Defence Against Terrorism Review
Vol.3, No. 2, Fall 2010, pp. 23-36

ISSN: 1307-9190

From Hit and Run to Invade and Stay: How Cyberterrorists Could Be Living Inside Your Systems

Alan E. BRILL
Senior Managing Director, Secure Information Services, Kroll, Inc., Washington DC, United States

Abstract: *In recent years, there has been a distinct – and troubling – change in the tactics used by cybercriminals (criminals who have the technical skills to exploit weaknesses in digital technologies)and the cyber-elements (the terrorist organization members or those willing to work for them – for money or reasons or belief – who have the skills to exploit those same weaknesses) of terrorist organizations. They have gone from a hit-and-run mentality to one that stresses gaining and maintaining long-term access to the digital infrastructure of their enemies. These new tactics have worked, and worked well. But the tools of these cyber-invasions can be used in many ways. The same techniques that can gather intelligence can be used to prepare to disrupt or disable technology-based targets, ranging from factories to military communication systems. In this paper, we try to describe the evolution of the threat, and suggest ways to meet the challenges posed by the Advanced Persistent Threat scenario.*

Keywords: *Cyberterrorism, Cyberdefense, Hacking, Networks, Security*

Introduction

The ever expanding use of computer systems to process information – including sensitive tactical, strategic and financial information – not only by governments, but by corporations and individuals as well, has created incredible targets for cyber-skilled terrorists and criminals to exploit in order to fund their activities and to collect information to assist them in carrying out their plans.

One of the characteristics of this expansion has been the adoption of these technologies by significantly more small- and medium-sized governments and corporations. For example, the ability to initiate and approve wire transfers online has created an opportunity to steal literally millions of dollars in minutes.

But to best exploit these opportunities, there has also been an evolution in some basic concepts that information security people took for granted for many years. It was assumed that those stealing sensitive information or money were going to get into the target system, exploit it and get out as quickly as possible. While there are unquestionably incidents where the hit-and-run mentality applies, the ability to get into a system, hide there and gain intelligence, cause problems, exfiltrate data and hide evidence of the wrongdoing is now the more common approach.

What this means is that where cyberdefense has been traditionally focused on the enemy who is outside and trying to get into the system, the new reality is that with the use of zero-day techniques (exploiting security weaknesses in ways that have never been used before) and the delays in making defenses to newly-discovered problems widely available (through anti-malware updates) it is necessary to assume that the enemy is already inside of your system. You have to be able to recognize and defeat the attacker from within your environment as well as the attacker outside your system perimeter.

In addition, the expansion of the use of automated banking and sensitive information processing by small and medium organizations means that the number of targets available to the cyberterrorist or cybercriminal has grown. And those new targets may not have the experience, the staff or the budget to protect themselves as thoroughly as larger organizations or governments can.

If the evolution of the threat is not realized, the advantage can turn to the attacker,but through understanding and proper planning, that advantage can be countered.

The Worst Case Scenario: The Enemy Inside Your Systems

If you are considering the security of a computer network used to process highly confidential and proprietary information, what is the worst case scenario?

- *Destruction of the network?* Most networks have plans in place for disaster recovery and continuity of operations.
- *Denial of service to users?* True, so-called distributed denial of service attacks can use thousands of computers to send enormous volumes of requests to a system to overwhelm it and knock it out of production, but there are effective safeguards available to rapidly detect and defeat most attacks of this kind.
- *Network Intrusion?* Yes, this is bad, but it is usually recognized, and you can investigate the cause and strengthen your defenses.

