A Trend Micro Research Paper

Suggestions to Help Companies with the Fight Against Targeted Attacks

Jim Gogolinski

Forward-Looking Threat Research Team

Contents

Introduction

This research paper provides some thoughts on how to configure a network in order to make lateral movement harder to accomplish and easier to detect, as well as how to prepare to deal with an infection. Given the advances attackers have been making, it is very unlikely that organizations will be able to keep motivated and patient adversaries out of their networks. In most cases, the best one can hope for is to detect targeted attacks early and limit the amount of information the attackers can obtain access to.

First and foremost, dealing with a targeted attack is a costly effort. To properly deal with an infection, a company should expect months of concentrated effort by a dedicated response team. In order for this team to be effective, the corporate network should have been configured in a way that provides all of the forensic data the team would need.

Before getting into specifics, this paper defines what a targeted attack is and provides a high-level timeline of a typical attack sequence. It then lists some possibilities as to why an organization may be a target for reasons other than its intellectual property.

This research paper also offers some suggestions on network tools and configuration settings to make it harder for an adversary to move laterally throughout a target network, and for an organization to protect its critical resources and help its response team with investigation and eradication efforts. It also states why ensuring adequate logging, network segmentation, and user account controls should be put in place.

In addition, this paper discusses the various levels of data confidentiality and why organizations need to protect sensitive data. It provides some insights into building a successful response team and why gathering threat intelligence is important in battling targeted attacks.

Finally, this paper provides reasons why running penetration tests against an organization's network is beneficial. Utilizing the right tools and implementing custom defense strategies can also aid organizations in the fight against targeted attacks.

Targeted Attacks

Defining a Targeted Attack

Several different types of attack can be perpetrated against computer users. One of the ways to classify attacks is based on the attackers' intent. Attackers may target a user for financial gain. In such an attack, the attacker is either looking to gather credentials that can be used to access online banking accounts or for some other way to extort money from a victim (e.g., installing a piece of ransomware in a user's computer).[1] In addition, corporate systems can be targeted for financial gain. Stolen databases of contact information can, for instance, be sold to spammers or other cybercriminals. Computers can also be targeted to become part of a botnet, allowing cybercriminals to remotely control a large group of infected computers.

A targeted attack can be considered a long-term cyber-espionage campaign against an organization. The attackers' goal is to gain persistent access to the target network. This may allow them to extract confidential company data and possibly place logic bombs that could damage a target network and infrastructure. The effect could be worse if the target organization is involved in handling critical infrastructure. Persistent access may also allow them to use a target's infrastructure as a false front to launch attacks against other organizations, giving new attempts an air of legitimacy (i.e., coming from a trusted partner). Furthermore, breached networks can be used as a stepping stone during attacks against other organizations, making it harder to trace the attack back to its originator.

Reasons for Attacking an Organization and Why Employees Should Care

When many people think of cyber espionage, they visualize government agencies and defense contractors as primary victims. While these are key targets, many other industries such as pharmaceutical, petrochemical, energy, manufacturing, and mining are also targeted. Some targets are even less obvious—legal firms, think tanks, and human rights organizations. Just because you think your organization is not an obvious target does not mean that it cannot be one.

1 Trend Micro Incorporated. (2013). *Threat Encyclopedia*. "Ransomware." Last accessed August 29, 2013, http://about-threats.trendmicro.com/us/definition/ransomware/index.html.

Industrial espionage can be detrimental to a company. One of the things one needs to keep in mind when considering a targeted attack is that corporate data can be just as valuable as military data. Stealing years' worth of research data is, for instance, much more economical than doing the research on one's own and can allow an organization that is just starting out in a new area to cut prices down due to nonexistent research and development (R&D) costs. Furthermore, disclosing that one's organization has been the victim of an attack can lead to decreased investor confidence, falling stock prices, and investor lawsuits.

