O'REILLY®

Cracking Security Misconceptions

Untangling Common Myths About Modern Information Security

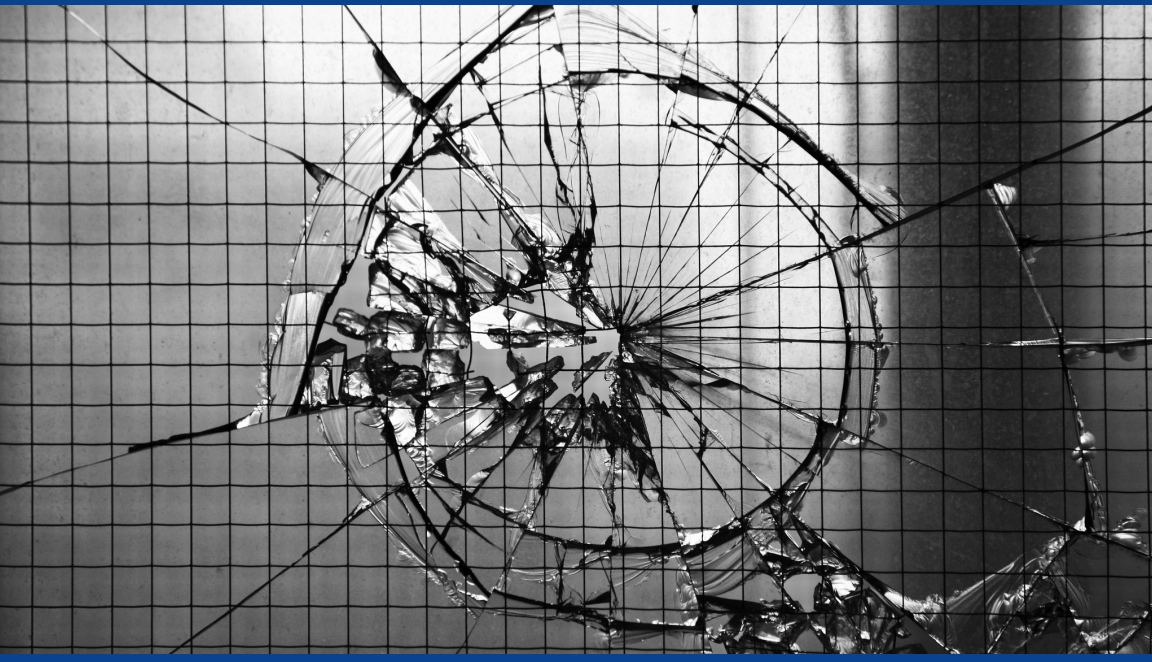

Andrew Peterson

3 Easy Ways to Stay Ahead of the Game

The world of security is constantly changing. Here's how you can keep up:

1. **Download free reports** on the current and trending state of security. **oreil.ly/Security_reports**

2. **Subscribe** to the weekly Security newsletter. **oreil.ly/Security_news**

3. **Attend the O'Reilly Security Conference,** the must attend conference for security professionals. **oreil.ly/Security_conf**

For more information and additional Security resources, visit **oreil.ly/Security_topics**.

Cracking Security Misconceptions

Untangling Common Myths About Modern Information Security

Andrew Peterson

Beijing • Boston • Farnham • Sebastopol • Tokyo O'REILLY®

Cracking Security Misconceptions
by Andrew Peterson

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*http://safaribooksonline.com*). For more information, contact our corporate/institutional sales department: 800-998-9938 or *corporate@oreilly.com*.

Editor: Courtney Allen
Production Editor: Colleen Lobner
Copyeditor: Octal Publishing, Inc.
Interior Designer: David Futato
Cover Designer: Randy Comer
Illustrator: Rebecca Demarest

September 2016: First Edition

Revision History for the First Edition
2016-09-06: First Release

978-1-491-95628-1

[LSI]

Table of Contents

Cracking Security Misconceptions

Introduction

Companies, governments, and organizations are failing to secure information in today's digital world, and the stories of those failures continue to mount. Crime has always been around. But the things a criminal can steal and the technology through which they can steal things has changed dramatically with the introduction of information technology. Cyber criminals, the people who use these new mediums to perform illegal activities, are finding ways to exploit faster than we can figure out how to defend against them. As a result, the criminals are winning and the defenders are by and large playing catch up.

So there's nothing we can do, right?

If you had asked me that question five years ago, back when my only understanding of cyber security was based on the stories I heard in the media, I might have said yes. But in the process of starting a security company with a number of leading security professionals, I've learned how far from reality my understanding of security was. And, the more I've shared those learnings with other nonsecurity professionals, the clearer it is that the misconceptions about the world of hacking are widespread.

If you're like most people I talk to, you're more aware of cybercrime than ever and you might even be incorporating security into your job responsibilities. So you are eager to learn! But here's the thing: I never had someone sit me down and reorient me to the real world of security because, unfortunately, security professionals are largely unaware of the gap in understanding that exists for those outside of

their world. They assume, like most people do, that everyone else knows the world like they do. Consequently, it's taken me years of direct experience to piece together lessons that represent a foundational understanding of the security challenges we face.

The world of information security needs the help and collaboration of nonsecurity professionals across their organizations to bring more attention and innovation to the problems that face the industry (and insider reports agree; see the following: 1, 2, 3, 4). To do so, you need to be equipped with an accurate understanding of the increasingly nimble and effective opponents we're all up against. In the following pages, I'll save you some of the trouble—and years—I went through getting up to speed by breaking down the most common misperceptions about security risk. Soon, you'll be informed and better prepared to join the fight.

Misconception #1: Hackers Are Criminals

All hackers wear black hoodies, have tattoos, work in dark basements with special computers, and methodically destroy whoever their target is for the day while listening to trance music.

At least that's what I used to think. In my defense, that's certainly the closest to what I've seen or read about in movies and books at that time. How was I supposed to know any different? And although many misconceptions about hacking and the world of cyber security persist via the media, the most basic one is that hackers are all dark, malcontented criminals.