I would like to suggest that the worst case scenario is an enemy who gains surreptitious access to your systems, who can stay below the radar of your defenses, and who can have access to your systems and your information for periods of months or years without being detected. A long-term intruder who can get in and stay in – and who you do not even know is there – is a nightmare scenario. But this nightmare is all too real. Consider the following:

> In 2008, the U.S. Department of Defense suffered a significant compromise of its classified military computer networks. It began when an infected flash drive was inserted into a U.S. military laptop at a base in the Middle East. The flash drive's

> malicious computer code, placed there by a foreign intelligence agency, uploaded itself onto a network run by the U.S. Central Command. That code spread undetected on both classified and unclassified systems, establishing what amounted to a digital beachhead, from which data could be transferred to servers under foreign control. It was a network administrator's worst fear: a rogue program operating silently, poised to deliver operational plans into the hands of an unknown adversary.[1]

The intrusion that U.S. Deputy Secretary of Defense Lynn is describing in that quotation went undetected for an extended period. In spite of extensive security measures, it spread through classified and unclassified networks;its operation over time was not caught. It was, perhaps, the ultimate infiltrator, providing valuable intelligence from an enemy who did not even know that it had been successfully attacked.

Looking at cases of intrusion across the public and private spheres, information security experts have seen an evolution in the thinking or adversaries. Where they were formerly content to follow a hit-and-run-and-hide strategy, their motive is now to get in, hide, and stay in target networks, prepared to exfiltrate data, or cause destruction at a time and of a kind (including the placement of false information, reduction or elimination of digital communication to forward deployed forces, etc.) chosen by the attacker. This is not to say that we have seen a complete shift to the new motive – far from it. Hit and run attacks to steal information, money or things of value (like credit card numbers) are still very popular. But anyone who ignores the new reality in which attacks can be very stealthy and very persistent is likely to find that the enemy is not just at the gates, but is already deep in their systems, watching, reading, and preparing for future actions.

Cyberattack Strategic Evolution

As data moved onto Internet-connected computers in almost unimaginable volumes in both the public and private sectors over the past two decades, hackers, criminals and terrorists kept pace. They understood the potentials for misusing the technology. As Scott Charney, formerly the head of the U.S. Justice Department's cybercrime section, now the Corporate Vice President for Trustworthy Computing of Microsoft, said in a television interview about hackers, "[a]t any given moment, there's a percentage of the population that's up to no good"[2] and he was right. Hacking and computer intrusions became issues that both governments and corporations had to worry about. Computer security became something indisponsible for governments and corporations moving into the new world of computers, and ultimately into the even newer world of the global Internet.

Starting out as just a technical issue of concern to information technology specialists, computer security has become something virtually everyone thinks about. From identity theft to the compromise of hundreds of thousands of credit card numbers, it has entered the global consciousness.

Over the years, as the technologies involved in computer networks and the Internet have evolved, so have the threats.

---

1 William H. Lynn III, "Defending a New Domain,"*Foreign Affairs*, Vol. 89, No. 5, p. 97.

2 Scott Charney, "Highway Robbery," *CBS 60 Minutes*, DATE.

Early threats were usually hit-and-run. Someone broke in, stole something – perhaps a blueprint or a set of documents – and left, attempting to leave behind as little evidence as possible. Perpetrators became skilled at manipulating internal system log files to cover their tracks.

But as networks evolved, new threats emerged, and some of them turned out to have the desirable characteristic (from the view of the attacker) of being able to be used over and over, either against multiple targets, or even against the same target. For example, as Internet-interfacing systems gained connectivity to traditional large-scale computer databases, the intruders discovered that many systems had a flaw that allowed them to manipulate the form of request that their computer sent to the Internet server of the target machine. It could contain instructions directed not to the web server, but to the back-end database. If the website's security was not all that it should be, the web server would pass the request along to the database, which could process the request and respond with data that it normally would not provide to an external user.

For example, in a recent case, a U.S.-based financial services firm was told that customer financial information was being misused for identity theft, and the pattern of victims showed that the data had to originate from their company. Our investigation showed that they had a number of customers on an old computer system that had been scheduled for replacement two years earlier. However, due to budgetary restrictions, the replacement of the system had been postponed multiple times. The old system still worked, but it only received attention from the programmers when there was a specific problem. It was a very low priority system from the company's viewpoint. But because it was obsolete and had not been updated in years – including its security – it was certainly a high-priority target for an attacker. The attacker targeted this system and used the technique described above – called a "SQL Injection Attack" to send questions to the database engine behind the website.

