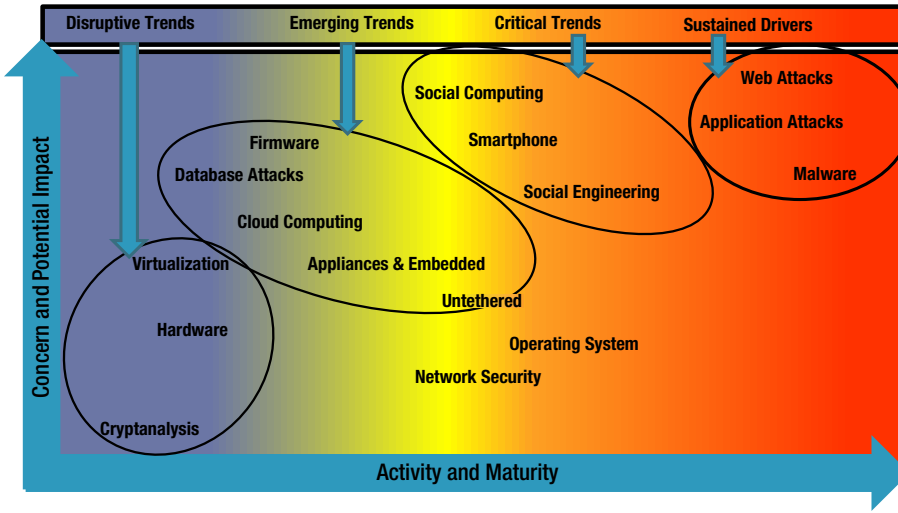

Figure 6-4. *Clustering areas of threat activity to highlight trends. Source: Intel Corporation, 2012*

These clusters are:

- *Sustained Drivers*. These are areas that already have a high impact or otherwise cause considerable concern. Typically, they are characterized by commoditized distribution and active exploitation by multiple threat agents. Today, examples include malware and web attacks.
- *Critical Trends*. These areas have begun undergoing active exploitation, with growing adoption beginning to shift toward commoditization. Current examples include social computing and smartphones.
- *Emerging Trends*. These areas have a low current level of exploitation, but considerable research and proof-of-concept activity. Examples include embedded and cloud computing.
- *Disruptive Trends*. These are areas with little or no active exploitation, but significant research activity and the disruptive potential to cause a major security problem. Frequently, they are discussed as theoretical risks, and because of this, many people in the industry would be caught off guard by a significant event. Examples include virtualization, an area in which potential threats and vulnerabilities have been exposed and a successful exploit could cause far-reaching damage.

We have found that clustering threat analysis information in this way enhances communication with stakeholders across Intel. Representing the information in

easy-to-understand charts helps to convey the key trends and their potential impact to a broad cross-section of people, helping them quickly assess whether they need to make adjustments to security strategy.

Understanding Threat Agents

Besides the product life cycle analogy, we also apply other techniques that help us think creatively about threats and identify risks we might otherwise miss.

Behind every threat is a human agent. To effectively plan our defenses, it helps if we can understand why and how these agents operate: their motives, typical methods, and targets. However, we realized several years ago that we lacked agreed-upon definitions of threat agents, as well as a clear understanding of which agents actually pose the biggest risks to us.

Some agents and their activities attract considerable publicity, resulting in the "TV news effect:" the most-publicized agents appear to be the biggest threat, so they often receive a disproportionately large percentage of limited mitigation resources. In reality, a wide spectrum of threat agents exists, some of which may be less well-known but pose bigger threats. For example, hactivists often want to publicize their activities as much as possible to draw attention to their cause. This publicity makes them appear to be a bigger threat than other groups, such as organized crime syndicates, which try to conceal their exploits.

In addition, terms often are used without clear agreement about what they mean. The phrase *advanced persistent threat* has become a buzzword whose exact meaning depends on who is using the term. It usually implies adaptive, long-term strategies employing a variety of stealthy techniques and used by attackers with considerable resources. However, it's important to remember that a variety of agents may be capable of generating this type of threat. To understand and predict their likely motives and methods, it would be more useful to clearly define the agents, whether they represent nations or other powerful groups, such as organized crime.

To solve these problems, we developed a standard threat agent library that provides a consistent, up-to-date reference describing the human agents that pose threats to our information assets (Casey 2007). The library helps risk management professionals quickly identify relevant threat agents and understand the importance of the threats.