An organization may also become a target not only for its own products or the information it holds but also because it is somehow connected to an ultimate target. Penetrating a network through a weak link then hopping around systems within that network or "island hopping" has been around for years.[2] The same is true of penetrating organizations. A less secure organization that is somehow connected to the ultimate target could be penetrated and used to gain access to the more strongly secured target. Value-added resellers (VARs) often have access to the networks of the companies they supply products and/or services to. An organization may only make simple widgets but sell these to a large integrator that uses them in its products. It may also have employees who belong to a group the target company's employees also belong to, be they professional organizations or even just special collaborative groups. Using the infrastructure, mostly through compromised email accounts and its connections to the "real" target, gaining access to the ultimate target organization's employees becomes easier for the attacker. If an employee, for instance, vaguely remembers working on a project at some point in the past with the sender of an email, he/she is more likely to trust it, especially if the context is related to that piece of work.

Typical Timeline and Real-World Examples

Once an organization has been targeted by an intrusion attempt, it can expect continued attempts until, at the very least, one succeeds. However, one successful network breach may not make attacks cease. In fact, attackers often continue with their attempts after a successful breach due to several reasons. First, the breach could be detected and that avenue of connectivity closed. Having a second entry point allows redundancy because as the saying goes, "two is one and one is none." Another reason is that different teams or subgroups from within the same attacking organization may be trying to gain access to a single network over time. These groups may be looking for different information and may not even be aware of the fact that other groups have already targeted the same entity.

2 Jesper M. Johansson. (January 2008). *TechNet Magazine*. "Island Hopping: The Infectious Allure of Vendor Swag." Last accessed September 11, 2013, http://technet.microsoft.com/en-us/magazine/2008.01.securitywatch.aspx.

A typical attack often begins with a spear-phishing attempt.[3] Spear-phishing emails are usually sent to a small number of targeted employees and crafted to make recipients believe they are legitimate. More often than not, they are made to look like they come from someone a user would expect to receive email from, possibly a boss or colleague. These emails likely contain a link to a malicious website or some sort of weaponized attachment. This attachment may exploit a vulnerability in Microsoft™ Word® or Excel® or an Adobe® product. It may also be a .ZIP file that when opened leads the recipient to believe it is some sort of document when it is actually an executable file. When a user unsuspectingly tries to open the file, it exploits his/her system. Often, to complete the ruse, doing so opens an innocuous document so the user has no idea what really happened. Note that threat actors can also use the same techniques to attack a user's personal (i.e., nonwork) email addresses because many employees use their work computers to read personal as well as corporate emails.

Even though spear-phishing emails are currently the predominant attack vector, an adversary can attempt to gain access to a target network via several other ways. All publicly facing computers can be probed for vulnerabilities. Threat actors can also attempt to socially engineer their way into a target organization's facilities and install malicious software, scan for open wireless access points (WAPs), or insert their own WAP into the network. They can also drop a malicious USB drive (i.e., thumb drive) in a target's parking lot in hopes that a curious employee would pick it up and plug it into his/her computer to see who it belongs to or what is in it.

Regardless of delivery mechanism, once a piece of malicious software runs, it usually attempts to communicate back to a command-and-control (C&C) server to receive further instructions. The software may immediately initiate communication or lie dormant on a system for hours in an attempt to remain hidden. One of two things usually happens when the software accesses the C&C server. First, the software may automatically download and install additional malware. This type of malware is called a "downloader." Second, the software can communicate back to the C&C server. A human monitoring the C&C server would then notice the new connection and initiate some sort of action. This type of software, called a "remote access Trojan (RAT)," gives a threat actor the ability to examine a system, extract files, download new files to run on a compromised system, turn on a system's video camera and microphone, take screen captures, capture keystrokes, and run a command shell.

3 TrendLabs℠ APT Research Team. (2012). "Spear-Phishing Email: Most Favored APT Attack Bait." Last accessed September 2, 2013, http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/white-papers/wp-spear-phishing-email-most-favored-apt-attack-bait.pdf.

At this point, the adversary will attempt to move laterally throughout a compromised network for two main reasons. First, he/she would want to gain additional persistent access points to ensure continued access. Second, he/she would want to locate whatever data could be valuable. To move throughout a network, threat actors would attempt to steal whatever user credentials they can find. These credentials may give them access to additional machines and data spread throughout the network. As they collect data, they will exfiltrate it out of the network to another location where they can collect and bring it back to their environment for further examination.

The adversary's initial reaction to contact by the attack software will typically occur after a very short period of time. After that, most adversaries may move at a very slow pace to try to remain undetected. If they think they have been detected, they may go dormant for weeks or even months before resuming activity. If an organization completely eradicates their presence from their network, the threat actors will start the attack cycle all over again.

