What Is an Advanced Persistent Threat (APT)?

We were at a meeting of Washington, DC, government officials and business leaders. A so-called consultant in cybersecurity (at least that's what his website said, and who are we to question the Internet?) spent half his presentation talking up the massive boogeyman of cyber danger that loomed for us all, repeatedly mentioning the new specter of "APTs." But fortunately, he spent the second half of his talk explaining how all that was needed to deter such threats was to be "good enough." He made a joke that it was like the two friends chased by a bear. As one told the other, "I don't have to

outrun the bear, just you." As long as you made sure your defenses were slightly better than the next guy's, he explained, the cyberattackers would give up and quickly move on. And, lo and behold, his firm had a generic package for sale that would satisfy all our cybersecurity needs. The presentation was slick, effective . . . and wrong.

APTs are "advanced persistent threats," a phenomenon that has gained more and more notoriety in recent years (Google reports the term as being used some 10 million times by 2013) but is still poorly understood. It illustrates the challenge in the policy world of calling attention to very real emerging challenges in cyberspace but also avoiding overreaction, hype, and hysteria.

If cybersecurity threats were movies, an advanced persistent threat would be the *Ocean's 11* of the field. It's not that APTs star handsome actors like George Clooney or Brad Pitt; indeed, they are more likely to be run by their polar opposites, clad in T-shirts instead of Armani suits. Like the high-profile heists in the movie, however, APTs have a level of planning that sets them apart from other cyberthreats. They are the work of a team that combines organization, intelligence, complexity, and patience. And as with the movie, they are quickly followed by sequels. No one knows how many APTs are out there in the world, but one cybersecurity firm CEO told us how, "Five years ago, I would get very excited, and very proud, if we found signs of an APT inside a client's networks. It was something that might happen once every few months. Now, we're finding them once a day."

An APT starts with a specific target. The team knows what it wants and who it is going after to get it. APT targets have ranged from military jet designs to oil company trade secrets. So while many of us would like to think that we are important enough to be targeted by an APT, the reality is that most of us don't rise to that level. But if you do, well, watch out; locking your windows like everyone else in the neighborhood probably isn't going to be enough. The bear in the sales guy's story actually doesn't care how fast your friend runs; it just wants to take a bite out of you.

The hallmark of an APT is its coordinated team of specialized experts, who each take on different roles. Much like a robber "casing" a bank or a spy observing a military base, a surveillance team engages in what is known as "target development," learning everything it can about the person or organization it is going after along

with key vulnerabilities. In this effort, online search tools and social networking have been a godsend to the attackers. Want to steal a widget and therefore need to know who the vice president of product development is? In the past, you might have sent James Bond to seduce the receptionist in human resources and then sneak into her files while she was sleeping off a night of shaken martinis and sex. Now, it's more boring. Just type the name into an Internet search engine and you can get everything from that executive's resume to the name of her daughter's pet iguana. As cybersecurity expert Gary McGraw notes, "The most impressive tool in the attackers' arsenal is Google."

It is this phase that also differentiates the attacks as "persistent." The reconnaissance and preparations can take months. The teams are not just trying to understand the organization of the target but also its key concerns and even tendencies. One APT, for example, was casing a major technology firm headquartered in Minnesota. Team members eventually figured out that the best way to crack the system was to wait until a major blizzard. Then they sent a fake e-mail about the firm changing its snow day policy; in Minnesota, this was something that everyone from the CEO on down cared about. Another effort, which American national security officials have blamed on Chinese intelligence and military units, gathered details not only on targets' key friends and associates but even what farewell they typically used to sign off their e-mails (e.g., "All the best" vs. "Best regards" vs. "Keep on Trucking") to mimic it for a spear phishing attack vector.

With the target understood, an "intrusion team" will then work to breach the system. What's notable here, though, is that the initial target is frequently not the main prize. An effective way into a network is via trusted outsiders, who often have lower levels of defense, or by targeting people in the network who have some access permissions to open the gates wider. For example, a series of American think tanks (including our place of work) were targeted in 2011 and again in 2012 by an APT that sought access to the accounts of scholars who worked on Asian security issues (they were interested not just in their files, but also their address books, which had the contact information for senior government leaders). But the APT initially went after other employees who had administrative rights and access to passwords.

These attackers frequently use spear phishing and faked e-mails, with some exploit hidden inside triggering a download of malware. In "Operation Shady RAT" (an APT we talk about later on in Part II), when the counterfeit e-mail attachment was opened, malware was implanted. This then created a backdoor communication channel to another outside web server that had, in turn, been compromised with hidden instructions in the web page's code, an effort by the attackers to cover their tracks.

