InformationWeek :: reports

Next

Reports.InformationWeek.com March 2012

Detecting and Defending Against Advanced Persistent Threats

APTs are a growing problem for enterprises big and small. Protecting your organization from these targeted threats requires constant vigilance, ongoing employee training and a concerted effort to align security systems to address every phase of an APT. Companies also need to develop a remediation and response plan if, despite best efforts, defenses are breached.

By Michael Cobb

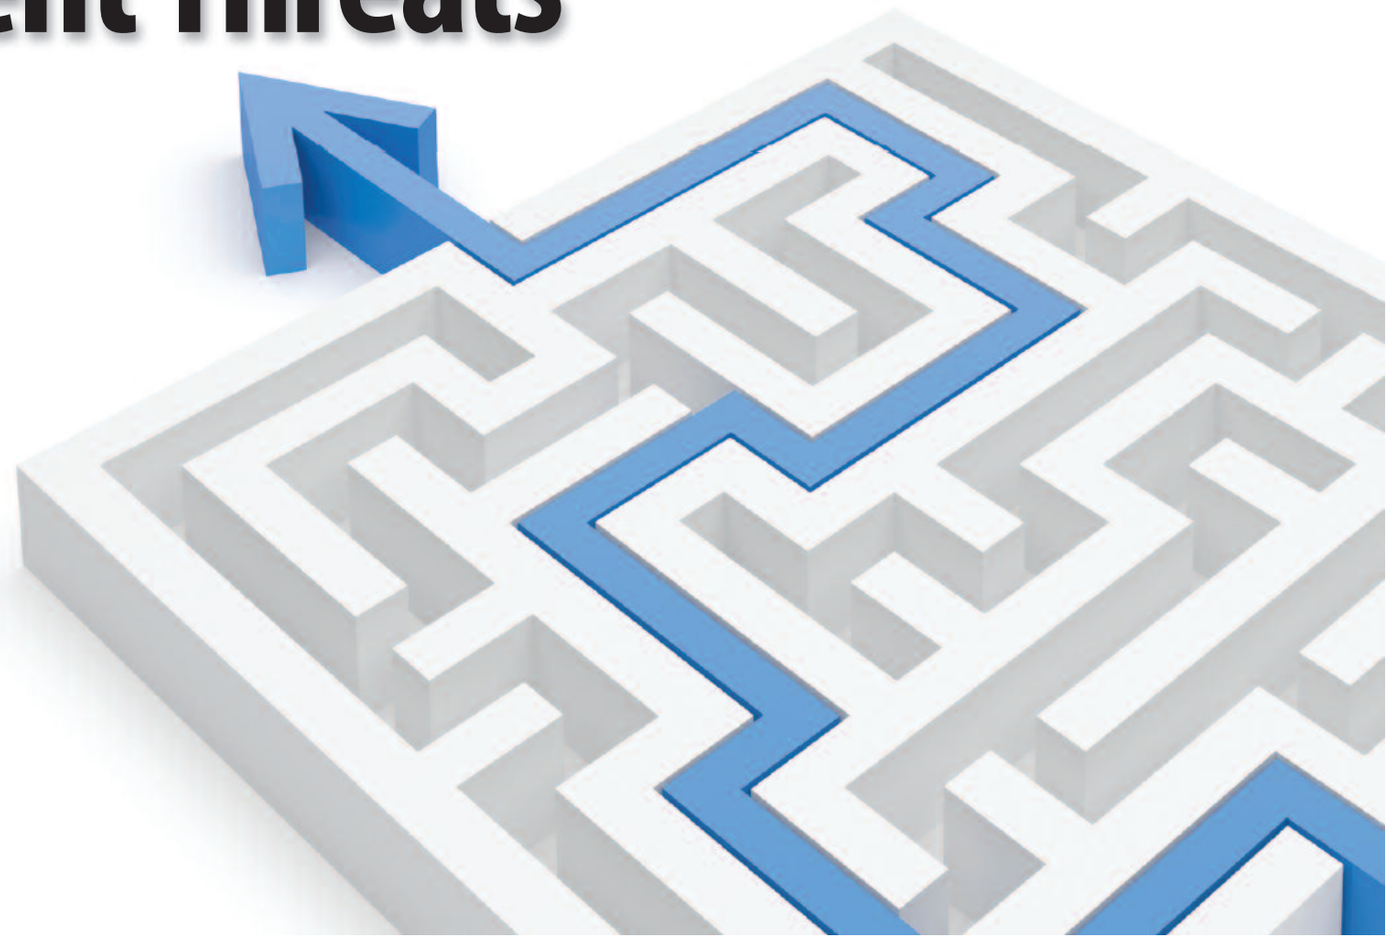

Presented in conjunction with

SECURITY dark READING
Protect The Business Enable Access

Report ID: S4390312

TABLE OF CONTENTS

ABOUT US

InformationWeek Reports' analysts arm business technology decision-makers with real-world perspective based on qualitative and quantitative research, business and technology assessment and planning tools, and adoption best practices gleaned from experience. To contact us, write to managing director **Art Wittmann** at *awittmann@techweb.com,* content director **Lorna Garey** at *lgarey@techweb.com,* editor-at-large **Andrew Conry-Murray** at *acmurray@techweb.com,* and research managing editor **Heather Vallis** at *hvallis@techweb.com.* Find all of our reports at *reports.informationweek.com.*

Michael Cobb
InformationWeek Reports

Michael Cobb, CISSP-ISSAP, CLAS, is a renowned security author with more than 15 years of experience in the IT industry. He is the founder and managing director of Cobweb Applications, a consultancy that provides data security services. He co-authored the book *IIS Security* and has written numerous technical articles for leading IT publications. Michael is also a Microsoft Certified Database Administrator.

EXECUTIVE SUMMARY