When you enter an address into a web browser – for example "www.google.com," you think of it as an address, but it can contain more. For example, if you go to the Google search engine and enter "cyberterrorism" as your inquiry, that is transmitted to Google as a web address that specifies the inquiry (along with other parameters). Google's servers can interpret this as a search request and process it appropriately. But in a SQL injection attack, the attacker is seeking access to data that is not intended to be released. By crafting inquiries that contain elements of Structured Query Language (SQL), a standard way of expressing data base inquiries, an intruder tries to send commands directly to the database, bypassing the web server and its security features. If the system is not designed with controls to prevent SQL commands from being executed, an intruder can start sending commands to see what information can be extracted.

By looking at the response to test questions, the intruder was able to figure out what data was available, and how to get to it efficiently. Over a period of days, the intruder sent tens of thousands of inquiries, all written with the SQL code embedded, and the system ultimately responded with the detailed and highly confidential data of more than 30,000 customers.

All of this was recorded on log files, but because this old system was regarded as being low-priority, the logs were never looked at by anyone.

Once we understood the attack vector (how the intruder was getting in – repeatedly – to the system and extracting data) we had to tell the victim company that the system could have been adequately detected against this kind of attack by the use of a free, open source web application

firewall program. The process of downloading, installing and tuning this program for optimum performance and putting it into live operation took one of our engineers working with the client less than 90 minutes to accomplish. While the attacks continued, they were no longer successful, and the ability to recognize the attacks in real-time was an important tool for tracking down the perpetrator.

In this case, the company never knew that there was a problem until confronted with it from an external source. Literally millions of dollars of damages could have been avoided had anyone taken the time to look at the software and assess its vulnerability.

What is vital to keep in mind is that this was not a one-time attack. The attacker was able to use the same attack against the same target over and over. The attack had the characteristic of persistence. The attacker wanted to be able to visit whenever he (or she) chose to in order to harvest more information. Granted, the actual method of attack – the SQL injection – was not very stealthy. In fact, everything that happened was logged, which enabled the investigative team to know with great certainty which records had been compromised and which had not. And it was not very sophisticated as these attacks go,but it had the advantage of working, and yielding tremendous returns for the time invested by the attacker.

Over time, the means of attacking networks in a way that permitted persistence of intrusion, have gained in sophistication and subtlety. These means became associated with attacks that appeared to be committed by nation-states, including hits against high-value targets (as in the attack described by Secretary Lynn in Foreign Affairs Magazine) and even was reported during the last U.S. presidential elections. According to well-known author Bob Woodward, both the 2008 Obama and McCain presidential campaigns had been hacked by very sophisticated hackers – perhaps even by nation-states. The attacks were detected and both campaigns were notified. They took steps to defend themselves in ways that they had not previously employed. According to Woodward, when then-candidate Obama recognized that the intruders could have done more than steal data – they could have destroyed data as well – he said "this is important."[3]

While Admiral McConnell, then serving as Director of National Intelligence was focusing his briefing on acts committed by nation-states, the same technical tools are available to terrorists and to those out for personal gain. Someone once pointed out that if you are shot through the heart, the age, sex, weight and motive of your attacker is irrelevant to you. Regardless of the answers to those questions, you are dead.

Admiral McConnell said that the U.S. intelligence community found the intrusion into the computers being used by the Obama and McCain campaigns because the intruders were "clumsy." Can we base our defensive strategy on the hope that the bad guys will always be clumsy or inept? What if they are not? What if they are really good at what they do?

Let us look at the latest evolution of the ongoing contest between cyberattackers and cyberdefenders. It's called the Advanced Persistent Threat – or "APT." While different writers disagree on how to define it (and even how advanced it is), it is a real problem, and one that those who manage national issues of defense against terrorism have to understand.

The objective here is not to make anyone a computer technician, but rather to provide an

---

3 Bob Woodward, *Obama's Wars*, Simon & Schuster, New York, New York, USA, pp. 9-10.

understanding of the motives, methods and problems associated with the latest evolution of the threat that we are facing, which may have enabled an attacker to read your emails, copy your documents and prepare your networks for disruption, even as you read this journal.

