Breaches Happen: Be Prepared

A SANS Whitepaper

Written by Stephen Northcutt

October 2014

Sponsored by
Symantec

Introduction

Computer viruses are yesterday's news; automated attacks that morph rapidly, concealing themselves through encryption and deceptive packaging, are the new hotness. Such camouflage makes purely signature-based detection obsolete and gives malware an advantage. At the same time, marketing companies, popular websites and governments are collecting vast quantities of web surfing and computer usage data, with new techniques such as "canvas fingerprinting," a method to distinguish individual users that uses a web browser's canvas element (a feature of HTML5) to identify individuals for phishing and other information-based attacks without the use of cookies.[1]

Another option for delivery of malware is targeted, web-based attacks made possible by commercial-grade attack kits. The Google transparency report summarizes the phishing and malware sites, derived from the experiences of approximately 1 billion users of Google Safe Browsing.[2] According to it, the most dangerous are the sites hosted by Korea Telecom; of the almost 500,000 KT-based sites scanned, 84 percent hosted malware.

A third vector for the advanced threat is a window of opportunity, such as that found between the time a vulnerability is discovered and when affected software can be patched. For instance, in July 2014 Microsoft had to issue an emergency patch to handle compromised SSL certificates, Adobe had to issue a patch for a Flash vulnerability, and scans conducted in June 2014 indicated almost 310,000 servers had yet to be patched for the well-publicized Heartbleed bug, two months following the disclosure of the bug's existence.[3]

The fourth and final vector we must consider is the universe of "thinglets": the subsystems that make up the so-called "Internet of Things." Thinglets may have their own CPU, memory, and firmware and can present unforeseen challenges to investigators. For example, current tools are unable to scan the firmware of a USB drive; attackers can change the behavior of such a drive to emulate a keyboard, bypassing any limits on drive attachment set by security policy.[4] Alternatively, one might accept only USB drives that are AES-256 encrypted, but as security researcher Chris Brenton points out—using research from SySS—the vendors may not take the implementation of encryption seriously, doing foolish things such as using the same encryption key repeatedly.[5]

[1] "Pixel Perfect: Fingerprinting Canvas in HTML5," Keaton Mowery and Hovav Shacham, 2012; https://cseweb.ucsd.edu/~hovav/dist/canvas.pdf

[2] www.google.com/transparencyreport/safebrowsing

[3] "Troubling Trends: Many Websites Still Not Patched for Heartbleed Security Bug," SpiderOak.com, July 3, 2014; https://spideroak.com/privacypost/cloud-security/troubling-trends-many-websites-still-not-patched-for-heartbleed-security-bug

[4] "BadUSB: Big, bad USB security problems ahead," *ZDNet*, July 31, 2014; www.zdnet.com/badusb-big-bad-usb-security-problems-ahead-7000032211

[5] "DLP & Encryption: Are They Mutually Exclusive?"; www.chrisbrenton.org/wp-content/uploads/2010/01/encryption-dlp-keynote.pdf

Organizations trying to protect themselves against advanced persistent threats by relying on stateful firewalls and traditional signature-based antivirus defenses don't stand a chance. As Eugene Kaspersky pointed out last year to the Canberra Press Club, "All the data is stolen. At least twice."[6]

So, in addition to antivirus and firewall technologies, IT security practitioners need a mix of tools as cited in frameworks such as the Critical Security Controls (CSCs) or the Control Objectives for Information Technology (COBIT).[7] They should begin by implementing well-understood best practices, starting with endpoint hardening to remove existing malware as well as close and manage vulnerabilities. Even then, they ought to have a plan for detection and response strategy if a breach should occur.

This paper describes how to start with improved malware reporting and gateway monitoring and how to combine this output with security intelligence from both internal and external resources. Intelligence comes in many forms, including reputation services and even honeypots that use valid email addresses to detect phishing. Forward-thinking organizations use these and other techniques promoted by frameworks such as the Critical Security Controls. The key is to—as quickly as possible—detect hostile activity, identify and locate affected systems and devices, and respond appropriately.