Several well-known advanced persistent threat (APT) attack examples exist such as the Google Aurora, ShadyRAT, DUQU, and FLAME attacks, as well as a fairly recent attack targeting Saudi Arabia.[4]

Infrastructure

What an Organization Can and Should Do

Without a properly configured infrastructure, organizations stand no chance against targeted attacks. Properly configuring a network may not provide any initial financial gain and would, in fact, more likely increase an organization's overall IT cost. In today's world, however, organizations cannot afford to have an improperly configured infrastructure. The actual and future costs of a breach and theft of proprietary information will far outweigh any amount incurred to properly configure infrastructure.

4 Valerie Boquiron. (January 19, 2010). *TrendLabs Security Intelligence Blog*. "Cyber Attacks on Google and Others—Who Is Really at Risk?" Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/cyber-attacks-on-google-and-others-who-is-really-at-risk/; Nart Villeneuve. (January 26, 2012). *TrendLabs Security Intelligence Blog*. "Top APT Research of 2011 (That You Probably Haven't Heard About)." Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/top-apt-research-of-2011-that-you-probably-havent-heard-about/; Karl Dominguez. (November 2, 2011). *TrendLabs Security Intelligence Blog*. "Zero-Day Exploit Used for DUQU." Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/zero-day-exploit-used-for-duqu/; Trend Micro Incorporated. (May 31, 2012). *TrendLabs Security Intelligence Blog*. "Update on FLAME." Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/update-on-flame/; Kelly Jackson Higgins. (August 22, 2012). *Dark Reading*. "Shamoon, Saudi Aramco, and Targeted Destruction." Last accessed September 11, 2013, http://www.darkreading.com/attacks-breaches/shamoon-saudi-aramco-and-targeted-destru/240006049.

Segmentation

Corporate networks should be broken down into as many segments as makes sense. Simply put, a network segment is a collection of workstations, servers, printers, and other devices that can access one another. In a home network or simplistic corporate environment, all of the components of an infrastructure may be situated in a single segment. As a network becomes more complex, it should be broken down into logical segments. Network segments are typically separated by firewalls that contain and control what traffic is permitted in and out of each segment. Network segments can be considered a series of secure rooms inside a large open building. The absence of segments can allow anyone who can enter the building access to everything inside it. Adding secure rooms makes the task of getting full access to the complete network much harder for anyone who can breach the perimeter. Getting in through the front door does not allow immediate access to everything in the building. Failure to properly segment a network can allow an attacker unrestrained access to the entire corporate infrastructure. In addition, network segmentation makes it harder for employees to access data they are not authorized to.

Networks should be broken down into as many logical segments as possible. They may, for instance, be segmented by function such as finance, marketing, engineering, sales, support, and production. They can also be segmented by geographic location or security level (e.g., unclassified, classified, secret, top secret, or contains personally identifiable information [PII]). Each segment should be made as secure as possible with a firewall preventing unwanted access by unauthorized systems. Unless a strong business case exists, a machine on the marketing network should not have direct access to the engineering network. Likewise, a machine from the engineering network should not have access to the finance network.

Logging

Logging and log analysis are key methods that can help detect a breach and allow a response team to understand where the attackers went and what they were looking for. In a large corporate environment, logging data provides insights into the health and activity of the network. Sharp analysts who understand the daily ebb and flow of network traffic may be able to detect a targeted attack early enough in the intrusion timeline to thwart it before it has a chance to take root and succeed. If the attack takes hold and is not detected early, the logged data is often all that the forensic response team has to work with to help recreate the attacker's activities. It goes without saying that logs need to be actively monitored and tools like security information and event management (SIEM) systems can help with this task.

What types of information should be logged? If an organization uses network address translation (NAT), records of internal mapping to physical machines should be kept. It will do an organization no good to know that 192.168.1.27 accessed a known C&C server if any machine in a large subnet could have used that NATted address. Web proxy and Domain Name System (DNS) requests should also be logged. User access (i.e., both log-ins and log-outs)—physical, Remote Desktop Protocol (RDP), virtual private network (VPN), and others—should also be noted. Both external and internal network traffic should be captured whenever sensible and legal to do so. If full packet captures (pcaps) cannot be obtained, at the very least, network flows should be recorded. Emails with full headers and attachments should also be archived for an extended period of time. Firewall drop rules should be logged as well, as these contain a lot of information on which systems may be attempting lateral attacks over nonstandard ports.