The APT Strategy

There is no globally agreed-upon definition of an Advanced Persistent Threat attack. In fact, writing in IT World's online publication, Kevin Fogarty has pointed out that "Advanced Persistent Threat has become the hot buzzword for an irresistible digital attack that should result in no blame whatsoever to the security, IT and business people involved – who, in fact, should get a raise and some time off for having endured such a harrowing experience."[4]

Fogarty is right. There is a tendency to use the term far too broadly. Traditional malware attacks are not by definition APTs. In fact, APTs can probably be better understood by members of the intelligence community than the technology community, since they use the traditional concepts of intelligence operations, aided, of course, by the latest in Internet and hacking technologies.

While APTs are closely associated with computer hacking, the perpetrators of APTs are willing and able to use the full range of intelligence resources,

They can conduct surveillance, both through online research and good old-fashioned feet-on-the-ground methods.

They use what the computer security community called "social engineering" – getting people to tell them things they should not. You might think that the fancy thumb drive you received is a gift from a vendor (after all, it has the vendor's name on it, and the accompanying letter certainly seems to be on their letterhead) but is it really from them, or is it a way to get you to introduce malware into the network? For example, in a number of recent incidents of fraudulent wire transfers involving the theft of millions of dollars, it was determined that the attackers used a two-stage strategy. In the first phase, they researched their targets to determine the names of financial executives and the bank the company used. They then called the bank and used what hackers call "social engineering," which is nothing more than getting an employee to reveal information that they should not divulge. They talked the bank's customer service people into giving them the credentials to enable them to access the online wire transfer system.

Once they had the credentials, they were able to access the bank's systems, get reports showing the numbers and balances of accounts, and when the moment was right, when the right amount of money was in the accounts, they initiated wire transfers, and millions of dollars moved around the world. According to logs maintained by the bank, the total time required to log into the system with the identity of an authorized financial officer, initiate the wire transfers, log in as a second authorized person, confirm the wires, and for the wires to be processed and completed was less than two minutes. In that two minutes, almost two million dollars was stolen.

---

4 Kevin Fogarty, "Advanced Persistent Threat is the Best Fake Excuse for Data Breaches" *IT World* (online) April 19, 2011, at http://www.itworld.com/security/157361/advanced-persistent-threat-best-fake-excuse-data-breaches (accessed April 27, 2011).

From the viewpoint of terrorists, using techniques like this one represents a way of getting tremendous amounts of funds quickly and efficiently, with little risk. By quickly transferring funds around the world, the trail that investigators must follow quickly grows cold.

Add to this that the cyberterrorist or cybercriminal can initiate all of this from anywhere in the world, and to the investigator, the attack seems to come from any of a million or more computers belonging to innocent people whose machines have been turned into 'zombies' by malware (malicious software) that enables them to be remotely controlled.

Cyberterrorists can put undercover agents in place, perhaps as employees, perhaps as vendors. For example, janitorial staff is in offices late at night, has a great deal of access, and is often minimally supervised. Not a bad disguise for an attacker to use to be in a position to put a USB memory device into a computer unwisely left running and unprotected to upload a virus into the network. Or they may walk in the front door as a temporary worker. We have seen cases where completely unvetted 'temps' (temporary worker) have been given access to highly sensitive information, particularly if it is discovered that the temp has a high degree of skill with complex graphics or spreadsheet programs.

Just because traditional hacking focuses on stealing data of value to the thief – like credit cards or antiterrorism plans – do not assume that trait fully defines what an APT is directed against; for example, a military, intelligence or political organization would have other data as the objective.