[6] "Kaspersky Claims Stuxnet Infected a Russian Nuclear Plant," *The Inquirer*, Nov. 11, 2013; www.theinquirer.net/inquirer/news/2306151/kaspersky-claims-stuxnet-infected-a-russian-nuclear-plant

[7] The Critical Security Controls: www.counciloncybersecurity.org/critical-controls; COBIT: www.isaca.org/COBIT/Pages/default.aspx

Protect and Respond

The longer an IT infrastructure is under attack, the greater the cost of recovery. One case that still boggles the mind was with the TJX breach in 2007, which exposed more than 45 million card numbers because attackers went unnoticed for 18 months before detection. That's not atypical, either: In a 2013 SANS survey on security analytics and intelligence, multiple respondents reported up taking up to 10 months to detect intruders, with the average being one week or less; the next-largest group did not even know the duration of attacks discovered in their networks.[8]

The experience of TJX and others was far better than that of DigiNotar, a security company that provided the Dutch government with digital certificates; it lost more than 500 certificates in a 2011 attack and ultimately filed for bankruptcy after industry heavyweights Google, Microsoft and Mozilla blacklisted all certificates issued by the company.[9] Unfortunately, the entire trust model for e-commerce is under stress from compromised certificate authorities (CAs), stolen and fraudulently obtained certificates, and fundamental problems with cryptography, highlighting just how difficult it is to detect and protect against costly breaches.

Detection Is Key to Prevention

To stay in business, organizations must detect fraud, attacks and malware as quickly and accurately as possible to respond decisively. The goals are to:

- Contain the infection
- Clean affected systems and recover data
- Repair the vulnerabilities discovered during response

[8] "SANS Security Analytics Survey" (September 2013), pg. 8; www.sans.org/reading-room/whitepapers/analyst/security-analytics-survey-34980

[9] "DigiNotar Certificate Authority Goes Bankrupt," *NetworkWorld*, Sept. 20, 2011; www.networkworld.com/article/2181294/security/diginotar-certificate-authority-goes-bankrupt.html

Know What You Have

Being prepared to detect and respond to attacks and attempted attacks starts with knowing your environment, no matter how complex, as described in the first two Critical Security Controls. Getting full visibility into your environment is not as easy as it sounds. Automated tools such as Nmap provide some visibility into devices, systems and users on the network, but they may fail to recognize other entry points such as:

- Wi-Fi networks
- Virtual server instances
- Rogue web applications with access to the data center
- Printers or other devices with network access

Most importantly, when assessing your network, its applications and endpoint systems, it is vital to know what sensitive data is stored, transmitted and processed across the environment—in particular, sensitive data that an attacker would want to exfiltrate from the protected network or data center. When visualizing your network, systems and applications, be sure to include known, trusted sensitive data uses and pathways to help you detect anomalous behavior indicative of data leakage. Endpoint security systems providing reports on access, data downloads, encryption and other data-related policies can also help feed this intelligence for more accurate detection and response.

Fleshing out a view of the threat landscape may require combining data from a variety of sources and is an ongoing effort for any IT security team.

Suspect Everything Else

If an active SIEM installation has not produced actionable information in the past 30 days, its administrator should review its configuration.

The result of this audit becomes the core of the organization's whitelist. Any attempt to reach any hardware, software or URL that is not on the authorized connect list should be flagged as suspicious. This flagging of the unknown will create a massive input that is best collected, processed and prioritized with a security information and event management (SIEM) platform. However, it is important to note many organizations with such systems don't use them properly and fail to tune them to changing conditions, according to the SANS Analytics and Intelligence Survey. If an active SIEM installation has not produced actionable information in the past 30 days, its administrator should review its configuration.

The system should also be looking for known-bad traffic, such as communications between systems that have no business communicating, unusual bursts of outbound traffic, and other signs of infection that should be detected at various ingress and egress points, endpoints and elsewhere.

Looking for Needles in Haystacks

Abnormal traffic can indicate suspicious activity; organizations should search traffic for unusual behavior such as:

- Traffic that isn't using ICMP, TCP or UDP
- Traffic to any port that is outside your network activity baseline
- Traffic (other than file transfer) with large packet sizes
- Inbound and outbound encrypted packets (hiding commands or payloads)
- Packets that say they are one size but are something else
- Connections to or from locations for the first time
- Traffic to or from the Tor network
- Systems with unexpectedly degraded performance
- System activity when the authorized user is not present
- System activity that doesn't match the normal application or user profile
- Multiple changes to DNS records within a 24-hour window
- Results from a DNS sinkhole

When attackers target your organization, it is at significant risk; the likelihood of a threat finding its matching vulnerability is high. Obvious weak spots are hostile or oblivious insiders, whether employees or contractors, who use their access inappropriately. Classic examples of insider incidents include the 2012 trade secrets breach at Toyota North America and a 2013 customer data compromise at Vodafone Germany.[10]

The danger presented by rogue insiders is why detection is a critical first step in preventing and reducing the impact of attacks.