The reality is that hackers—and the activities they perform—span the gamut from safe to legal to criminal, and the people in the industry come in all shapes and sizes (though, to be honest, the black t-shirt is a bit of an industry uniform). There's a wide gulf between how hackers are portrayed in the media and what hackers really are. Let's begin by breaking down the basic groups involved in the industry, which you can see in Figure 1-1.

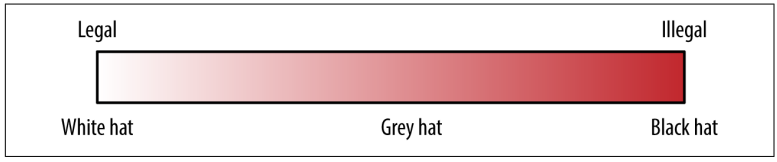

Figure 1-1. The range of hackers: white hat, gray hat, and black hat.

White Hat

White-hat hackers are the so-called "good" hackers, named after the good guys who wore the white hats in westerns. They're usually computer security specialists who test and assess the security that goes into systems and networks. They have the intention of helping organizations fix vulnerabilities instead of exploiting them and often have permission from the system's owner, which makes their activities legal.

Companies typically hire this type of hacker, who are usually seen as ethical, in order to make their systems less vulnerable to any future attacks. These hackers have driven many of the advances made to online security over the past two decades, such as security improvements in email, credit card processing, ecommerce, and even Internet-connected health devices.

Penetration testing is one example of white-hat hacking. Either an internal group or (more often) a contracted company is tasked with looking for holes that a hacker could exploit in a company's systems. Their objective is to find security weaknesses, test compliance standards, and deliver a report with the findings.

Many companies also have started embracing white-hat hacking with bug bounty programs. In the past, if a white-hat hacker found a vulnerability in a given system or website and were to report the security flaw to the company, she didn't know how the company was going to react. It could either be welcomed as help or just as easily be seen by the company as an illegal and unauthorized attack for which the company could, and often did, seek legal action against the hacker. A bug bounty program makes the intentions of the organization clear by providing a process and guidelines for white-hat hackers who have found a vulnerability to safely report it. Often, there are rewards of public recognition or even cash compensation to the person reporting the vulnerability as a show of gratitude for helping to make their system more secure. Companies such as Bugcrowd, HackerOne, and Synac are helping their clients adopt these

forward-thinking security bug bounties, making them easier and more cost effective than ever before.

Security Conferences

In one sign that hacking has become a legitimate industry, many conferences are devoted to it. Conferences can be a great way to learn more about hacking. They have keynote presentations, hands-on activities, and competitions. Here is where white-hat hackers show off the latest attacks they've performed. The original conference, DEF CON, is the largest, but security-related conferences continue to grow every year. You can find national conferences, international conferences, local conferences, or conferences that specialize in a certain type of hacking; look for one that suits you. Try this list or search for "hacker conferences" to find the most recent and relevant.

Black Hat

Black-hat hackers are named after the bad guys who wore the black hats in the classic western films. The main difference between white- and black-hat hackers is their intent. Black hats use the same methods as white hats, but their purpose is to breach Internet security measures for their own personal or monetary gain. Often they use social engineering techniques such as phishing to gain information that allows them to gain access to a database. For example, they might steal credit card numbers or social security numbers to sell to identity thieves, or they infect a web application and database with malware to destroy data.

Most of their activities fall into the illegal realm because they don't have permission and they're out to cause harm or make money. Think of them almost as the 21st-century equivalent of an old-fashioned bank robber.

One way to distinguish between white-hat and black-hat hackers is that white-hat hackers like to *raise awareness of a problem* or *improve* security systems, whereas black-hat hackers like to *exploit* holes in security systems.

Gray Hat

Just like in real life, not everything is black and white. Gray-hat hackers fall in between white-hat and black-hat hackers.

They hack without permission, and might disclose their findings publicly instead of privately for the target to fix. They might not be out for monetary gains; nor are they always predisposed to help companies secure their networks.

Unlike black hats, gray hats aren't typically malicious; they mostly hack because they're interested in how a system works. They might hack an iPhone to bypass authenticating it with a phone company, for example. Many times, however, these activities still fall in the illegal realm because they don't have permission.

For example, as cars go digital, they have become a popular gray-hat hacking target. It can be fun and safe—making horns honk or turning on and off radios in a lab or garage—or it can creep into the malicious and dangerous realm—disabling a transmission or accelerator of an innocent driver on the freeway. Whether it's for fun or not, hacking done recklessly, irresponsibly, and without the consent of others classifies it as grey hat. Sometimes, the hacker also gains something from the hack: an increased reputation, a consulting job, or money by selling the vulnerability on the black market.

Wrap-up

White-hat, black-hat, and gray-hat labels aside, the hacker community is growing more diverse in a variety of ways. Although hackers started as a group of self-taught tinkerers, it has matured to the point that a number of universities even offer Computer Security degrees (though only a few). As security stories have become more mainstream, so too has the community. If you attend a security conference you'll encounter people from all over the world, men, women, young, old, engineer, businessperson. In subsequent sections, we'll continue to uncover different classifications of hackers but it's important to understand that being a hacker can mean many things; but it does not mean that you are a criminal by definition.

Misconception #2: Hackers Must Be Geniuses

Hackers are all such off-the-charts geniuses that defenders have no chance to stop them, right? How else would hackers be able to find

loopholes and backdoors that allow them to break into someone else's system other than being overwhelmingly smarter than those trying to defend it?

It's easy to believe this misconception.

Frank Abagnale, Leonardo DiCaprio's character in *Catch Me if You Can*, wouldn't have been nearly as fun to watch (and secretly root for) if he weren't so darn clever to continually outwit Tom Hanks' FBI agent character, Carl Hanratty.

And what PR group—not to mention legal group—would want the story of their company's data breach to be about how easy it was for the hacker? Instead, they want to make sure everyone believes that they were compromised because of highly sophisticated and never-before-seen hacking methods that they couldn't have possibly predicted or defended against so as to save them from lawsuit and embarrassment.