Advanced persistent threats are exactly what their name says: The threats are advanced, requiring a high level of expertise to develop and pull off, and they are persistent, lying in wait for just the right opportunity. There is certainly no silver bullet for preventing and defeating APTs—they are a new attack methodology built to circumvent current perimeter and endpoint defenses. However, training employees, providing robust user credentials, and hardening servers and workstations will help stop the spread of malicious code. Robust logging will increase the chances that APT-related activity is detected and will certainly give emergency response teams better information for identifying and remediating the attack. In this report, we examine the six phases of an APT and recommend how to protect your company from this growing problem.

Detecting and Defending Against Advanced Persistent Threats

One of the most insidious cyberthreats is the one that lies in wait. These exploits, commonly known as advanced persistent threats, are sophisticated, custom exploits with the express objective of gaining access to a targeted system and remaining undetected for an extended period of time. An APT's success requires considerable resources and expertise—hence the term "advanced." "Persistent" doesn't mean a continual barrage of attacks launched in the hope that one may succeed, but instead the relentless pursuit and development of a successful attack methodology. These exploits are developed by skilled, motivated, organized and well-resourced programmers working with a well-defined road map. These attacks can take many months to develop and even longer to successfully deploy. In this report, we will dissect an APT—from reconnaissance to data extraction—to identify where your organization might be vulnerable and how you can close the gaps.

You've likely been hearing a lot about APTs lately. Hydraq and Stuxnet, as well as the recent attacks on RSA SecurID and Lockheed Martin, would be classified as APTs. However, Conficker and attacks by hacker groups Anonymous and LulzSec would not. Conficker did not target a particular organization, and while attacks by Anonymous and LulzSec had specific targets, the groups made little or no effort to keep the attacks under wraps.