[10] "Toyota Says Fired IT Contractor Hacked into Company Secrets," *Automotive News*, Aug. 27, 2012;
www.autonews.com/article/20120827/OEM06/120829918/toyota-says-fired-it-contractor-hacked-into-company-secrets;
"Vodafone Accuses 'IT Insider' of Data Theft," *Contractor UK*, Sept. 13, 2013;
www.contractoruk.com/news/0011226vodafone_accuses_it_insider_data_theft.html

Intelligence Is Key to Response

The next step is using what you know about your network, along with detected event information and external intelligence to organize your response. Unfortunately, having such a plan in place—and an incident response (IR) team to support it—is not as common as it should be. According to the 2014 SANS survey on incident response, 43 percent of survey respondents admitted they lack formal IR plans, while 55 percent lack a formal IR team.[11]

One of the biggest barriers to having such a plan, according to the survey, is lack of people and processes. Once the remediation process is underway, including outside help can go a long way to identifying what went wrong and how to prevent future incidents. Honest assessment of vulnerabilities requires the ability to face up to mistakes and the will to prevent any repeat of those mistakes.

Golden Rules of IT Security

SANS Senior Fellow Dr. Eric Cole has suggested four simple "golden rules of security" for companies wishing to improve their IT defenses:[12]

1. **Know your systems.** Companies must understand their systems and what they are trying to protect.
2. **Follow the principle of least privilege.** Servers, devices and users must receive the least amount of access they need to do their jobs and nothing else.
3. **Defend in depth.** No single measure is going to make an organization secure; it must deploy multiple levels of protection.
4. **Always try to prevent, but never fail to detect.** Organizations must build their networks in such a way that, in cases where they cannot prevent an attack, they can detect it before it causes major damage.

Teach Paranoia

If something smells fishy, it probably is. Wise security practitioners focus on making their detection tools work better for them. Although the jury is out as to the value of security awareness training, it is another layer to help in detection of events that automated systems may not collect and collate. The Transportation Security Administration's slogan "if you see something, say something" makes sense. Every employee, especially system and network administrators, can be the eyes and ears of the organization.

However, staffing the skills needed to sniff out and respond to incidents requires a curious person with the ability to write pattern-matching regular expressions; admins with experience using Perl or other scripting languages are ideal candidates for this role. Don't load them up with too many operational duties; instead, let them look for patterns using automated pattern-matching tools, the SIEM system, packet sniffers, endpoint monitoring tools, access points and whatever else your response team uses to detect, scope and remediate threats.

[11] "Incident Response: How to Fight Back" (August 2014), pg. 2; www.sans.org/reading-room/whitepapers/analyst/incident-response-fight-35342

[12] "New Riders Interview with Eric Cole," Pearson Education website; www.pearsoned.co.uk/bookshop/article.asp?item=435

When your people find something that needs attention, reward them. Several organizations, including SANS and the Naval Surface Warfare Center, Dahlgren Division use an "On the Spot" award. It doesn't have to be a lot of money; if it buys a steak dinner for two, that's enough to be meaningful. To really make it work, present it with an email of praise that will make that employee's mother blush. The combination will motivate employees to find something else that presents a hitherto-unknown threat traipsing around the network, and, with luck, other system administrators will want to get in on the act.

Utilize Technology

Analysts and automated tools have a lot of "noise" to sort through. Security gateways of every kind are valuable tools for both protection and detection, as they represent a choke point for traffic both normal and hostile, sending alerts to administrators and feeding data to the SIEM system as directed. In one form or another, they have been in use for years: from the original proxy firewalls focused on lower-level networking protocols, to unified threat tools (such as antimalware, antivirus or IDSes) to modern web application firewalls that examine traffic at the application layer. Reducing the raw data from these tools to the bits that actually represent a threat requires careful filtering and even more careful baiting of traps for attackers.