A summary of an article in the U.S. Air Force's Strategic Studies Quarterly[5] describes a possible scenario for a cyberwar in 2020. The nature of the war did not involve conventional acts, but involved disrupting military networks and injecting false information into the networks.[6]

In an interview with the author of the original article, Dr. Christopher Bronk, he downplayed the popular visions of an 'electronic Pearl Harbor' in which critical infrastructure, such as the electrical grid, is knocked out. Such attacks cannot be ruled out entirely, but it is unlikely that a nation-state would launch one because of the catastrophic response it would trigger, he said. Instead, Bronk said, cyberwar will be an effort "to get inside the other guy's decision making process rather than shutting it off entirely."[7] Of course, actions that a nation-state might hesitate to do for fear of retaliation might seem perfectly reasonable to a cyberterrorist who does not have the physical infrastructure of a nation-state to worry about. The factors that might demotivate a nation-state in regard to certain offensive actions may be totally irrelevant to the cyberterrorist. In fact, the cyberterrorist would probably hope that those defending a national infrastructure would believe that certain types of attacks would be less likely because of the potential for retaliation, and would do less to guard against them.

For those in the antiterrorism community, it is self-evident that you do not want any adversary – whether a nation-state or a cyberterrorist deeply embedded over a long time in your systems – able to steal plans and other data, enter false information and cause disruptions to your communications, command and control, logistics and other infrastructures.

---

5 Christopher Bronk, "Blown to Bits," *Strategic Studies Quarterly*, Spring 2011.

6 JaikumarVijayan, "What a Cyberwar with China Might Look Like," *Computerworld*, April 18, 2011.

7 Ibid.

For the rest of this paper, we will look at the anatomy of an advanced persistent threat attack, and suggest how cyberdefense strategy is evolving from a model focused on perimeter defense to a model focused on more defense-in-depth.

Anatomy of an Advanced Persistent Threat Attack.

The defining feature of an APT attack is persistence. Regardless of other motives – to steal sensitive data, to put tools into place to cause on-demand damage to a network, to be prepared to inject false or misleading information into an information system or whatever operational objective the perpetrator may have (and which can change over time) – the ultimate objective is to gain unauthorized entry into your network and to maintain that access over an extended period. It is not a hit-and-run attack. As the intruder, you donot want to be noticed. You want to fly under the radar and do so for as long as possible.

Unlike more traditional attacks which can be described in standard ways (for example, the SQL injection attack discussed earlier in this paper), the APT perpetrator can choose whatever means of infiltration that will work. For that reason, APTs frequently start with surveillance and intelligence gathering;the perpetrators want to know as much as possible about your network and your information security features as they can find out. They want to know the operating systems you use, the applications and database management systems that have been implemented, and as much as they can learn about employees and the possibility of infiltrating a terrorist sympathizer (someone who is not an actual member of a terrorist cell, but who is willing to act on their behalf – either because they believe in the aims of the terrorists, or they are willing to be employed by them) into your environment.

The more intelligence that can be gathered, the less you are an 'unknown.' The information that is collected can be of great value in figuring out how to best gain initial entry into your systems environment and then how to dig in for the long run.

Sometimes there is much more information available than a company may know. For example, a Google or Bing search may turn up information about your organization's networks in the form of press release from a vendor. We have seen cases in which a vendor white paper described an agency's key systems architecture (with a network diagram) and a description of the systems environment including operating systems and database management systems in use, and that information was used to enable a successful attack on the system.

For an intruder, time spent researching the target is time well spent. There should be no doubt that those responsible for the intrusion into the US defense systems described by Secretary Lynn did their homework and understood that using a specially-prepared USB device was, for them, an effective way to gain access to the initially targeted network.

Once a would-be APT perpetrator has completed planning, the actual work of the attack can begin.

The following scenario is very typical. While there will be countless variations, there is an anatomy to these attacks.[8]

Many of the attacks that we see start out as an email message. Of course, the email is carefully tailored to make the recipient feel safe in taking the action that the attacker wants. The user could be induced into opening an attachment which contains a computer virus or other form of malware. A user could be induced to click a link to visit a site that will – in addition to whatever else it has in response to the click -- upload malware into the targeted computer. (The automatic uploading of malicious software by simply visiting a suitably-built website is called a "drive-by infection.")