The concern about APTs is increasing because of an escalating number of incidents and the severity and extent of the

Figure 1

Advanced Persistent Threat Life Cycle

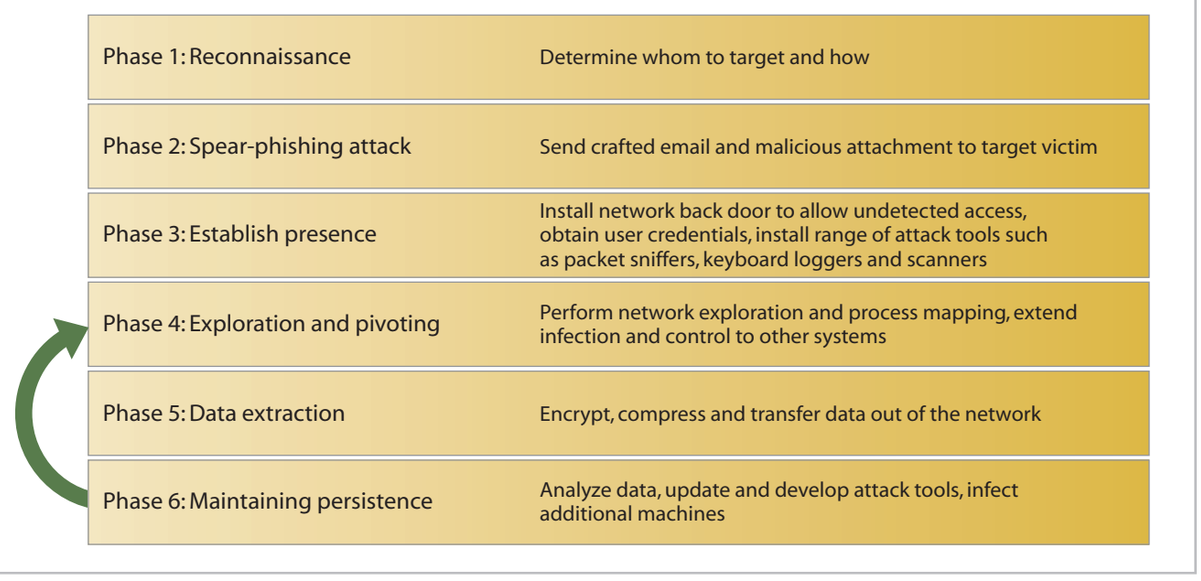

Data: Michael Cobb

damage they cause. Cisco Security Intelligence Operations has reported a significant increase in the number of unique instances of malware it's finding, an indication of APTs under development or being deployed. And although big and well-armed companies such as Google, RSA, Sony and Lockheed Martin have been hit, there are signs that APTs may be going after smaller and less-well-protected organizations to get to their eventual targets.

The nature and perceived goals of the APT attacks that have been uncovered so far point mainly to government involvement. For example, the goal of the Stuxnet computer worm was to disrupt Iran's nuclear weapons development, while information stolen by other APTs has correlated with the need for intelligence related to upcoming corporate or government negotiations and acquisitions, or as a means of accessing resources of national interest.

To combat the threat of APTs, it's important to understand the different phases of an APT attack and the defenses required for each.