Filtering

One key skill is knowing how to screen the data being collected to reduce false positives, to identify targets and to stop the attack. In an interview for this paper, TJ O'Connor, author of *Violent Python*, explained how, as part of the 2010 NSA Cyber Defense Exercise, cadets from the United States Military Academy whom he coached created some expert proxies for detecting anomalies in content passing through them:

> *By removing some of the more active content from potential malicious files, the cadets were able to retain the functionality of the original content but still render any potential malicious pieces neutral … One of the more simple checks was to remove JavaScript that automatically loaded inside a PDF document without user intervention.*

O'Connor's West Point team also expanded the checks to verify that executables weren't backdoored, PowerPoint presentations didn't contain macros and HTML pages didn't contain cross-site scripting. This shows a good use of inquisitive, talented individuals and technology to filter out the known bad from the known good.

Advice from the Honeynet Project

Several technical controls are available from The Honeynet Project.[13] One may also consider a few simple tricks:

- Include a fake acronym in some of your key business documents that you can scan for with your IDS/IPS/DLP/NGFW.
- Plant an incorrect name or email address for your comptroller, since attackers commonly target that position.
- Base a few bogus profiles around email addresses with corresponding Facebook or Twitter accounts identifying them as members of the organization. This can be a great way to look for sophisticated custom phishing.
- Create a file labeled "Reduction in Force Plan August 2014" or "Bonus Proposal December 2014" and track access attempts. These titillating titles can help identify untrustworthy insiders with privileged access.

Observing

Honeypots and related tools, such as "honeycode" (hidden in applications) "honeynets" (groups of honeypots) and "honeytokens" (e.g., dummy email accounts or files used as bait), can help organizations improve their response to detected events as well as improve detection of future, similar attempts.

Foremost, the "honey" has no business use, so investigators can consider any activity in such tools hostile. Of course, this approach will generate false positives, simply because with a large enough data sample, one can expect every possible combination to occur from time to time. Nevertheless, they also provide early detection of an advanced threat and a safe environment to observe its attempted behaviors. Many intelligence vendors use honeynets and sensory technologies to detect and examine malware types in the wild, enhancing their detection products or services with the newly gained understanding of malware behaviors.

Share and Reuse

Retailers, financial service providers and airports belong to industries that have come under consistent—and persistent—attack. These and other industries have learned, firsthand, the value of sharing information.

Consider joining and participating in your industry's information sharing and analysis center (ISAC); if another member of your industry shares a signature or detection method, that's actually one more tool in your toolbox. When your organization does the same, the whole industry can strengthen its defenses.

This is a practice that many "security intelligence" vendors are also taking, deriving information from industry groups and from their own sensors and honeynets, relaying the indicators of new attacks to their clients. This area is growing but not well utilized; according to the 2013 SANS analytics survey, 32 percent of respondents utilized internally gathered intelligence to enhance their detection, while nearly 40 percent used no external sources of intelligence information.[14]

[13] www.honeynet.org

[14] "SANS Security Analytics Survey" (September 2013), pg. 12; www.sans.org/reading-room/whitepapers/analyst/security-analytics-survey-34980

Protection and Defense Cycle

Prevention, defense, response and repair are the basic tenets of the Critical Security Controls (CSCs), and all must take place to provide an intelligent response, as shown in Figure 1.