It could be something other than an email. Perhaps, as in the case discussed by Secretary Lynn, the adversary could use something like a USB device. We have seen cases where an adversary dropped a USB memory device attached to a key ring with a couple of house keys in the parking lot of a building. The perpetrators counted on the fact that a Good Samaritan would find it, and in an attempt to be helpful and find out who lost it plugs it into a computer. Seconds later, the network is infected. Some people believe that APTs always use previously unreported security defects (called "Zero-Day Attacks") but this is not the case. For high value targets (including sensitive government targets) it is worth it to an attacker to use a zero-day attack – if they have one available. But obtaining a zero-day defect is not easy and can be costly, so attackers will often target a known security defect that may not have been patched, or rely on 'social engineering' to get an authorized person to do something foolish.

Once the email attachment is clicked, or the employee visits the linked site, or the USB device is installed, the malware immediately takes control of the employee's computer, establishing a digital beachhead.

The malware works to establish a connection via the Internet to a server where it can send information and from which it can get instructions (sometimes called a 'command and control site.') Once this communication channel is in place, a human attacker uses a 'back door' established by the malware to enter the compromised computer. Once inside the compromised system, the attacker can take a stealthy look around, and can typically determine the privileges the machine has on the network, what else is on the network, and what part of the organization the compromised machine belongs to. The key to doing this with a degree of stealth is using tools that are normally on the computer that you are attacking and that are normally used by authorized systems administrators. A systems administrator seeing the use of the 'netstat' or 'nbstat'tools, for example, is unlikely to immediately think of an intrusion, as they are normally used for authorized purposes.

Once the intruder has gained access and looked around, the typical next step is to try to get the passwords of authorized users. Of course, virtually all systems maintain passwords in an encrypted form, but that does not faze the attacker. Using one of a number of hacker tools designed to find and steal encrypted password files. (An example of this is the program 'pwdump.') Most of the time, the intruder finds it easier to export the encrypted password list, and process it on a computer – or a network of computers controlled by the attacker.

---

8 The anatomy of an attack is in part derived from Christopher Day, "An Approach for the Detection of the Illicit Use of Legitimate Network Access Credentials by an Intruder," presented at the American Academy of Forensic Sciences 63rd Annual Scientific Meeting, Chicago, IL, February 24, 2011, Session B5.

So now the attacker has the encrypted password file. Assuming it is a strong encryption system, at first glance, it would appear to be virtually impossible to crack the code and get to the original passwords. Unfortunately, this is not the case. Through the use of what are called Rainbow Tables, attackers can often begin identifying passwords within minutes. While the exact way that these work is complex, the idea is simple. Assume for example, that your secret password is "curiosity." When the password is processed through the password encryption system, it comes out as "3v8qr@7dps^4". What if I now take a list of every word in the dictionary, and run it through the same encryption engine used by the system. Somewhere on the resulting list is going to be an entry which says, in effect, "curiosity" = "3v8qr@7dps^4." When I match "3v8qr@7dps^4" from the compromised laptop with the same entry in my table, I know that the matching password must be "curiosity." It is more complex in practice, of course, and there are ways of defending against Rainbow Table attacks, but the reality is that in most cases these defenses arenot in place, and the attacker will quickly succeed in recovering useable passwords. If any of those belongs to a privileged user who has more rights than a normal user (for example a domain administrator) all of those rights are compromised, we have a happy attacker, and the attack continues.

Using the compromised passwords, the attacker moves through the network to access other systems, and to compromise them. Typically, the attacker will choose to compromise a number of network devices and will bring in various tools to create multiple entry points into the network (which the intruders want, because if you find and close their original entry point, they will have many others to permit them to continue to compromise the network). Having multiple entry points also protects against the day when the original user changes their password. That one may not work, but with multiple entry points and compromised passwords, there are always more that will still do the job for the intruder.

When this is accomplished, the intruder can enter the network at will, and if the system provides remote access (using supposedly secure tools like Virtual Private Networks, webmail or enterprise application portals, the intruder can freely use them through their compromised entry points and credentials. This is typically the point when the attackers are secure in their access, and they can focus on their objectives (stealing documents, planting false documents, erasing log files showing what they have been doing, building code structures to render a network unusable or whatever other tasks are set by the attacker's controllers.)

In fact, at this point, it is fair to say that the bad guys are not actually hacking into the system. They have what the systems sees as valid credentials, and they simply log in like any other user to do what they want. It is the adversary's ability to get to the point of masquerading as a valid user that makes APTs so difficult to detect and eliminate.