Figure 2

Defenses Against Advanced Persistent Threats

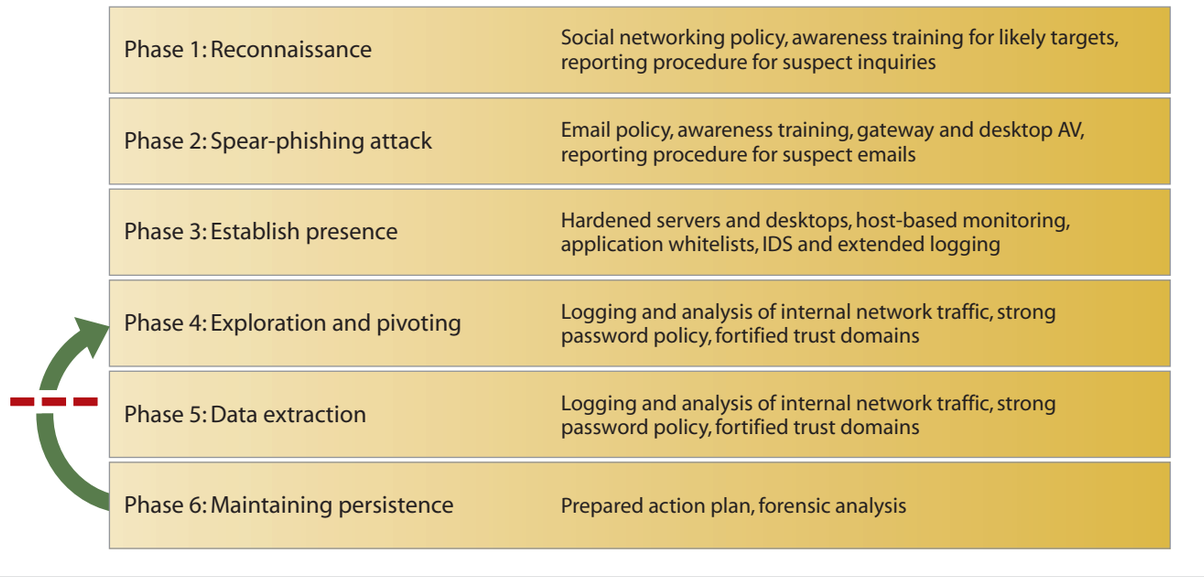

Data: Michael Cobb

Phase 1: Reconnaissance

During the first phase of an APT, the attacker searches for a weak link among an organization's network and data systems.

The most common technique for introducing malware into a victim's network is a phishing campaign. Instead of trying to break through network perimeter defenses, it's easier for attackers to focus attention on company employees. Infected email and email attachments are still very successful methods for gaining an illegal presence on a machine. Communication via social media is also being used to successfully gain a

foothold in the organization.

To make a malicious email appear real and thus fool the recipient into opening it, potential victims—usually senior executives or employees with access to sensitive data or high-level privileges—are researched using publicly available information found on Google, Facebook, Twitter, LinkedIn, and other social networks and company websites. Increasingly, geolocation is also being used to help build a picture of victims. Such data provides insight into not only a person's whereabouts but also habits and lifestyle. Armed with this type of information, an attacker can more easily script a personalized email that references the victim's recent activities. This greatly increases the chances that the victim will assume the email is from someone he or she knows.

All of this underscores the importance of warning employees about the dangers of overpublishing their lives on the Internet. Public sources of personal information enable an attacker to craft very convincing emails, often referencing recent company meetings and referring to colleagues the recipient knows. Attackers may also have other sources of intelligence, such as telephone interceptions and government records. It's important to keep staff informed of the latest phishing techniques being used and to make them aware of how their personal and geolocation data can leak into the public domain and be misused.

Phase 2: Spear-Phishing Attack

Many employees, even those who are savvy about security, have been tricked into opening an email or email attachment or following a link. As seen from the reconnaissance phase, it can be very difficult to identify an APT spear-phishing email. Making matters worse, a hacker may also send spam to hide the real payload among the general noise of unwanted emails to help divert attention from the real attack.

Stopping these attacks in their tracks requires constant vigilance by employees. Your main defense is security-awareness training that covers how a phishing attack works and how and why certain employees may be targeted. Employees should be trained and retrained in safeguards such as checking to make sure that an email comes from a trusted party before opening it. Yes, these kinds of practices make electronic communication a little less instantaneous, but they also make it a lot safer. With the rise in the use of social networking, employees should also be trained to exercise caution with both their personal and business accounts. It's important that this type of training is ongoing, to address the churn we have seen in privacy and security processes on public social networks such as Facebook. Employees should also be trained to report any suspicious emails so that network surveillance can be stepped up.

It's also important to keep email gateway filters, as well as intrusion detection and protection systems, up to date so they can spot known malicious patterns of system and network behavior. The attack on RSA SecurID, for example, involved a customized remote administration tool, known as Poison Ivy, embedded in a Microsoft Office document. Poi-

son Ivy had already been used extensively in many other attacks.

Other recognized attack vectors include supply chain compromise. (For example, RSA was targeted in order to successfully infiltrate defense contractor Lockheed Martin.) Check service-level agreements covering the management of (or access to) your information processing facilities by external parties. You should also take reasonable steps to check that affiliated third parties are implementing any required security measures. It is good practice to re-evaluate agreements with third parties to ensure that they cover any new risks and that their staff is aware of your particular security requirements.