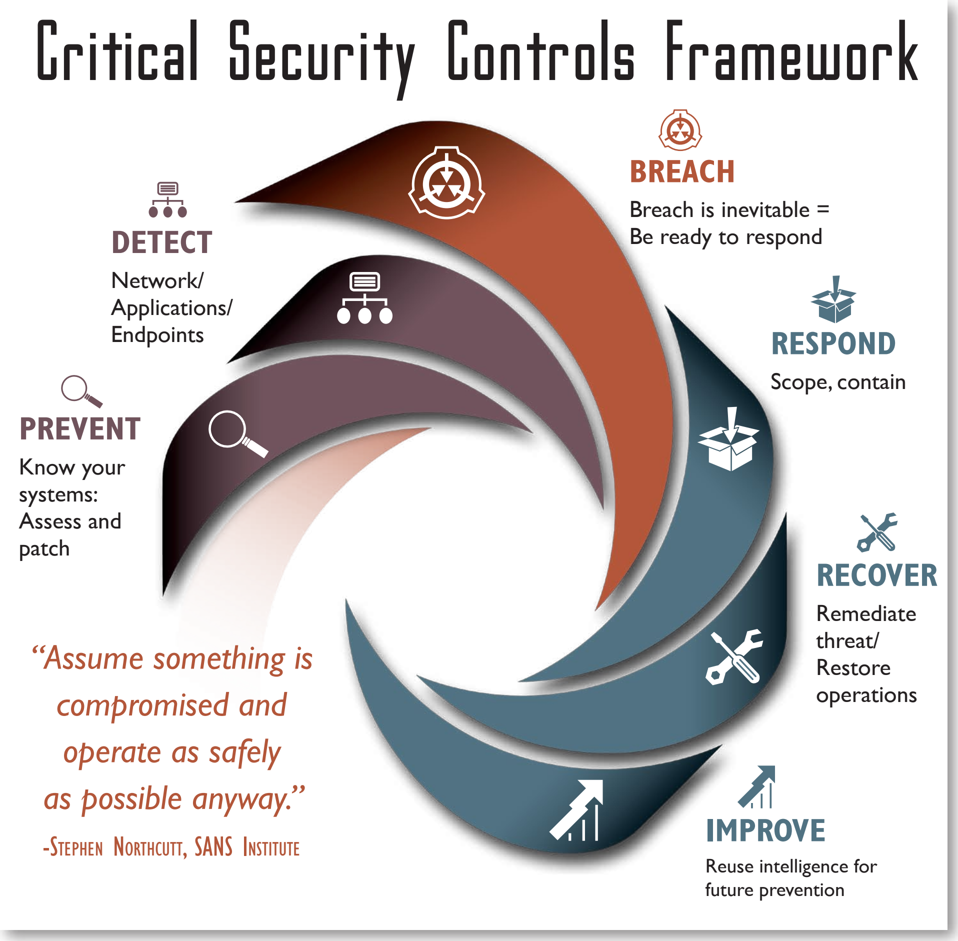

Figure 1. Prevention and Response Cycle

Ultimately, a coordinated system for prevention and response should be self-learning, so that incidents and the vulnerabilities they exploited should not be a problem in the future. Intelligence gathered through this cycle can also lead to improvements in the risk management program. For an example, take this hypothetical scenario of an automated intrusion response:

Although the IT administrators hadn't realized it, ACME Medical Center had been under attack for a week; what finally sounded the alarm was that the egress monitoring system noticed a lot of encrypted traffic traveling out of the network, originating from a device not associated with a specific user. The response team searched through the medical center's SIEM system, correlating network activity with the host of that data. From that, they determined that a single authorized user accessed the data host six times, all within normal business hours. An examination of the user's endpoint system concluded that the endpoint was sending the purloined data to a network printer—the system from which the outbound data originated.

Figure 2 depicts this theoretical response.

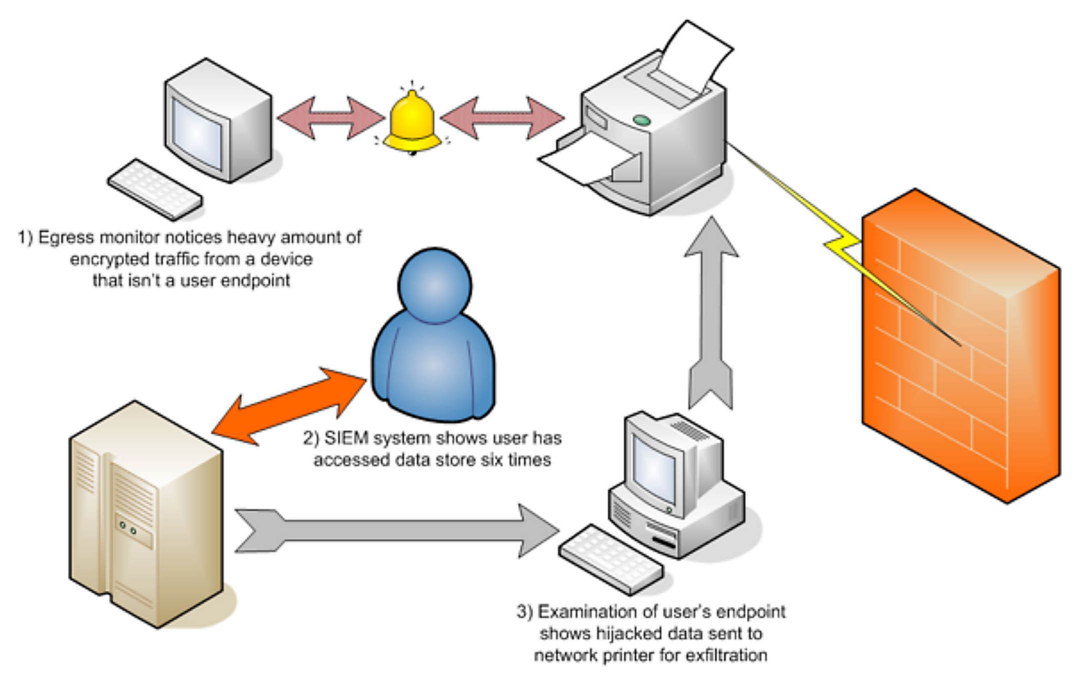
1) Egress monitor notices heavy amount of encrypted traffic from a device that isn't a user endpoint
2) SIEM system shows user has accessed data store six times
3) Examination of user's endpoint shows hijacked data sent to network printer for exfiltration