The library acts as a collection point for information about each agent, making it easier to share information across Intel. It includes profiles of agents such as disgruntled employees, opportunistic employees, industrial spies, and politically motivated attackers. The library also catalogs agents' typical targets, objectives, skill levels, current activity, and exploit outcomes. As part of our regular threat assessments, we determine which agents pose the biggest risks to Intel. We then can use the information about their typical methods and exploits to help plan our strategy. The library helps us understand why specific events and attack trends occur and what might happen next.

Playing War Games

We conduct war games a few times a year. War games are intense role-playing exercises in which Intel employees take on the role of attackers and attempt to compromise key

assets using any feasible methods (Casey and Willis 2008). We have found war games are particularly valuable for analyzing threats that could have major consequences but whose vulnerabilities are not well understood.

This technique provides the most comprehensive method of assessing threats to key assets, because the people playing the role of our adversaries are essentially allowed to use any method to achieve their goals. However, because of this, it is also resource-intensive and should be used selectively.

A typical war game takes one and a half days and might involve eight to ten Intel staff from a variety of roles, such as factory workers, business process leads, salespeople, and technical experts.

The game focuses on a target or scenario, such as disabling a key facility or stealing Intel's trade secrets. We can use war games to examine potentially catastrophic events with a low probability of occurrence, but a high probability of causing damage if they do occur. The team members are instructed about the threat agents involved and draw on archetypes from Intel's threat agent library. Led by a facilitator, the team takes on the attacker's perspective and postulates ways to achieve the attack's objectives.

Because the team can propose any attack method, they often identify risks that might be overlooked using conventional methods. For example, a malicious group might attempt a devastating attack by purchasing a small but essential technology provider and inserting malware into their products in order to infect their customers. After each game, security analysts examine the results to determine how to address newly identified vulnerabilities.

At Intel, we also examine the cyber consequences of large physical events as part of our disaster recovery planning. These could include earthquakes and tsunamis that damage data centers, or even solar flares that disrupt the communications that the business relies on. Exercises can include drills that last a day or more.

A large company like Intel can justify the considerable effort involved in conducting these exercises because of the enormous potential benefit of mitigating the threats.

But smaller companies can also benefit by considering extreme events and formulating response plans. If you prepare for the extreme, you'll be more prepared to deal with everyday events. Planning doesn't need to be as resource-intensive as war games. It can be as basic as bringing team members together to discuss likely scenarios and responses. This method enables members to get a feel for what it would be like to work together should an actual disaster occur. Considering these extremes can also provide motivation for introducing simple yet effective measures to reduce the risk that catastrophes will occur. You might realize it is worth increasing investment in user education to reduce the risk of social engineering compromises, or becoming more diligent about analyzing logs and network traffic to identify patterns that indicate botnet activity.

Trends That Span the Threat Landscape

I've described some of the methods that can be used to analyze emerging threats. Now I'd like to turn to some of the key themes we have identified in emerging threat analysis conducted by Intel's information security group. These themes paint a broad-brush picture of threat and vulnerability trends spanning multiple technologies across the threat landscape.

Trust Is an Attack Surface

As the technology industry erects new technical defenses, attackers seek to bypass these controls by exploiting user trust—typically using social engineering techniques such as phishing.

If an attacker can win a user's trust with a sufficiently convincing e-mail or fake web site, the user will make it easy for the attacker by clicking a link or downloading a file. These actions usually undermine even the most rigorous system-level controls, initiating a chain of compromises that ultimately can result in major damage.

In 2011, this breach of trust was a common theme linking every major reported compromise. The initial stages of each of these compromises involved employees who trusted an external communication such as a targeted phishing attack.

Whenever users place their trust in a new technology, attackers quickly follow. Studies have shown users trust social media services more than other information sources—a user is more likely to click a link if it appears to have been sent by a social media "friend." Exploiting this trend, attackers have spread malware via social computing circles of trust such as friend networks.

Attackers have also been quick to take advantage of the trust users place in their smartphones and in other appliances, such as game consoles. The exploitation of trust also extends to the relationships between systems. Once configured, communications between systems often operate autonomously, without manual oversight. Smartphones are set to automatically update applications from trusted app stores; other systems blindly trust firmware updates and dutifully install them. This automation provides convenient opportunities to insert malicious code, abusing trust without the need to directly involve the user.