Phase 3: Establish Presence

Once an attacker manages to infect a victim's machine, the process of installing a full range of attack tools and exploring the network begins. Organizations often fail to identify this type of activity because perimeter defenses and traffic analysis tend to concentrate on inbound traffic. Meanwhile, the attack tools used are often common system administration utilities, installed using valid credentials, so they don't get flagged by host-based antivirus scans. APTs tend to use common network ports 80 and 443, with the majority also using encryption. Process injec-

Figure 3

Advanced Persistent Threat Data Extraction

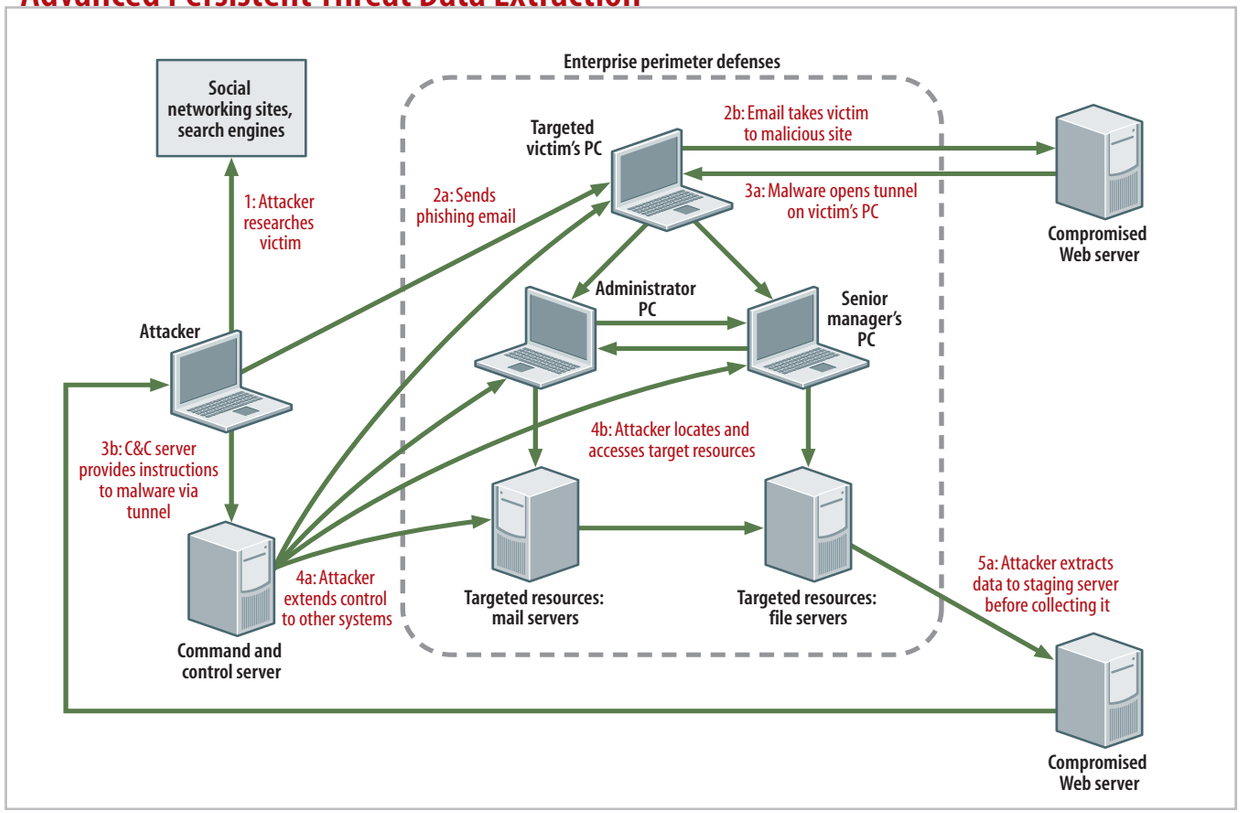

Data: Michael Cobb

tion and service persistence are also used to allow surveillance to go undetected.