The reality is different.

The task of a defender is much more difficult than the task of an attacker. A defender needs to keep an eye on, and defend against, every possible way she could be attacked, whereas an attacker only needs to know one way in from among the many possible doors.

This imbalance has been exacerbated over the past 15 years for both defenders and attackers. Here are some of the key components:

- Defenders are working at companies and organizations that, starting with major investments in IT infrastructure in the early 1980s all the way to today with the rise of Software as a Service (SaaS) tools, have all been adopting technology to work more effectively and efficiently. The result of which is an increased technology landscape for hackers to attack and defenders to defend.
- The Internet makes it possible for these technologies and services to be accessed anywhere in the world. The sheer number of potential attackers against a given organization has increased exponentially.
- The tools, techniques, and education available via even a simple Google search in some cases to teach and enable hackers to attack have become more prevalent, more automated, and dra-

matically cheaper to the point at which, in many cases, they're completely free.

What this has resulted in is an increase in vulnerabilities (or undefended ways into an information system) and an increase in attackers. In particular, the number of inexperienced, unsophisticated attackers (commonly referred to in the industry as script kitties) has grown significantly. These attackers rely heavily on tools and techniques developed by others instead of having to come up with new, specialized tools for each organization they target.

So, even though some attackers might be as smart as you'd expect, they certainly don't have to be geniuses anymore to be successful (and often times they aren't). Hackers vary greatly in regard to skills and experience. Here are a couple of real-world examples of highly publicized hacks that were relatively unsophisticated.

United States Department of Justice

In early 2016, a hacker accessed the US Department of Justice servers. These servers require a two-factor authentication to gain access, a feature that offers a higher level of security. How did the hacker get access to the information on these supposedly safe servers? He did it by using a simple social-engineering attack: he called the help desk, where a helpful employee gave him the second authentication code. With that code, he easily had access to the servers, where he downloaded several gigabytes worth of data, including the US Department of Homeland Security employee directory. Of course, policies are in place directing employees to not to give out that information over the phone; an employee needs to go in person to show identification to prove the code is needed. But, the hacker was able to exploit someone wanting to be helpful and who was persuaded to make an exception to the rules.

Target

In 2013, Target's customer names, credit and debit card numbers, expiration dates, and security codes were stolen from its secure servers. Hackers installed malware (also known as a computer virus) on Target's systems, which gave them access to 40 million debit and credit card numbers entered in at the point-of-purchase systems. Multiple Target security systems had flagged the unauthorized malware. But the flag had to be reviewed by a person who would

instruct the system what action to take. No person reviewed the alert and no action was taken. The malware was not craftwork of a genius. In fact, it was particularly ordinary and it was easily identified by multiple internal tools. The breach only happened because of a breakdown in process, not because of brilliant tactics.

Both the Department of Justice and Target hacks demonstrate the true security landscape. Attackers do not need to have an exceptional intellect or rely on discovering the one highly sophisticated technical back door to protected data. Instead, they can use basic, off-the-shelf tools or simply find the right person that will unwittingly let them through the proverbial front door.

Misconception #3: Hacks Are Sophisticated and Complex

When it comes to the world of hacking, it's important to understand not just who a hacker is, but also the actual hack itself. Similar to the assumption that hackers are all geniuses, many people assume that their methods are similarly complex and sophisticated. But just as hacker skills range in sophistication, their methods do, as well.

The most unsophisticated examples typically arise because of human error. Take, for instance, the password. After each major login/password breach, analysts review the data and find people use the same passwords. So much so that upward of 5 percent of people use the same 100 passwords. This means that if you wanted to try to hack into someone's account, you have a 1 in 20 chance of getting in by just trying the top 100 passwords. As a quick aside, this is very easily stopped if companies occasionally reviewed the most common passwords and didn't allow users to set them.

Although there are many attack techniques that span sophistication levels, the following is a basic breakdown of some of the most common types of attack categories, including examples.

Social Engineering

Social engineering differs from other attacks because it depends on human interaction. Here, the hacker manipulates people into performing an action or divulging confidential information. The hacker relies on people's natural inclination to help. It's usually easier to trick someone into giving information rather than hacking for it; for

example, fooling someone into revealing a password, rather than attempting to brute-force it by running a computer program that tests hundreds of thousands of password options automatically.

An example of social engineering is *phishing*. The hacker sends an email that appears to come from a legitimate email address from a trusted organization (a popular choice is a bank), claiming the recipient needs to update a username and password, and provides a convenient link to click. The email might come from a domain like wellsfargo-alerts@passwordrecovery.com that makes it looks official even though it doesn't come from the Wells Fargo domain. It looks exactly like past emails from Wells Fargo all in an attempt to get the recipient to think it's real. If the recipient clicks the link, she goes to the phisher's site, which is designed to look legitimate, not the trusted website, and provides her private information for the hacker to scoop up and then use to gain access to the actual account.

Network Attacks

A network attack is when a hacker performs an intrusion on a network infrastructure or host system. The hacker analyses the network address of the targets, takes advantage of open ports or vulnerabilities, and collects information. These attacks can be passive (in which information is gathered, but not changed) or active (in which information is altered); they can occur from within an organization or from outside.

An example of a network attack is a man-in-the-middle attack. Often seen as MITM, MitM, MIM, MiM attack, or MITMA, a man-in-the-middle attack is when the hacker relays communication between two other parties using the opportunity to capture or modify the data (see Figure 1-2). The two parties believe they're communicating with each other, when in reality, the hacker is intercepting and potentially altering the messages.

The hacker completely controls the messages for his own purposes. This could be to gain financial information being relayed to a bank, login information to a website, or any messages encrypted by a public key.