In the near future, we anticipate trust will become a commodity that is bought and sold. The digital reputation of systems and services will become critically important. In the past, tokens of trust, such as digital certificates and social computing credentials, were stolen for immediate use. In the future, they will be stolen so they can be sold in underground markets. The value of these tokens depends upon the access they grant and the other circles of trust they can be used to penetrate. Already, attackers are using stolen digital certificates to sign their malware in an attempt to avoid detection by operating system defenses.

I expect social engineering attacks will continue to present significant risks because they exploit human weaknesses and will adapt to take advantage of new technologies. So we, as security professionals, need to focus on the role of users as part of the security perimeter, as I discussed in Chapter 5. To reduce the risk to the enterprise, we need to make users more security-aware and influence them to act in more secure ways.

Barriers to Entry Are Crumbling

Our adversaries gravitate toward the path of least resistance. They tend to select targets that are easy to access and analyze, and they typically use the most readily available and cheapest tools.

They are much less likely to use methods with high barriers to entry such as the need for specialized expertise, expensive hardware or software, or access to extensive compute capacity. However, several of these barriers have begun to crumble as a result of trends

such as cloud computing, lower-cost communications components, and commodity malware toolsets. This trend ultimately is likely to result in new types of attack.

A key factor is that security researchers are sharing not only their knowledge but also the tools they design as part of their research. Recently publicized tools, such as rogue base stations and Bluetooth sniffers, provide attackers with more accessible, low-cost ways to intercept network traffic. Researchers have uncovered vulnerabilities in *femtocell devices* (miniature, low-cost cell towers) that can be used to take control of the devices, lowering the barriers to attacks targeting cell phone data traffic.

Using a laptop and open-source software, a highly skilled researcher demonstrated the ability to create a base station to locate and communicate with a smartphone, then crash the mobile device and install rootkit or backdoor software on it.

Ultimately, lower barriers to entry mean increased risk to enterprises. However, because several of these areas are still at the research stage, it can take time for them to mature into active exploitation.

The Rise of Edge Case Insecurity

Each day, the environment becomes more complex with millions of new devices, each running its own operating system and collection of applications. This complexity generates new edge cases—problems or situations that occur only in unexpected or extreme situations.

Edge cases can include unlikely interactions between two familiar objects. A hacker team recently demonstrated that, with a popular smartphone, a paperclip (used to pop out the phone's SIM card at the critical moment), and a little patience, it's possible to gain access to contact information, phone call logs and voice mail, e-mails, and other information stored on the phone.

Overall, the growing number of third-party plug-ins and widgets introduce edge cases that are hard for developers to anticipate even if they use secure design techniques.

Interoperability between programs has resulted in a new category of hybrid attacks where malicious objects are concealed in innocent-looking ones to thwart detection. One proof of concept in 2011 demonstrated it was possible to conceal a fully functioning Trojan in an e-mail plug-in.

Some of these hybrid attacks have shown they can circumvent new security features. As web browsers and search engines try to protect users from malicious links, attackers are responding by hiding links in image search results, where they cannot be detected using standard tools. Research into network intrusion methods has discovered over a hundred methods of evading detection by manipulating traffic to remain functional but undetectable by typical tools.

There is no silver-bullet solution for eliminating edge-case insecurities. It's unlikely even the most rigorous testing could ever uncover them all. The best approach may be to exercise caution when adopting new technologies with the potential to generate edge cases.

The Enemy Knows the System

The technology industry has often relied on security through obscurity—the idea that if attackers can't see the insecurities in code or other technology, they won't exploit them.

Over time, it has become clear that security through obscurity is poor security. To quote the maxim coined by Claude Shannon, one of the founders of modern computing: "The enemy knows the system."

It's now relatively easy for attackers to get access to the same tools enterprises use, such as web hosting services and smartphone application development tools. Hackers can now more easily engineer malware and attacks that take advantage of these elements. The fact that static platform controls tend to become less effective over time (one of the Irrefutable Laws of Information Security noted in Chapter 1) is partly due to the ability of malware authors to pretest their malicious code against technical controls.

Even the success of social engineering demonstrates that the attackers' knowledge of the target greatly increases the likelihood of successful deception. Today, competitors and other threat agents learn a great deal about a company and its employees by simply searching information publicly available on web sites or social media accounts.

Because we cannot assume insecure technology is safe just because it is hidden, we need to design with security in mind. The ineffectiveness of security through obscurity is also an argument in favor of standards and open-source solutions. This idea may initially seem counterintuitive, but the fact that open source is exposed to public scrutiny requires it to be secure. At a minimum, we should ensure devices are rigorously tested against industry standards because the attackers will do so.