The malware used in these attachments is often quite sophisticated. The emphasis is on stealth, so the authors will not only try to hide from traditional antivirus defenses, but burrow deep into networks and operating systems to avoid discovery, attempting to impersonate legitimate network traffic. Like other businesses, APT groups often conduct dry runs and even "quality assurance" tests to minimize the number of antivirus programs that can detect them. But e-mail is not the only way in. Other APTs have, for example, used networks like Facebook to find friends of individuals with a high level of privilege inside a targeted network. Then they compromise these friends' instant messaging chats to sneak in. Perhaps the most interesting example of this use of social networking tools saw senior British officers and defense officials tricked into accepting "friend requests" from a faked Facebook account that claimed to be Admiral James Stavridis, the commander of NATO. Who wouldn't want an admiral as a friend; imagine their disappointment when it turned out to be a hacker!

Once the team is in, they branch out like a viral infection, often with more personnel joining the effort. They jump from the initial footholds, compromising additional machines inside the network that can run the malware and be used to enter and leave. This often involves installing keystroke-logging software that tracks what people are typing and a "command and control" program that can direct the malicious code to seek out sensitive information.

At this point, the target is "pwned." Now at the mercy of the attackers, an "exfiltration team" works to retrieve the information the APT was targeting all along. Here is another APT hallmark: Instead of the usual criminal ethic of "Grab what you can get," they go after very specific files. Frequently, the attackers don't even open the files, suggesting that their reconnaissance was so thorough that they didn't need to review what they were targeting. Instead, someone draws

up a specific list of collection requirements and the team is disciplined enough to stick to it.

Not all APTs just copy and exit with the data. Some add technology to allow them to steal new secrets beyond what was inside the network or even gain control. French officials, for example, have accused APTs linked to Chinese intelligence of gaining access to the computers of several high-level French political and business leaders and then activating microphones and web cameras so that they could eavesdrop on conversations. Even more nefarious are those that don't simply steal data but also alter files, which as we explore later can have major consequences. This ultimately shifts the APT from an act of crime or espionage to an act of sabotage or even war.

The exfiltration phase, when massive amounts of data leave the network (such as when an entire e-mail file exits), is actually when many successful APTs are detected. This "phone home" phase makes for an anomaly in network traffic that is hard to mask.

Exfiltration teams therefore use all sorts of tricks to sneak the information out and then hide their tracks. One common tactic involves routing data through way stations in multiple countries, akin to a money launderer running stolen funds through banks all over the world. This not only makes it more difficult to track them down, but also routes the APT's activities through different countries and legal jurisdictions, ultimately complicating prosecution.

What makes APTs even more of a challenge is that when a target finally realizes it has been attacked, the pain is not over. Finding which machines inside the system have been infected can take months. Even worse, if the effort is truly persistent—say if the target has some sort of ongoing value to the attacker—there might be an additional unit in the APT whose very job it is to maintain an electronic foothold in the network. Rather than focusing on what information to steal, this unit will monitor internal e-mails to learn how the defenders are trying to get them out. In one case, an American company hired a Pentagon-qualified computer security firm to clean its infected network after being targeted by an APT. Despite this, a few months later, a thermostat and printer in its building were caught sending messages to a server located in China. With their e-communication compromised, the defenders' response is often to go old-school. They will literally yank hard drives out of their

computers and post handwritten signs in the hallways about password policy changes.

APTs are a nightmare scenario for any organization. Most don't know they've been targeted until it is too late. And even if they do find out, it is often impossible to prove who's behind it. Indeed, that's why APTs may be the most controversial of all the threat vectors. Except in cases where the attackers are sloppy (our favorite example being when a high-ranking Chinese military officer employed the same server to communicate with his mistress and coordinate an APT), there is little actual proof that would stand up in a court of law or sway a country's position. What we are often left with instead are suspicions and finger-pointing, which makes APTs so poisonous for diplomatic relations, as we've seen over the last few years between the United States and China.

How Do We Keep the Bad Guys Out? The Basics of Computer Defense

It is by far the world's largest zoo. In 2013, it held more than 110 million different species painstakingly collected from "the wild." And yet it is a strange sort of zoo, where you can't actually see the animals. The reason is that they only exist in the virtual world.

The McAfee "malware zoo" is what the computer security firm calls its collection of the various types of malicious or malevolent software (known as "malware") designed to wreak havoc on Internet users. Its growth illustrates the seemingly insurmountable scale of the problem. In 2010, McAfee thought it impressive that it was discovering a new specimen of malware every fifteen minutes. In 2013, it was discovering one every single second!

If we think of each type of malware as a unique threat, these numbers are overwhelming. Instead, we must understand why there are so many "unique" threats. The answer reflects the cat-and-mouse game that attackers and defenders play. Since the early Internet, attackers have tried to exploit vulnerabilities, and the defenders have sought to deny them. The adversaries, in turn, adapted and altered their patterns of attack, changing it into an evolutionary game.