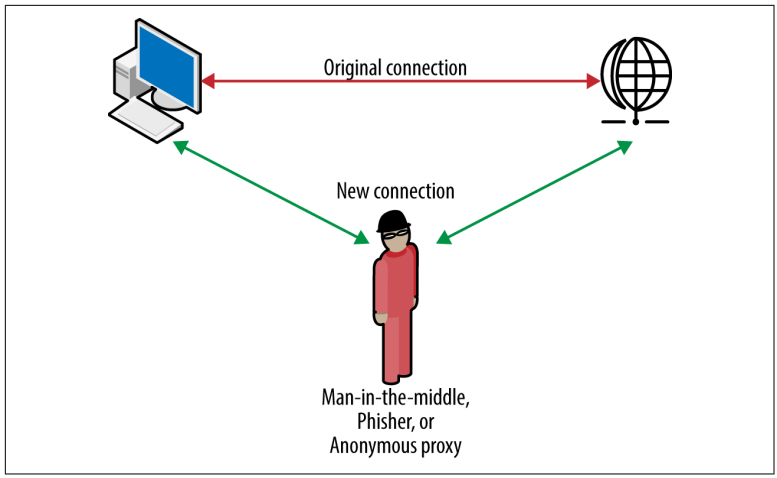

Figure 1-2. A basic man-in-the-middle attack.

Web Application Attacks

A web application attack happens when a hacker targets vulnerabilities to a service that's connected to the web (website, mobile application, etc.). Software that used to be installed on a desktop (for example, Microsoft Excel) is rapidly moving to the Internet (Microsoft Office 365 and Google Spreadsheets are run in a web browser instead of a local computer) so that you can access and run it on your computer, phone, or tablet anywhere in the world. Unfortunately, this also means that hackers can easily access it anywhere in the world, as well. As a result, this type of attack has grown in frequency. The application layer, which is easily accessible from the Internet, makes it a particularly soft target.

A *SQLi*, or SQL injection attack, is an example of a web application hack. A hacker exploits a code flaw (also known as a security bug) in a web application with malicious SQL statements that make the application potentially return any data that's available in that website's database (passwords, credit cards, addresses, etc.).

Although this type of attack typically results in stealing a copy of the data to sell, attackers can also use SQLi to tamper with data, such as voiding transactions or changing an account balance. And, in some cases, the hacker can even take over as administrator and controller of the data.

Endpoint Attacks

The endpoints of a network make them the least secure; every time someone connects a mobile device to a corporate network, plugs in a USB drive, or downloads an email attachment, a potential hole in the network is created when not done securely. As soon as a hole is created or identified, a hacker can take advantage and install malware onto a network.

An advanced persistent threat (APT) is a type of network attack that relies on vulnerable endpoints. The point of an APT is for the hacker to stay undetected for as long as possible, keeping access to steal a large amount of data. The hacker must continuously rewrite code to stay undetected, making this type of attack time consuming and sophisticated.

Wrap-up

Although this list of attack categories includes the most common methods, it is by no means all-encompassing, nor are attack types static in nature. As long as there's data worth stealing, there will be people attempting to get at it by whatever means necessary. The bar for how easy or unsophisticated the hack is that's required to break into a system, however, is dependent on how well defended it is. And, unfortunately, that bar has been dipping to the point where the most basic techniques can be successful. Luckily, many security professionals are working hard to push the bar back up.

Misconception #4: Hackers Have No Reason to Attack Me

The next misconception we'll tackle is whether a hacker is only interested in big, well-known organizations with terabytes of information available to steal. The stories that make it into the news—the Target, Home Depot, Department of Homeland Security hacks—perpetuate this fallacy. These organizations had servers that contained massive amounts of valuable data that was worth stealing.

The vast majority of people and organizations assume that hackers are purposefully targeting everything they attack and, therefore, the less well known you or your organization is, the less likely you are to be hacked.

This belief was born decades ago when many of the techniques and tools to carry out a hack were typically expensive, homegrown, or hard to access. Today, though, the world of hacking tools and techniques has never been more accessible, affordable, and—most important to this misconception—scalable. This means that attackers can automate and replicate attacks by using computers instead of having to do it all manually. As a result, hackers can aim their attacks against a much wider group of targets. The minimal time, effort, and expense makes the payoff of even smaller targets with shallow pockets worth it to a hacker.

Today, it doesn't matter if you have a well-known brand, you're running your own personal website for fun, or you're somewhere in between. If you have any type of data worth stealing, you have to consider yourself a potential (and even likely) target for an attack.

Are Hackers Out to Ruin My Business?

You do hear of the rare story of a company going out of business because of an attack, whether because of financial reasons or a damaged reputation that can't be recovered. But generally, hackers aren't out to sabotage your business. They're more interested in copying your data, and maybe leaving a backdoor into your system so that they can come back again. You might not even be aware that you've been hacked.

This doesn't mean that hackers don't attack specific companies or organizations anymore. It still happens. But exploring the various motivations of today's attackers is a critical part of understanding what they're after in the first place.

Motivated by Knowledge

When hacking first became popular in the 1960s, it was simply to gain knowledge. Computers were expensive, physically large, and typically only owned by the largest enterprises or universities. Hackers "broke in" to access these computers to learn new technology, get computers to perform new activities, or output interesting data. Access to technology is a lot easier today, but hackers are still motivated by the challenge and the joy of exploring technology to do interesting things.

Motivated by Money

Most hackers are just in it for the money (around 80 percent in fact, according to the latest Verizon Data Report). They fall into three types of categories:

White-hat hacking for rewards
: Many companies have instituted bug bounty programs or hacking contests, compensating hackers for disclosing found vulnerabilities.

Smash and grab
: The same way a burglar can smash in a window or door to break into your house, a hacker can exploit an "open door" into a security system. They get in and steal what they can without any concern for setting off an alarm. They're gone before anyone responds. Or, they leave backdoors in the system so they can have long-term access and steal data over time.

Espionage
: These hackers are typically highly skilled and well resourced (not a casual hacker) and engage in industrial or national-security espionage—hacking to steal intellectual property, confidential or sensitive information, or personally identifiable information (PII) for the benefit of another hacker, company, or government. This hacker could also be someone who poses—with fake credentials—as an employee with the sole purpose to obtain confidential information to give to the real employer or government, or sell to the highest bidder.