So now they are in your system and they have probably deployed some additional tools to maintain access. For example, if you find them and close down the malware they are using, they may have planted a little program deep in your system that mostly does nothing. But once every few weeks, it wakes up and checks to see if you are still infected. If you are, it goes back to sleep. If not, it attempts to start the infection chain again – after you believe you have eradicated the problem. Some other malware is designed to re-start itself whenever the computer it is infecting is re-started (These are examples of just how persistent these attacks can be.)

Seeing an ATP in Your Network

Let us assume that for whatever reasons, your perimeter defense did not detect the initial penetration of your network and that no one noticed as the intruders moved through your system to establish a persistent presence. How can you determine if you have been compromised?

While there are no magic answers – new intrusion and new defense techniques will have evolved between the time this is written and the time you read it – but here are some ways that systems security specialists work to detect and defeat these attacks.

- Look for the communication between the malware and those controlling it. This command-and-control traffic can often be recognized either because it is operating through unusual communication ports or because traffic is going to places where you would not expect it to go. Log analysis can help you to see periodic traffic, often indicative of this type of traffic, as the malware 'beacons' or 'calls home' at scheduled intervals. If you detect internal systems sending "here I am" beacons or listening for traffic on unusual ports, is a strong indicator of a persistent attack.
- Look for files and what a security expert could identify as software tools placed in locations that do not make sense relative to the normal operation of the system. Seeing files out of place is a strong indicator that they are being moved as part of a scheme to eventually export them. If you find a cache of executable programs that you do not recognize, particularly in a strange location (for example in the recycle bin) you should be very suspicious. Clearly, having intrusion detection tools and file integrity tools in place in your network is important.
- Using log file analysis (which assumes you have comprehensive logging and that you maintain the logs for a long enough period) you may detect that legitimate credentials are being used in parts of the system where you would not expect them. Activity outside of normal expected ranges should be investigated.
- You can also analyze how valid credentials are being used. Are sessions using valid credentials originating from unexpected places? If an employee is known to be on vacation in the Canary Islands, it would be very suspicious if the person's credentials were being used from somewhere 10,000 miles from the Canary Islands. It is also suspicious when credentials are used at odd times, or where one credential appears to be addressing resources that the person normally would not use. Similarly, if one user credential is sequentially used for sessions one hour apart, but the sessions originate in Washington, Istanbul, Sao Paulo and Vladivostok, you have an intrusion (or an employee who can travel at speeds that Superman would envy.)

Ultimately, most of the analysis that you can do is based on having the right logs. Maintaining Windows event logs, IIS (Internet Information Server) logs, firewall logs, web server logs (Apache logs as an example), Syslogs and Windows RAS (Remote Access Server) logs (and their MAC, Linux, or other operating system equivalents) as well as logs relating to specific applications– and securing them as soon as there is a suspicion of problems (so that they cannot be erased to cover an intruder's tracks) is important, and there should be protocols for locking down systems where necessary. Storage space for logs has become inexpensive. Having logs that are detailed and which are held for a long enough period to be useful is now practical in almost all situations.

Given the potential for an intruder to remove log files to destroy evidence of what was done, consideration should be given to having logs written to locations for which the logging application or system has read-only access. Alternatively, any of a number of log aggregation systems can be used to move logs to a protected environment. Information released about the SONY PlayStation Network intrusion indicates that the intruders erased log files, and that made the investigation more complex.[9]

Toward a Defense in Depth Strategy

We live in a world where the software we use is so complex and consists of so many lines of computer code (sometimes millions of lines of code) and where the pressure to quickly release new versions of programs to satisfy customer or competitive demands is so great, we can no longer assume that the code we run is secure. We know that intruders have also developed techniques that are specifically designed to defeat the controls at the perimeter of our networks. The best firewall will not help if an employee can plug in an infected USB device brought from home. Building a strong perimeter is important, but it isnot enough.