Figure 2. Tracking and Tracing an Attack to Block It and Repair Damage

With an accurate view of the situation, remediation teams can specifically target the affected systems and network segments. They should correlate incident data against vulnerability assessments to ensure that the exploited vulnerabilities are identified and remediated. Finally, administrators and managers should ensure that the data collected in the response—from initial probe through remediation—improves their blocking of and response to similar attempts.

Drafting a Breach Plan

Every response plan needs to include set actions, from detection to validation, to containment, remediation and future prevention. These are critical components of the CSCs, COBIT and other frameworks, as well as the SANS Institute's PICERL model.[15]

Using these and other guidelines, here are some ways to establish your response plan:

1. **Have a Jump Bag.** The first step of response is preparing a "jump bag" that contains the plan, contact lists, tools and procedures that are needed for recovery. If you do not have this, you are not ready to respond. (For what goes into a useful jump bag, see the 2011 Squidoo post on the subject by Todd Edmands.[16]

2. **Show Due Diligence.** The second point is that it is imperative to be able to show you made every effort to achieve due diligence. The key to this is adoption of a data security framework; among the best known are the CSCs, COBIT 5, and the ISO 27000 series, and the National Institute of Standards and Technology's preliminary cybersecurity framework shows promise.[17] Being aligned with a credible framework shows that an IT organization is pursuing due diligence.

3. **Identify the Problem(s).** Proper detection, as stated earlier, is key. This involves well-tuned IDS, firewall and perimeter devices, as well as application security, endpoint security, correlation using a SIEM system and, finally, strong reporting and alerting. To reduce the noise of false positives, consider whitelisting, blacklisting, using third-party intelligence for indicators of compromise, and sandboxing with honeypots to capture maliciously behaving code.

4. **Validate Your Findings.** It is critical to establish "ground truth" as quickly as possible. For example, if an IDS or similar system reports something suspicious, one must take the time to validate it, starting with a "sanity check." An automated system could report a Windows machine running Internet Explorer is showing signs of compromise, but when the incident responder investigates and finds it is a Mac running Safari, the odds are the IDS needs tuning to eliminate the false positive. One should also beware of false positives from tools that rely on comparing strings of data against signature databases; if you are looking for the string **`Xyzzy`**, that pattern will eventually occur—there is so much traffic and only so many possible characters, that it has to turn up sometime. However, it might not be the **`Xyzzy`** that indicates hostile activity.

[15] From the steps of incident response: Preparation, Identification, Containment, Eradication, Recovery and Lessons learned; featured in the SANS course "SEC 401: Security Essentials Bootcamp Style"; www.sans.org/course/security-essentials-bootcamp-style

[16] "Managing a Computer Security Jump Bag," Squidoo.com; www.squidoo.com/jumpbag

[17] COBIT: www.isaca.org/Knowledge-Center/COBIT/Pages/Overview.aspx;
ISO 27000: www.27000.org;
NIST: www.nist.gov/itl/upload/preliminary-cybersecurity-framework.pdf

5. **Contain the Damage.** This is where the best incident response folks shine, because they maintain their calm even when the storm is at its height. Two of the most important keys to containment are limiting the degree of exfiltration of intellectual property and limiting the number of infected hosts. Security gateways can be key tools to limit exfiltration. Simply normalizing the average amount of traffic a given internal host sends, and at what times of the day, can be a very effective step. Additionally, focusing on encrypted traffic can be a signal your organization's data is leaving the building, and should be contained by stopping it at the egress point. One simple thing organizations can do to limit infection is to make your response systems aware of the fact that desktop computers do not talk to other desktop computers; they talk to servers. A simple access control statement in the switches that does not allow a user VLAN to access another user VLAN may slow or even prevent the spread of infection. Couple that with a script that screams for help if the log file of access-denied events starts to grow exponentially, and you have a prevention and detection system that still advances the kill chain by stopping an attack's spread at the source.