Motivated by Politics

Hacktivism—a combination of hacking and activism—is the practice of hacking for a political agenda. This kind of activity can encompass everything from cyberterrorism to technological hacking to affect social change. They're solely out to create awareness and/or create fear and chaos by disrupting critical infrastructures. You'll find a wide spectrum of hackers, from promoting awareness of social issues to stopping or impacting a political movement (hacking a political candidate's website), and everything in between.

The Sony Pictures Entertainment Hack

In late 2014, a hacker group called "Guardians of Peace," or GOP, hacked into the Sony Pictures Entertainment (SPE) servers using malware and released confidential information, such as personal employee information, confidential emails, and unreleased films. Their motive: to force SPE to pull its movie *The Interview* from theaters. In the end, not enough movie theaters would show it—after the GOP threatened to bomb theaters—and Sony released it directly to video.

The hacking happened over a period of a year, with GOP first using simple social-engineering tricks to test SPE's security, and then moving on to more sophisticated techniques. No one at Sony noticed the repeated attacks on its security system; no security alarms were triggered. Initial attempts to blackmail executives went unnoticed or ignored. No one at SPE took it seriously until *after* GOP started leaking data, when it was too late to properly respond to the security breach. It's still unclear who perpetrated this attack and for what reason but it was illegal (black hat) and appears to be meant to embarrass and hurt Sony as opposed to be for financial gain.

Motivated by Revenge (Insider Threat)

Hackers motivated by revenge act maliciously to cause harm to an employer or person. They can act out of greed, anger, ideology, loyalty, adventure, blackmail, ego, or just have general problems with their work situation or interpersonal relationships.

A hacker motivated by revenge might be a longtime employee who feels overlooked by management or not compensated enough. He could download sensitive data to a USB drive or upload it to the cloud to get the data out, post it to the Internet via public sites like pastbin.com, and point the press to it in order to make damaging information about the company publicly available.

Wrap-up

When you know why hackers would be motivated to attack you or your business, you can proactively set up defenses to prevent a hack; barring that, you can make smart decisions to mitigate a hack in progress. If you keep in mind that hackers can easily attack anyone

today—from the large corporation to a personal website—as well as the motivations they might have to attack you, you realize that you need to stay on the offensive.

Misconception #5: There's No Money in Hacking

The most common reason I hear from friends and family about why they don't invest in improving their personal data security, such as using unique passwords or enabling two-factor authentication on their personal accounts, is that they don't think their personal information is valuable. "Is there really someone who would pay for my passwords?" they ask.

The short answer is yes. Here's the typical explanation I give them of why.

There's a lot of things an attacker can steal but let's say she is after your login and password for a given website. For starters she could steal not just your login credentials but the site's entire database. Most likely you use the same login and password she just stole for other sites—maybe even every site that requires a login/password—including your bank accounts, ecommerce accounts, and social-media accounts. The attacker now has access to not just all of your online accounts, she has access to all of the accounts from the entire list she stole.

So is a singular login and password worth enough? Maybe so, maybe not. But attacks typically target large datasets. And data used in bulk combined with automated programs to hijack and take actions on your accounts can be used for meaningful financial gain.

Have You Been Hacked?

If you want to know if your login and passwords have potentially been exposed, go to *https://haveibeenpwned.com/* and type your email address. The site will search a database of 1.3 billion known hacked emails and passwords and let you know which (if any) known hacks your information has been publicly exposed.

Hackers can absolutely make money from stealing your login information. But, in addition to that, there are many ways that both

black-hat and white-hat hackers make money, and almost none of them are commonly understood. Here's a quick breakdown.

Hackers can make money on the black market:

Sell the data
: When hackers have information they can sell, such as credit card or Social Security numbers, one place they head is to the *dark web*. The dark web consists of websites that are visible to the average person, but its Internet Protocol (IP) address is hidden through *Tor*. A Tor encryption tool bounces an IP address through several layers of encryption so that the address appears as another address. The dark web and any transactions done with it are not indexed by search engines.

Use the data to get other information
: Hackers can also use data they find to gain access to other information. Hackers take advantage of the fact that people tend to use the same username and passwords for multiple websites. When a hacker gains access to a stolen list of usernames and passwords from one site, they can use that information against other sites to gain access to email accounts and then use that information to find bank accounts. They can then sell the information to someone else or transfer the funds for their own financial gain.

Extortion
: Another way hackers can make money is to exploit a vulnerability in a system (even a temporary one), and then blackmail or extort money from the owner of the system to fix it, or threaten to sell the information to another hacker so that they can exploit the vulnerability.

Zero-day market
: Zero day refers to the window of time between when a vulnerability is exposed and when security vendors release patches to shore up the vulnerability. During that time, hackers can sell the security vulnerability either to a government or a business competitor on the zero-day market. The value is in the vulnerability itself that other hackers can exploit until a patch is released and installed.

Hospital Ransomware Extortion

Hospitals have become popular targets for attackers via ransomware (see the following: 1, 2, 3). Ransomware is a generic term for software that takes complete control of someone's computer or, in this case, an entire network of computers. When an attacker locks out hospital employees—doctors, nurses, patients, and administrators alike—from their computers, they severely disrupt patient care, which has both dire health impacts for patients and enormous amounts of lost money for the hospitals and doctors. Oddly, attackers have asked for relatively small amounts ransom ($10,000–$20,000, for example) compared to the overall financial impact the hacks cause. As a result, hospitals are advised to simply pay the attacker to fix the system, rather than try to fight them as delays could mean lawsuits or patient deaths.

Now let's talk about how hackers trading on the non-black market can make money:

Working for a government's intelligence agency
: These hackers break into the systems of its foreign adversaries, and occasionally its allies.

Working for a security company
: Many hackers might start off with black hats and turn white hat by consulting or working for a security firm that offers companies security advice. Penetration testing companies are one type that help businesses find weaknesses in their security.

Working for an organization that has an in-house security team
: As mentioned earlier, more companies are hiring their own in-house security teams to secure their data and internal systems. The demand for these skills in-house is rapidly growing and the compensation packages are increasing in kind.