Key Threat Activity Areas

Threats are evolving in many technology areas, from embedded systems to cloud computing. I'd like to discuss a few areas experiencing significant developments with implications for enterprise IT.

The Industry of Malware

Malware has become a profitable industry that increasingly resembles the legitimate software market, with market leaders, mergers, licensing agreements, real-time support, and open source. The organized business activity in this market reflects the extent to which well-crafted malware has become a viable career pursuit for members of the criminal underground.

Today, malware development and malware use may be distinct activities carried out by different groups or individuals. Malware authors are producing standardized toolkits, which have made life much easier for would-be attackers. These attackers can now simply buy or acquire a toolkit rather than expending the effort to identify vulnerable web sites and develop their own exploits.

The Zeus malware family provides a useful case study showing how complex this industry has become and how hard it is to accurately track developments. Sold mainly in underground forums, Zeus has been used extensively for theft by creating botnet nodes. During 2011, a code merger was reported between Zeus and another popular crimeware kit, complete with assurances of future support for the customers of both products. Around the same time, Zeus toolkit source code was made publicly available. Since then, multiple new variants have appeared and been used for a variety of attacks. At one point, security researchers attempting to monitor Zeus exploits discovered a server

they believed was the hub of a Zeus botnet. However, the server was the equivalent of an espionage honey pot, allowing the botmasters to turn the tables by spying on the researchers who were attempting to analyze the hub.

The Web As an Attack Surface

The Web continues to present a huge attack surface. And this attack surface is growing rapidly with the number of connected devices expected to expand to a billion or more. These include nontraditional devices such as appliances and control systems, cars, and the "smart" grid. Each of these is a potential source of risks.

For a glimpse of the probable future, consider the history of embedded devices in the enterprise environment. Companies have a history of deploying specialized devices without engineering security controls that reflect the risks these devices can introduce. Often, businesses deploy off-the-shelf devices without taking steps to harden them because of the perception that specialized devices are "dumb" and do not have a full set of capabilities.

In reality, the exact opposite is generally true. Devices marketed for a specific function are often capable of much more. Printers contain processors and may be capable of acting as file servers, for example. Furthermore, support for these devices is often outsourced, introducing a further source of potential risk in the form of external support technicians who enter the premises for monthly service visits. As a result, embedded devices can introduce as much risk, or more, to an organization as a traditional computing device since they lack security controls and administrators are generally unaware of the danger.

For attackers, embedded devices may become the path of least resistance. Embedded devices are always on and often poorly monitored. These devices store, transmit, and manage credentials and data, yet their default passwords are rarely changed. Some are initially configured to send data outside the network perimeter. Many can be remotely administered through web interfaces, making them viable points of attack. Furthermore, organizations often outsource on-site support of printers and other devices to an external supplier, who sends technicians to service the devices on a regular basis—introducing another potential source of risk that must be considered.

Security focus areas include printers and industrial control systems. In a recent example, researchers demonstrated they could replace printer firmware with fake updates capable of stealing information on documents sent to the printer, then forwarding this information to an external address. The vulnerabilities in industrial control systems were exposed by the widely publicized Stuxnet malware, which was used to sabotage systems with the apparent purpose of hampering Iran's uranium enrichment capabilities.

The incorporation of computer-based control and automation technology into the existing electrical power infrastructure—resulting in the "smart grid"—is another source of potential vulnerabilities. The US government has warned of increasing threats to the grid, noting that many embedded systems lack adequate security controls and are susceptible to known techniques such as cross-site scripting attacks (US GAO 2012).

Embedded devices, including medical equipment, safety systems, and locks, increasingly include wireless capabilities, so exploitation doesn't even require a physical network connection. Security researcher Jerome Radcliffe, a diabetic, remotely disabled

his own insulin pump live on stage at the Black Hat conference in Las Vegas. Executing the attack required less than 60 seconds. In another celebrated example, researchers demonstrated a vulnerability in control systems at federal prisons that could allow an outsider to remotely take them over and perform functions that include opening cell doors.

We might also see logical attacks as precursors to physical attacks. On a macro scale, a nation state might attack another nation's cyber infrastructure before staging a physical attack. This approach might also be applied at a more personal level. A burglar might remotely disable an Internet-connected alarm system before sneaking into a house, or perhaps even use the system's video cameras to watch the owners and note when they leave the house unattended.

Smartphones

Smartphones are attracting almost as much malicious interest as desktop and laptop platforms. The adoption curve for smartphones is steep, with no end in sight. I expect the growth curve of smartphone malware to be at least as rapid.