6. **Eradicate the Threat.** The antimalware companies have made great strides into "deep cleaning," but it is very hard to know you have removed all of the malware, especially malware located in thinglets, mobile phones and other hard-to-monitor systems and endpoints. Log files can help determine what the attackers changed so the response team knows what needs to be repaired. Think about keeping logs for at least a year. It sounds radical, but organizations may want to think about protocols that let the organization continue to operate even if infected with malware.

7. **Close the Vulnerability.** System and network vulnerabilities can be relatively easy to fix by applying a patch, tweaking a configuration, retiring a program or adding new firewall or IDS rules. Human vulnerabilities, of course, are much harder to patch.

8. **Coach and Follow Up.** Consider putting a qualified "cyberbreach coach" on retainer. According to Michael Hoehl, a global IT security officer interviewed for this paper, senior executives are "keenly aware that loss of consumer data can have serious career continuation consequences."

9. **Retain Legal and PR Support.** In addition to a coach on operational response, organizations may want to consider keeping qualified and specialized legal counsel and public relations experts on retainer for data breaches because the potential for additional legal action is staggering. A quick visit to the Privacy Rights Clearinghouse website[18] will convince you that every industry sector has suffered data breaches, and one ought to assume that third parties—whether those are credit card thieves or government intelligence agencies—have access to every online action. According to the Ponemon Institute, "critical to controlling costs [related to a data breach] is keeping customers from leaving."[19] That is where a skilled public relations effort comes into play. In general, acknowledge an incident has occurred, but be careful to double- or triple-check any facts that you release and explain what you're doing to protect your stakeholders. Do so in a timely fashion in accordance with your state, local and industry regulations.

10. **Re-Examine the Plan.** After-action reports and similar postmortem analyses are critical parts of the process, because during the initial response you can be moving fast enough that perhaps not everything is done by the book or as planned. During the analysis, you should revisit all stages of your response and remediation looking for ways to improve. Even though response requires moving quickly, IT organizations should train responders to make notes on any errors or process problems they encounter. From a response point of view, these tend to fall into two basic categories:

    - **Organization problems.** These issues either allowed the incident to happen or made remediation difficult.
    - **Response issues.** These can range from lack of a tool or procedure to simply an incorrect analysis. In the lessons learned phase, you can examine these notes carefully with an eye to improving the process.

[18] www.privacyrights.org

[19] "Ponemon Institute Releases 2014 Cost of Data Breach: Global Analysis," Ponemon Institute press release, May 5, 2014; www.ponemon.org/blog/ponemon-institute-releases-2014-cost-of-data-breach-global-analysis

Conclusion

Attackers need to find only one weakness to get into an enterprise and begin to spread. Defenders need to plan for the inevitable breach and have a plan in place. Before an incident happens, organizations should have answered some important questions, including:

- What is our response plan?
- What personnel do we have and whom do we call?
- Do we involve law enforcement?
- Do we engage a breach coach or other outside organization?
- What is our response to blackmail?
- What is our response to a sensitive data breach?
- What is our remediation plan?

These and many other questions must have answers before a breach occurs. Then, once the investigation ensues and the breach is remediated, these questions should be asked again, retrospectively, to tune that plan for future problems.

Ultimately, the response plan should involve processes for removing the vulnerabilities involved in the event, and creating intelligent responses to future similar events. This creates a full-circle cycle to detect, prevent and respond as needed in today's complex threat landscape.

About the Author

Stephen Northcutt founded the GIAC certification and is the former president of the SANS Technology Institute, a postgraduate college focusing on IT security. He is the author or co-author of *Incident Handling Step-by-Step*, *Intrusion Signatures and Analysis*, *Inside Network Perimeter Security* (2nd edition), *IT Ethics Handbook*, *SANS Security Essentials*, *SANS Security Leadership Essentials* and *Network Intrusion Detection* (3rd edition). He was the original author of the Shadow intrusion detection system before accepting the position of chief for information warfare at the Ballistic Missile Defense Organization. Stephen is a graduate of Mary Washington College. Before entering the field of computer security, he worked as a Navy helicopter search and rescue crew member, whitewater raft guide, chef, martial arts instructor, cartographer and network designer.

Sponsor

SANS would like to thank this paper's sponsor:

Symantec™