The advantage of defending a computer system is that once you know what might attack you, you can just tell the computer

what to watch for and how to avoid it. Traditional antivirus software relies on detecting these "signatures." The programs scan all files on the system as well as incoming traffic against a dictionary of known malware, looking for anything that matches these signatures of malice.

This classic approach has a few glaring flaws. As the number of attacks grows over time, these definition files and the time it takes to search them also grow. Most of these old signatures don't represent current threats, and as the threats proliferate, it becomes a losing game. One study found that only 0.34 percent of signatures in common antivirus programs were needed to detect all the malware found in all incoming e-mail. And yet prudence dictates that we still must look for the 99.66 percent of the old malware, just in case the attacker gets sneaky and goes back to them.

The bigger problem is evolution. Authors of malware have fought back against the traditional antivirus approach by taking a page from biology. Just as some viruses such as HIV and the flu change their protein coatings to avoid detection by the human immune system, malware creators change the outward appearance of their attacking programs. The very same attack can be made into very different signatures, disguised by programs that automatically generate new features. This gives rise to enormous numbers, like those reflected in the "zoo" statistics cited above and renders the old approach of detection by signature less useful. One analysis found that, over an eight-day period, while over one hundred thousand new signatures of new known malware were added by a major antivirus vendor into the list of what to scan for, only twelve new detections resulted. Twelve detections for processing one hundred thousand signatures may seem paltry, but this reflects the camouflage techniques used by malware authors more than ineptitude on the part of antivirus companies. Inflating the number of signatures may make the malware problem seem bigger than it is, but it would be equally wrong to conclude that malware is not a problem at all.

Security vendors have thus had to change how they detect malicious code. Modern antivirus don't just screen, they use "heuristic" detections to identify suspicious computer code behavior based on rules and logical analysis. Static analysis breaks apart the computer code and looks for patterns associated with the behavior of an attacker. Virtual machines and other sophisticated

defenses dynamically simulate the code operation to determine whether the file examined will misbehave without putting the actual system at risk. Just as police bomb squads test suspicious packages, virtual "detonation chambers" can cause an incoming piece of malware to mistakenly think it is inside the machine and detonate prematurely.

If securing a modern operating system is difficult, an alternate approach tries to prevent the bad stuff from reaching the computer over the network. The simplest form of network defense is a "firewall." Taken from the concept of barriers built into cars or buildings to prevent fires from spreading, computer firewalls are like filters that reject traffic based on specific rules. Firewalls can prevent external computers from connecting to the firewalled machines except under preset circumstances or prevent certain applications on the computer from opening network connections.

Firewalls are filters that only permit valid activity on the network; the next layer of defense is a set of sensors that look for invalid behavior. "Intrusion detection systems" exist at the computer level or on the network. They detect attack signatures and identify anomalous behavior ("That's funny, the janitor's computer doesn't usually open an encrypted connection with Moldova at 2 in the morning"). These systems alert administrators to potential attacks and keep logs for detailed forensic analysis. Most detection systems now have some intrusion prevention capacity as well, which closes suspicious network connections and throws away anonymous traffic. Like antivirus software, these systems come with a price; in addition to typically costing more money, they cost time and performance resources inside the machine, especially if a system must evaluate all incoming traffic on a large network in real time.

While you heard earlier about a "zero day" that goes after a newly found vulnerability, most attacks attempt to exploit vulnerabilities that the vendor has already discovered and attempted to ameliorate via a code update or "software patch." The presence of a patch indicates that the vendor has found a vulnerability, identified a threat mitigation, and perhaps most importantly, built a fix into the existing code.

The problem is that many users don't always pay attention to these security updates and leave the vulnerabilities unpatched.

This why we have seen an evolution from the makers sending out simple notifications of new updates, which puts the burden of action on the user's part, to automatic downloads and even installation of the patch. But other costs have come along with this shift. Modern software is very complex, and fixing one small vulnerability could affect countless other processes in the program. Horror stories describe patches turning brand-new smartphones into rather snazzy paperweights or breaking entire enterprise networks. Since every large organization runs different software and has different configurations, patch management is an important part of any IT support.

Just as your body's defenses aren't just skin-level, serious threats are countered not only by keeping them out. There are measures to protect what's valuable even if threats get in. In the case of cyber espionage, for example, if you cannot prevent attackers from accessing the data, you can limit their ability to understand the data through encryption.

The last line of defense is akin to the strategy that nuns use to police Catholic school dances. The nuns often stuff balloons between teenagers dancing too closely, creating an "air gap" to ensure nothing sneaky happens. In cybersecurity terms, an air gap is a physical separation between the network and critical systems. Such practice is common with critical infrastructure, such as with power companies, and was even attempted by the Iranians to protect their nuclear research from cyberattack.