Those organizations that are doing the best job of defending their networks recognize that network defense is a complex and always changing problem. They also realize that implementing a defense-in-depth can be expensive, in terms of technology and of people.

Here are some of the defensive measures to consider. This is not an exhaustive list by any means. And the selection of the right measures for use in a particular situation must flow from the sensitivity of the data, the likelihood that it will be targeted, and the architecture of the network and physical environment within which it is running.

- Control the end-points. You cannot have good security if you do not exercise some control over the equipment your people use. If you provide desktop or laptop computers, as well as smart-phones or tablet computers to your staff, you have a right to control what runs on them. The best advice is to lock them down. Do not let users add on any software without specific permission (based on need and testing the integrity of the proposed software) and donot let them plug in unauthorized USB devices. They will not like you for these decisions, but without them, the job of securing the network becomes much harder. Consider, for example, that even if you lock down the central network, without end-point control, information stored on a laptop computer, for example, can be at risk from malware on the local machine.
- Harden your devices. We see many cases in which an intrusion was made possible by having the wrong settings on a server, router or firewall or even an end-user machine. Organizations like the US National Institute of Standards and Technology provide a number of very useful guides to setting up network devices in a way that improves security.[10]

---

9 In a letter sent on May 3, 2011 to the Subcommittee on Commerce, Manufacturing and Trade in response to U.S. Congressional inquiries, Kazuo Hirai, Chairman of the Board of Directors of Sony Computer Entertainment America said “Among other things, the intruders deleted log files in order to hide the extent of their work and activity within the network.” (Page 4)

10 The NIST Security Configuration Checklists can be accessed at no cost at http://checklists.nist.gov.

- Make sure your logs are turned on, are creating detailed records and are saved – securely – for a sufficient period.
- Make sure vulnerable applications are protected by Application Level Firewalls. These programs filter the information that is passed to an application to avoid an intruder's attempt to smuggle in inquiries or commands disguised as regular transactions. Programs such as WebKnight (which is open-source) can be easily installed and configured, and are often a practical solution to protecting an application against SQL injection attacks.
- You should have the appropriate monitoring solutions in place for your sensitive networks. From simple intrusion detection systems to more complex intrusion prevention and data-leakage prevention systems, to more comprehensive packet-capture and traffic monitoring/analysis technologies, it has become vital to design the right level of surveillance into sensitive systems. It's also necessary to have people trained to monitor those security systems, and who have the proper training and experience to interpret alerts, quickly react to problems, and evolve the protection appropriately.
- Keep training your people. While there's no 100% protection against an employee or contractor reacting to a falsified email or other human error, good training can certainly diminish the risk.
- Know your people. Everyone who can access a network with sensitive data should have the appropriate background checks. This not only applies to employees, but to contractors, temporary workers and vendors as well.

Conclusion

There are no easy solutions to the highly sophisticated persistent threat attacks that every organization faces. Terrorist organizations in particular are highly motivated to establish persistent access to the networks of adversaries, targets and those who defend against their attacks. Understanding the problem, and taking the right defensive actions is vital if we're to stay one step ahead of those attacking our networks.

References

Day, Christopher, "An Approach for the Detection of the Illicit Use of Legitimate Network Access Credentials by an Intruder", presented at the American Academy of Forensic Sciences 63rd Annual Scientific Meeting, Chicago, IL, February 24, 2011, Session B5.

Bronk, Christopher, "Blown to Bits,"*Strategic Studies Quarterly*, Spring 2011,pp. 1-20.

Fogarty, Kevin, "Advanced Persistent Threat is the Best Fake Excuse for Data Breaches" *ITWorld Online*, at www.itworld.comprint/157361(accessed April 27, 2011).

"Highway Robbery," 60 Minutes, Scott Charney (guest) CBS, WCBS-TV, New York, 1998.

Lynn III, William H., "Defending a New Domain," *Foreign Affairs*, Vol. 89, No. 5, 2010.

Woodward, Carl, *Obama's Wars*, Simon & Schuster, New York, NY, USA, 2010.

Vijayan, Jaikumar, "What a Cyberwar with China Might Look Like,"*Computerworld*, April 18, 2011.