Ideally, you should be monitoring internal and outbound traffic because the installed attack tools will need to make outbound connections at some point. Even though APTs are custom-developed, they usually include code snippets from previously identified malware that can be caught by analyzing internal network traffic and host-based information such as registry entries. The security tools in use today are typically up to this task, but security professionals need to make sure they are following up on alerts, even to the point of code disassembly and analysis.

The concern about APTs is increasing because of an escalating number of incidents and the severity and extent of the damage they cause.

Phase 4: Exploration and Pivoting

With the right tools in place, an attacker can begin to analyze the layout and map the defenses of the network, checking for the existence of vulnerabilities and programs that may be of help. On Windows machines, the registry will certainly be examined, as will processes and open ports. To help with analysis, the attacker will most likely take screen shots, dump password hashes, capture traffic and perhaps even tap a built-in microphone. The data gleaned from these activities is then examined to determine how best to proceed: quietly continue to capture more data, looking for additional access points and vulnerable systems, or install additional tools to try to take over a specific resource. Analysis of compromised systems by security organizations including Mandiant and SANS shows that multiple tools are often used in parallel and that the tools usually mutate to avoid detection.

But even if the tools change, the goal never does: to secretly acquire information to enable access to better-protected data. The process of using one compromised system to bypass firewalls and other restrictions to attack other systems on the same network is called "pivoting." Once pivoting has been put into place, an attacker may have more control over the system than the system administrator does.

IT security teams need to identify critical host and network resources because these will be any attacker's primary targets. Data of different classifications should certainly be kept in different security domains, but, to make life even harder for an attacker, different types of data should be held on different servers. For example, research and development files should not be kept on the same file server as plans for a major acquisition.

Extensive logging is vital for spotting and tracking down an infection. For example, a zero-day attack used as part of an APT to compromise the initial host might go undetected, but having network sensors logging and analyzing internal network traffic for unusual activity can trigger an alert that may in turn uncover an infection.

To enable your security team to fine-tune monitoring and alerting for a particular network segment, a baseline of well-known, expected behavior must be established. This way, unusual and potentially malicious traffic

can be more easily detected. For example, port scanning or egress traffic initiated by an internal host that isn't in response to an external request can be flagged as an immediate indicator of unexpected and suspicious behavior. SQL statements that are 10 to 20 times longer than normal are another obvious indicator of malicious activity—usually, an attempted SQL injection attack. In any case, organizations need to know what is normal to identify what is abnormal.

If an attack is detected, the last thing you want to do is panic and rush into an unplanned response.

It's important that all events are logged for all systems. Organizations should consider implementing logging on all DHCP and DNS servers, as well as on VPN concentrators. In addition, make sure that Windows application, system and security event logs are appropriately sized and logging. Host-based antivirus and intrusion prevention programs should be set up to log events, and all internal traffic should be logged. Centralizing key logs makes integrating them into a security information event management, or SIEM, system a lot easier. With that said, any centralized log storage should be thoroughly secured.

It's a good idea for any company (but especially for companies susceptible to APTs) to schedule a new security audit to reassess risks to its data. An audit will reveal potential gaps in existing protections and highlight the need for additional security controls. These may include database firewalls that feature SQL grammar analysis and application whitelisting controls that can, for example, prevent custom tools from executing and stop valid tools such as netstat from being run on nonadministrative desktops.

Finally, keep up to date with security advisory services so you know what you need to be looking for.

Phase 5: Data Extraction

For attackers to be able to use the data they have been so carefully collecting, they need to get it out of the network. This phase of the operation—data extraction or exfiltration—is the most challenging for an APT because data has to travel in one form or another out of the network. Attackers use a variety of tricks to accomplish the task, including peer-to-peer networks, encryption, onion routing applications and steganography.

Data-loss-prevention technologies can make the extraction process a lot more difficult for an attacker. DLP rules can be set to send alerts based on the characteristics of outgoing traffic and stop certain file types from leaving a given network segment altogether. However, because these attacks are very sophisticated, additional outbound traffic analysis is essential. In the same way that attackers are continuously monitoring and probing your network to achieve their objectives, you to have continuously audit the traffic leaving your network to ensure its validity.