Just as in legitimate software markets, malware authors are likely to maximize the value of their code by using tools that allow their software to run on multiple devices. They are increasingly targeting applications, a trend also seen on other platforms. A unique aspect of smartphone application attacks is the focus on application marketplaces, which present a convenient centralized location for disseminating malware. Attackers have purchased copies of applications, incorporated their malicious content into the otherwise legitimate software, and then redistributed their code under a new name or as a "free" version of the original. On one smartphone platform, autodialing malware was found in more than 20 applications. Variations of a Trojan were found in dozens of applications and are believed to have been downloaded by at least 30,000 users.

A further development is the use of smartphones as bridges to traditional networks, resulting in the potential for enterprise network attacks that originate from within mobile networks.

In the future, we could see greater exploitation of location-based services to deceive users. Because smartphones contain location sensors such as Global Positioning System (GPS) chips, knowledge of the phone's location can be used to present targeted ads and useful information. For example, a user in a supermarket aisle might be presented with online coupons for products on nearby shelves. But this information could also be exploited to present fake coupons that are all the more convincing because they suggest that the sender knows the user's preferences.

Attackers could also exploit other smartphone capabilities to take advantage of the fact that the devices are carried into confidential meetings and other highly sensitive situations. As security expert Dmitri Alperovitch recently observed (2012), "with remote control of a CEO's mobile phone, an advanced persistent adversary could activate the microphone to record private negotiations."

Current trends in the mobile platform space indicate attackers are most interested in stealing personal data. This trend is partly due to the increasing use of smartphones for financial and banking transactions, which provides new opportunities for identity thieves and other criminal groups. As a result, it is now important that smartphone hardware and software developers focus on protecting personal data. Software developers should

adopt the same discipline and commitment to following secure design principles as traditional platform developers. Today, more and more people are becoming app developers—creating software and posting it online for others to use. One has to question how much security testing and validation has been applied to these applications. As users move more of their everyday activities onto smartphones and other small devices, the consequences of poor or insecure designs will have greater impact on individuals and their employers.

Web Applications

Web applications, primarily comprising client browsers and server-based applications, continue to be heavily attacked. In our threat analysis model, we characterize this area as experiencing full exploitation activity and moving toward commoditization. There is also considerable research in this area, suggesting the number of attacks will continue to grow.

Attackers have adopted new techniques to hide their intentions and deceive users long enough to achieve their aims. As web browsers and search engines try to protect systems from malicious links, attackers are instead obfuscating their links in image search results, where they may not be detected.

Techniques for hiding messages within images have been used within the security realm since long before the invention of information technology. Now, this technique, known as *steganography*, is being used to hide malware and botnets on publicly used image hosting sites.

Search poisoning has also become a common method. Attackers using search poisoning tend to focus on events and topics of popular interest, optimizing their web pages to achieve high search engine rankings. After a search query, the victim clicks a link among the search results. They are redirected multiple times and eventually land on a page that is used as a vector to deliver malware.

Conclusion

In this chapter, I've outlined some of the real threat trends and described methods information security groups can use to analyze the threat landscape as it continues to evolve.

No doubt, new and more-sophisticated types of exploitation will continue to emerge, and we need to stay aware of them. As Mustaque Ahamad, director of Georgia Tech Information Security Center, noted recently (2011), "We continue to witness cyber attacks of unprecedented sophistication and reach, demonstrating that malicious actors have the ability to compromise and control millions of computers that belong to governments, private enterprises, and ordinary citizens."

Yet, as we try to make sense of the deluge of news about attacks and vulnerabilities, it's essential to retain a sense of perspective. Most threats do not take place using exotic, obscure methods. Instead, they take the path of least resistance, exploiting well-known vulnerabilities. Therefore, business can mitigate many of these threats by implementing basic, established security measures. To put it another way: when you hear hoof beats, think horses—not zebras.

Social engineering will continue to be a key attack method because it takes advantage of user trust and is hard to prevent using technical controls. Therefore, as I discussed in Chapter 5, we need to continue to focus on educating users to become more security-aware. By doing so, we can reduce the risk to the enterprise.

Ultimately, while doing our best to prevent compromises and breaches, we must remember we cannot control the threat actors and their exploit attempts. For all organizations, some level of compromise is inevitable, making defense in depth as essential as ever. Losers ignore the trends. Winners survive by being able to predict, prevent, detect, and respond.