Bug bounty programs
: Many companies offer compensation to hackers who find vulnerabilities in their software. Officially, these are called bug bounty programs. In the end, bug bounty programs are a way for companies to embrace white-hat hackers, acknowledge their capabilities, and are an inexpensive way to find and patch vulnerabilities in their systems.

Hackers can make a living in a variety of ways, ranging from the illegal to the legal. When you understand how they can make money by hacking, you can begin to understand how they attack. And, ultimately, build a more effective defensive strategy against the attack.

Misconception #6: Big Organizations Are the Most Secure

I recently attended a talk given by *New York Times* bestselling author Marc Goodman to a group of well educated, albeit nonsecurity, professionals about security issues and a comment came from the crowd that went something like this:

> So you're telling us about all these threats to our data security but we don't need to worry about that with our financial, health, and government data right? I assume they have teams of experts working to protect us.

To which Marc replied with a wry smile and a pause. He then went on to give a nuanced and thoughtful response that boiled down to "it depends."

I had a similar initial reaction to the comment (but it was more of a smile and a sigh) because it represented another common misconception I also had about data security in the past. I assumed that big organizations that have the money and the means to protect the most sensitive data—and have the most to lose—have the best data security.

The problem is that having great security is a natural subset of having great technology. The majority of the biggest organizations and companies of the world are not technology organizations at their core (though this has been changing over time).

Nike makes shoes, not websites.

The National Parks Department manages amazing nature experiences, not online experiences.

Mayo Clinic strives for the best medical care possible, not the best access to your digital health records.

By contrast Google, Microsoft, and Facebook are technology businesses whose products fully revolve around technology and therefore have huge teams devoted to building that technology.

The point is that the key to cyber security is the "cyber" part. It's technology that has paved the way for the new data security problems that we face today, and it's technology that is the key to solving these problems, as well.

And whereas the Nikes, National Parks Departments, and Mayo Clinics of the world historically haven't seen themselves as digital tech organizations, they all have been increasing their investment in technology at some level. Understanding the various levels is critical to understanding why some of these large, important organizations don't have the type of data security that you might expect.

Here are a few common examples of how technology is built and maintained (see also Figure 1-3):

Fully outsourced with minimal ongoing support
: This has been a common scenario primarily (but not exclusively) for government contracts where a technology project is defined to achieve a stated policy goal (like a database and a website to access that data for a specific department). There's a budget defined and approved for the project. A group of government contractors bid on the project. One contractor gets picked. The contractor builds the technology, gets paid, and provides minimal support on that technology over the next 15 years per the terms of the contract. The government entity has no in-office resources to continually maintain and upgrade the technology that was built for them. The contractor is contractually obligated and incentivized to spend as little effort possible to maintain that technology. The technology degrades. It's not actively defended and it's not actively monitored. So even assuming the organizations care about the security of the data in the system that was built, they have no warnings or alerts to identify a problem when the system breaks and no one to fix the problem.

Fully outsourced with support and maintenance
: This is basically the same scenario with more explicitly defined ongoing maintenance work included. This set up is more common for businesses that are more committed to long-term success of their projects than political projects built around election cycles. The maintenance is focused on the functionality of the system for business gain (e.g., the checkout function needs to be fixed to complete transactions) and rarely includes security

services as they tend to drive maintenance costs up. Systems are still typically unmonitored and undefended.

In-house, small team
: Organizations turn a major corner when they decide to bring technology creation, maintenance, and support in-house; they plan to make a meaningful investment of time and resources into technology on an ongoing basis. In other words, technology is no longer a project—it's a part of business. You'll see this in forward thinking, nontechnology centric businesses and organizations that believe an investment in technology will have a meaningful impact on achieving business goals over time. These groups don't tend to have a security specialist on their team. But the best teams set up an amount of ongoing system monitoring to identify problems that need addressing (inclusive of some types of potential security issues). They're most likely still dramatically under-resourced but they have a chance to identify and fix basic issues.

In-house, large, strategic team
: Most large organizations and even some government organizations are starting to understand that they need to make a strategic investment in technology to stay relevant in today's economy. Regardless of whether their product is technology or not, these companies understand that they need to go online to reach their users and that data about their users is a key leverage point for their future business. These organizations likely have full-time employees—if not even teams of full-time employees—devoted to data security and stopping hackers. These security professionals are building and continually evolving their company's security strategy, tools and processes to try to keep up with an evolving attack landscape.

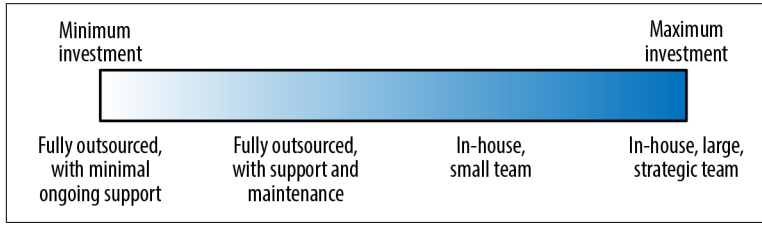

Figure 1-3. The spectrum of options for building technology capacity within an organization.

US Digital Services

In 2014, the White House started its Digital Services group, dedicated to building a world-class technology team within the government. Its goal is to make it so that citizens can access government services as easily as buying a book from Amazon or paying a utility bill online.

The Digital Services group is an in-house team consisting of some of the country's top technologists. It partners with the nation's leading civil servants on the most important federal services in a range of capacities, from consulting to building to maintaining technology. The group has worked on numerous projects so far, but here are a few examples: it has worked with Veterans Affairs to launch Vets.com, Citizenship & Immigration Services to streamline the immigration process, and the US Department of Education to launch College Scoreboard to assist high school students with college decisions.

Regardless of where a company's technology investment falls on the scale of options, most nontechnology companies continue to lag light years behind the Googles of the world in terms of both technology and security. The good news is that the majority of the best and biggest companies have all either built their own in-house technology teams or they're beginning to. Additionally, these organizations have begun to recognize the importance of security to their users and have begun building dedicated in-house information security teams within their technology groups. That said, large organizations that have in-house technology and security teams vary drastically in how much and how long they have invested in building these teams. Some companies have cutting edge, highly effective technology and security teams, whereas most are either focusing on the very basics or are perpetually trying to play catch up to the innovators. This is why Marc responded to the initial question about how all big organizations have experts protecting us with "it depends."