Attackers may make use of innocuous-sounding domain names as staging posts, or fast-flux to circumvent defense mechanisms such as IP-based access control lists. This makes it more difficult to identify the

attacker's own network, but just identifying the existence of such traffic in DNS logs could indicate compromise and possible exfiltration.

Because these attacks are carried out over an extended period of time with lapses between malicious activity, centralizing device logs helps SIEM technology find correlations among disparate and non-consecutive events. Many SIEMs can also place devices on a watch list based on a predefined rule set that automatically classifies certain behaviors and increases scrutiny of the suspect host. Solutions that provide visualization of captured traffic can quickly highlight infected devices probing network neighbors for other systems that the attacker is looking to exploit.

Phase 6: Maintaining Persistence

It takes time for an APT attack to fully explore and remotely manipulate a system, which means that an attacker's presence must remain hidden from standard defenses for many months. Developing the code for the next stage of the attack takes time because the code has to be rigorously tested to ensure that it will work and that it will not reveal the attacker's presence. This requires patience on the attacker's part and comprehensive and regular analysis on yours to try to catch the bursts of activity that occur as the data is sent back to the attacker and updates are installed.

Remediation Planning

If you implement defenses for each phase of an APT attack and tackle the weak points that enable it to persist, you will reduce the chances of any attack successfully establishing itself on your network. However, you need to be prepared to act if your defenses are breached. If an attack is detected, the last thing you want to do is panic and rush into an unplanned response. With APTs, you can't simply install the latest AV signatures and scan infected machines to stop further infections. However, you also can't afford to take too long to develop a remediation plan. If you do, the infection will spread, becoming more damaging and time-consuming to fix. This means you need a preplanned response. A senior manager, preferably at board level, needs to be in charge and accountable for the remediation plan. This will ensure the cooperation and coordination of all key stakeholders. It will also ensure that sufficient resources are allocated to the emergency response team, or ERT. Being prepared means having up-to-date documentation of the enterprise infrastructure, including lists of all DNS and DHCP servers, Internet points of presence, VPN concentrators and Windows domains, including the group policies enabled on each Windows workstation.

The remediation plan should be executed once the full scope of the compromise is understood and the team has identified all compromised APT hosts. The ERT should be made up of trained security personnel who understand the network infrastructure and are familiar with how it looks when it is operating normally. This will make identifying indicators of compromise and separating the APT from other malware quicker and

easier. You may well need specialized outside help if you don't have the in-house personnel, technology and ability to dissect the tools and techniques used by APT intruders.

If you've identified your organization as a potential target, you need to start a discussion among your industry peers to determine how best to coordinate efforts against this cyberthreat. It may even be necessary to establish contact with government agencies such as the U.S. Cyber Command, which is tasked with coordinating the U.S. military's response to APTs.

Because of the serious nature of an APT infection and the potential for it to spread to your partners, it's important to take the bull by the horns and come clean if a data breach occurs. Keep clients, stakeholders and regulatory authorities up to date on the situation. While embarrassing, these attacks are becoming part of the IT landscape, and sharing the experience can only help others become more resilient.

MORE LIKE THIS

Want More Like This?

InformationWeek creates more than 150 reports like this each year, and they're all free to registered users. We'll help you sort through vendor claims, justify IT projects and implement new systems by providing analysis and advice from IT professionals. Right now on our site you'll find:

Database Defense: The biggest threat to your company's most sensitive data may be the employee who has legitimate access to corporate databases but less-than-legitimate intentions. And while the incidence of insider data breaches has decreased, external attacks often imitate them—and do serious damage. Follow our advice to mitigate the risk.

Understanding Software Vulnerabilities: To protect company and customer data, we need to determine what makes it so vulnerable and appealing. We also need to understand how hackers operate, and what tools and processes they rely on. In this report, we explain how to ensure the best defense by thinking like an attacker and identifying the weakest link in your own corporate data chain.

Four Steps to Virtualization Security: From access controls to separation of duties, we present four key steps to managing risk in your virtual environment.

PLUS: Find signature reports, such as the *InformationWeek* Salary Survey, *InformationWeek 500* and the annual State of Security report; full issues; and much more.