The problem with air gaps, much like the abstinence that the nuns try to enforce, is that it often doesn't work in practice. Giving up control of operational infrastructure involves sacrifices in efficiency and effectiveness. Power companies that don't link up, for instance, may be less vulnerable, but they can't run "smart" power grids that save both money and the environment.

Similarly, maintaining an air gap is often unrealistic, as the Iranians discovered when their supposedly air-gapped systems still got infected by the Stuxnet virus. At some point, old data needs to come out, and new instructions need to go in. Systems need to be patched, updated, and maintained. Indeed, the National Cybersecurity and Communications Integration Center has conducted literally hundreds of vulnerability assessments of private American business air-gapping attempts. Not once did it find an

operations network successfully separated from the firm's other computer enterprise networks.

Finally, some advocate the old slogan that the "Best defense is a good offense." In the cyber world, this is known as "hackback," and a number of firms have emerged to go after the attackers' own computer networks. Like vigilantism, it both feels good and may teach the original attacker a lesson.

But the hackback business has two major problems. The first is that the question of who has the "right" to carry out cyberattacks is unclear, which means that "cyber Blackwater" firms are "skating on thin ice" legally, says Nate Fick, CEO of the cybersecurity firm Endgame. The second is that it's not yet clear that hackback is even that effective over the long term. Alex Harvey, a security strategist for Fortinet, explains that "Breaking in and shutting them down isn't hard, but a new one will just pop up. You'll get a couple of minutes of peace and quiet."

The bottom line in cyber defense is that it is a hard task, with various options that are far from perfect. But the only other option is to close the zoo and let the malware animals run free.

Who Is the Weakest Link? Human Factors

In 2008, a US soldier was walking through a parking lot outside of a US military base in the Middle East when he spotted an unwrapped candy bar lying on the ground. Without knowing who had left it or how long the candy had been on the ground, he decided to take the bar inside the base and eat it for lunch.

Sounds absurd and even a bit disgusting, right? Well, substitute a USB flash drive for that candy bar, and you have the story of what started Buckshot Yankee, one of the largest cyber breaches in US military history. In 2008, a foreign intelligence agency left flash drives in the parking lot outside a US base, a tactic known as a "candy drop." A soldier saw one of the drives in the dirt, picked it up and thought it a good idea to plug the drive into a computer on the US military's Central Command network. The drive uploaded a worm named agent.btz that scanned computers for their data, created backdoors, and linked to command and control servers. The Pentagon spent the following fourteen months cleaning the worm out, all because one soldier didn't have the common sense to apply the "five-second rule" to how he treated his computer.

The real battle in cybersecurity is not just about high technology. It is also driven by the human factor, the fight over our behavior. To use a metaphor from the wars of old, the walls protecting a network don't always have to be knocked down or tunneled under by an attacker. Sometimes, as in ancient Troy, oblivious defenders simply open up the gates without recognizing the attacker's deception. While there are varied advanced threats in cybersecurity, many of the most successful take advantage of good old-fashioned human error. Examples range from the executive at an IT company who found a malware-ridden CD in the men's bathroom and popped it into his computer to see what was on it (again, think of the comparisons: would you pick up a comb you found beside the urinal? A sandwich?) to the employee at a defense company, who used his business network to file-share music. Besides sharing proprietary rock songs online, he also unintentionally shared designs for the electronics of the US presidential helicopter with Iranian hackers.

It is for this reason that many IT experts believe that if a network has any kind of sensitive information in it, all users need to be regularly certified in cybersecurity basics. This means everyone, from junior staff all the way up the leadership. This is not just about enforcing the five-second rule for things you plug into your computer. Clever attackers take advantage of our innate trust to convince us to click links, open attachments, or give out our passwords to strangers over the phone. Since 99 percent of the time our phone calls and e-mails are not malicious, it is hard to be constantly vigilant.

Even experts can and will be fooled. Professor John Savage founded the Computer Science department at Brown University in 1979 and has devoted much of his distinguished career to studying information security, including a stint advising the Department of State. Yet while teaching a class on computer security, one of his students successfully tricked him into clicking a link in an e-mail and entering his password. (We subsequently hired that very enterprising student as a research assistant for this book.)

We will go more into this in Part III, but the goal is to recognize the key part that human behavior plays in enabling threats,

and then build constant awareness, reinforcing it with new training. If users fail to learn the lessons of proper caution, then their access privileges should be revoked. Indeed, some companies like Lockheed Martin even have "red team" programs that every so often try to trick their own employees. If the employee opens a link in a suspicious e-mail, for example, it links the offender to a refresher course on cybersecurity. Better we learn our lesson this way than download real malware, Trojan Horses, or any other Greek-borne cyber gifts.