Misconception #7: If I'm Compliant, I'm Secure

Imagine that you're at a forward-thinking, large organization that not only has its own in-house technology team, you also have the

means to hire security professionals. So how do you go about creating a great security program?

Historically, the most common approach has been to build your program around achieving compliance standards of various forms. You've probably heard of some of these audits and checklists; PCI (for online payments processing), soc2 (accounting report for publically traded companies), and HIPPA (for handling patient medical information) are a few. They're well known, official, and are industry-regulated security standards. So it's understandable that people have the impression that being compliant equals being secure.

But that's not the case. Compliance doesn't make any organization inherently secure. Some of the components of a compliance audit and checklist can help to improve security systems. But compliance is a business function that, in the best case, creates security benefits. Virtually all of the recently publicized data breaches of large companies, including Target, were companies that had passed a variety of compliance audits. But the checklist didn't stop the attackers.

As Target's then-CEO, Gregg Steinhafel, wrote in an email statement, "Target was certified as meeting the standard for the payment card industry (PCI) in September 2013. Nonetheless, we suffered a data breach. As a result, we are conducting an end-to-end review of our people, processes, and technology to understand our opportunities to improve data security and are committed to learning from this experience."

The more accurate reality about compliance for those actually implementing it is that it's simply a series of hoops that you must jump through to conduct other business. For example, PCI certification is necessary for any ecommerce business wanting to process its own payments. An auditor is hired and the team spends months building documentation, systems, and processes to pass inspection. But after the certification is attained, the documentation, systems, and processes are often discarded or ignored until the next audit comes around.

Compliance is clearly necessary for doing business. But it's not what great security practices are built around.

Security isn't a set of checkboxes. It's literally a fight against a human opponent in a digital landscape who doesn't need to follow the same

attack pattern every time, who can evolve her tactics, and who doesn't play by the rules. The technologies being built and used across organizations are numerous, complex, and changing at an incredible rate. To stop attackers, security teams are charged with the task of securing all of them at all times. To stay ahead, it's imperative that companies empower their security teams to move beyond the checklist mentality of compliance and onto building a highly agile, innovative and, above all, attacker-focused defensive infrastructure.

Misconception #8: There's Nothing I Can Do to Stop Hackers

Amidst the many stories of fear, uncertainty, and doubt ("FUD") told around the security industry, you will find many saying things like, "There's no way to build a 100 percent secure system." The response that often accompanies this myth is, "If we can't build something 100 percent secure, why try at all?"

With an increasing number of data breaches being reported at the companies that should be the most secure, it's understandable why you can lose hope with your own security.

Some in the industry call "all or nothing" responses like this, security nihilism. A better question is "What percentage of my system is secure?" or more plainly, "What do I need to do to be successful at security?"

This sounds like a basic question but it's an important one for the future of building effective security programs.

So far, most of the industry judges the success of a security program based on whether you've been hacked.

If you haven't been hacked, you are successful.

If you have been hacked, you failed.

There are two flaws to this standard of judging security programs.

Flaw #1: The Gray Scale

Being hacked isn't binary. It isn't a matter of being hacked or not hacked. The reality is that it's a gray scale (Figure 1-4). What does it

mean to be hacked? It's easy to go down the rabbit hole, so I'll give you a few examples:

- You're running an ecommerce website with 100,000 accounts and an attacker takes over 1 of them. Does that mean your company has been hacked?
- One of your employees mistakenly sent an Excel file of all your employee HR data to the wrong person (or posted to an incorrect email list). Is that a data breach?
- A burglar stole a laptop of an employee but you don't know whether that person gained access to important files or data on the laptop. Is that a hack?

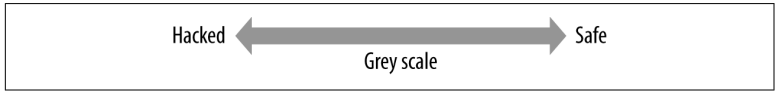

Figure 1-4. Hacking is not binary; most situations fall into the gray scale.

Would any of these examples be newsworthy data breaches? Unlikely. They all, however, represent some level of data breach.

Partial data leaks, inadvertent data sharing, and unknown data access are all part of the gray scale. Good security professionals understand the gray scale of hacking. Most nonsecurity professionals do not. But they need to if they want to track the effectiveness of your security efforts.

Is the Website Down or Not? What Security Can Learn from Web Operations

InfoSec needs to innovate and mature to improve data protection. Luckily, the recent evolution to the field of web operations provides very relevant learnings for security.

Web operations is focused on the stability, performance, and availability of web applications and services. Like security, most people outside of the operations team have been unaware of the challenges they face. In the past, the success of its work used to be simplified into a binary: is the website up or down? If the website is up and all services are available, success! If it's down and users can't access the site, failure. This success rating system is based on the false assump-

tion that it's possible to maintain a perfectly functioning system that never breaks. In practice, the site breaks regularly; sometimes on purpose for maintenance or system upgrades, sometimes because someone on the team made a mistake.

The industry evolved. Organizations have invested heavily in understanding the challenges web operation teams face and have developed solutions. Entire tool sets have been created to identify and fix problems as they arise. And more measurable and realistic goals have been set for the teams to encourage process changes and innovation.

Now, web operations monitor for problems and measure how fast they are identified, how quickly they're remediated, and how soon their users are notified. Even though consumers and management alike have now learned that it's unrealistic to assume systems are always available, they also now expect prompt and transparent communication when problems arise.

Information security needs the same evolution that web operations has had. And we don't need to reinvent the wheel. We too need improved tooling, process changes, and increased awareness of the challenges from the broader organization.

Flaw #2: No Knowledge of a Hack

Most companies don't know when they're hacked. The annual Verizon Data Breach Investigations Report report shows that up to 90 percent of companies that report and acknowledge that they have been hacked found out about the hack by an entity outside of their organizations. This means that if you asked people inside the company whether they had been hacked or not, the vast majority would have said they had not been hacked and they would have been wrong.

Let's look at the prior examples:

- Do you know when someone has an account hacked? Maybe the account owner self-reports it to you when he sees some odd activity. But how is the company supposed to know that happened without being told?
- Do you know when your employee accidentally sends sensitive data to the wrong people?

- Do you know what the burglar actually gained access to in a stolen laptop? And do you know what he did with whatever he gained access to?

Being hacked doesn't always mean that you have a massive data breach that you hear about in the press. The definition can be as simple as a single account being taken over. So, if you take the hacked or not hacked approach to measuring the success of your security team, the answer will mostly likely always be some form of yes, you have been hacked and you're failing. The real question is how does that form of measurement help your organization and how do you know if you're improving?

A Way Forward

Most security programs have been almost singularly focused on identifying potential vulnerabilities in their own systems before the attackers do. Although this isn't an inherently bad approach, technology changes are making it so that bugs and vulnerabilities are being created and identified much faster than their teams are fixing them. That means you have vulnerabilities that you know are there, waiting to be exploited. And the real question becomes: Do you know if and when they are actually being exploited?

This is why a new approach is necessary.

The most successful and modern security programs are focusing on answering—and measuring when possible—these three questions:

What are the areas in our organization that are susceptible to attack?

> If an attacker is looking to target you and your organization, where are the areas he could try to exploit? This is known as your organization's *attack surface*. To build a successful defensive strategy, you need to know what you're trying to defend in the first place. Identify the areas and organize them into categories like "network security," "application security," "physical security," and so on. There must be a process to review and update your organization's attack surface as it adopts new technology and/or discovers new exploitable areas. Your attackers aren't static; they're flexible and evolving. Your defenses need to be flexible, as well.

How successful are we at detecting when attacks on those systems are happening and how do we improve our ability to detect attacks?

After you have your attack surface identified, the next question is how do you know whether someone is attacking you there? No system is impenetrable, so what you need to be worried about most is developing a detection strategy that will alert your team when a security problem exists in the first place. The best teams invest in building both detection capabilities and also testing frameworks to continually measure and improve their ability to identify attacks in real time.

But how can you test and measure your ability to identify real attacks? Simulate real attacks!

You can simulate real attacks by either hiring a team of security analysts to attack your systems (commonly known as *pentesting*) or by electing a group of current team members to attack your own systems (commonly known as *red teaming*). The simple, yet critical, difference in this exercise is that instead of having these teams generate a list of vulnerabilities in your system as has been done in the past, you use their simulated attacks as a test of your defensive ability. How many attempts were your defensive teams able to detect and how quickly or effectively were you able to stop them (if ever)? This approach enables you to establish detection metrics based on what your defensive team identified versus what the offensive team was able to exploit (see Figure 1-5). Running new simulated attacks regularly provides you with a continuous testing framework for your attack defense.

How quickly are we able to minimize these attacks and remediate any problems that arise?

The last piece of a great defensive strategy hinges on how reliably and quickly you can fix problems that have been detected and identified. Because most attacks require a series of steps—known as an *attack chain*—to be successful, a good defender can identify the problems early in the chain and fix the flaws before they're exploited. The faster your teams can fix the problems they identify, the harder it will be for the attacker to succeed.

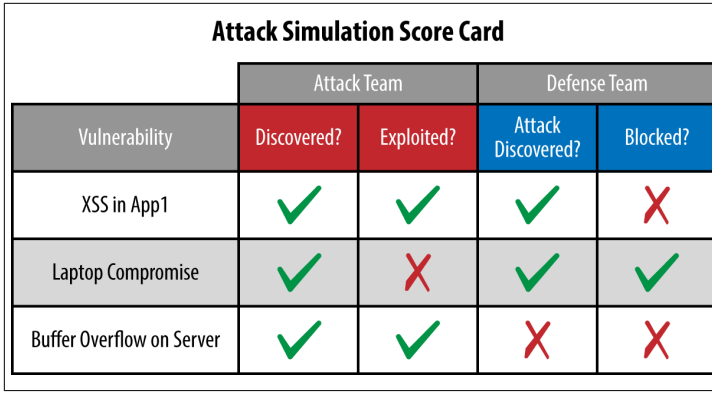

Figure 1-5. Create a scorecard to keep track of how you're defending attacks. And compare scorecards over time to measure progress.

Adopting this framework for building your security practice might mean that when you start measuring, your organization scores poorly. The goal is never to be perfect at security because that's truly impossible. Security is always a tradeoff between costs and risk reduction. But what this approach will do is give your organization a meaningful and measurable place to start. And, most importantly, it gives you a path to be able to track and improve.

Conclusion

At the end of the day, hackers will try anything to get at the valuable information your organization is protecting—be that by sending fake emails to the accounting department, stealing laptops from your sales team, or hacking into your engineer's code. If your organization wants to take security seriously you need to involve everyone in the company—security and nonsecurity professionals alike.

Now that you have an accurate and heightened awareness of the real security challenges facing your organization instead of the common misconceptions out there, you can be a part of the solution. Identify who is currently leading the charge and make sure they're asking the right questions, their bosses are setting measurable achievable goals for them, and they have the resources to achieve those goals and defend your data.

About the Author

Andrew Peterson (@ampeters06) is the CEO and Cofounder of Signal Sciences, an information security company based in Venice Beach focusing on attack detection and protection for websites and mobile applications. Previously he led a multidisciplinary product development group at Etsy in Brooklyn focused on global growth. Prior to Etsy he worked in Tanzania with the Clinton Foundation to improve the data quality of the Tanzanian National Health Information System. Before the Clinton Foundation, Andrew worked with Google's AdSense and AdX sales and product teams in Mountain View, CA. He holds a B.A. in Science, Technology, and Society from Stanford University with an emphasis on Human Computer Interaction through the Stanford d.school.