Logging should be done at a very verbose level, as it is better to have too much information than miss key data. Ideally, all of the systems in an organization should synchronize their clocks with a central time server so that all of the log file timestamps are identical across encompassed geographic areas.

Log data should be centrally located within an organization as well. Organizations with networks in different geographic locations need to determine what level of centralization makes the most sense although they should have a way to obtain all of the logs from a central location (within a short period of time) if required.

Organizations should keep log data for as long as possible. At the very least, a rolling year's worth of logs should be made immediately accessible. Since the majority of logs are text based, they can be easily compressed for longer-term storage. These logs should also be backed up and stored offline in case an attacker manages to access a centralized logging system or its associated repository.

One of the key elements to keep in mind when devising a logging strategy is to store data in a format that is easily searchable and allows for correlation. A comma-separated values (.CSV) file allows for easy searching and parsing. Many devices will, by default, log to a proprietary binary format. If possible, try to configure everything to log to pure text and if that is not possible, implement a process to export the data into a textual format.

Although properly configuring a logging infrastructure costs slightly more, the cost will be paid back many times over in the course of a single targeted attack investigation. Imagine how much quicker an investigation can proceed if searching through DNS logs to find machines that attempted to access a known malicious C&C server is possible. Internal IP addresses can then be correlated with individual machines on the network. Once those machines have been identified, searching through network flow logs for other machines these accessed can be done. Having all data available to a response team allows much quicker and shorter investigations and lower overall remediation costs.

Some would argue that logging will require too much storage space. Given the decreasing cost of storage and the high cost of resources to perform analyses, organizations cannot afford to discard data. At least three months of uncompressed data can be kept while the rest can be compressed and archived. When considering an archival strategy, keep data in reasonably short time durations for easier and quicker decompression and searching.

User Accounts and Workstations

As a general rule, users always want access to everything. They want to be able to go everywhere and want access to all data available; even if they do not need it now, they just might in the future. And as most people have computers at home, they are used to being able to administer their own machines and install whatever software they want. Unfortunately, what users do on their home computers is not applicable to the corporate environment. Forcing users to request access to new data sources may lead to more work for IT staff but, in the long run, makes an environment much more secure.

As previously mentioned, adversaries try to get as many sets of credentials as possible. An organization should make that as hard to do as possible while ensuring that any compromised account has as little access and as fewer privileges as possible. These will help minimize the damage caused. It is important to note that if the attackers are able to obtain valid password hash credentials, there is no defense against pass-the-hash techniques.[5] The user account and workstation policies outlined in this paper will not make IT staff popular with the employees and may result in more work for IT staff who will have to deal with inevitable problems that will arise. Again, organizations have to weigh these against the cost of data loss and remediation efforts.

5 Bashar Ewaida. (January 21, 2010). "Pass-the-Hash Attacks: Tools and Mitigation." Last accessed September 11, 2013, http://www.sans.org/reading-room/whitepapers/testing/pass-the-hash-attacks-tools-mitigation-33283?show=pass-the-hash-attacks-tools-mitigation-33283&cat=testing.

User accounts should work under the least-privilege model. Secure passwords should be required and, if possible, the use of two-factor authentication, implemented. Passwords should be frequently changed and users should not reuse passwords across accounts (e.g., domain and VPN log-in). If possible, check passwords against rainbow lists and run those not on the lists through password-cracking software. Users whose passwords were found or cracked should be forced to immediately change them. Local administrator accounts should be disabled or removed and domain administrator credentials should never be allowed to remain cached. If administrator access is required, the administrator should remotely access the machine. After completing the work, the machine should be rebooted. Password policies should require account lockout after a number of failed tries with either a very long reset timeout or, if possible, that the locked-out user calls the help desk to reset his/her account.

User workstations should be locked down if at all possible. Each computer should be kept fully patched and have full logging enabled. If the environment permits, an organization should consider implementing a white-listing solution on user workstations and laptops. The Trend Micro™ OfficeScan™ Intrusion Defense Firewall plug-in applies application control filters to alert system administrators to or block specific traffic such as instant messaging and media streaming, removing bad data from business-critical traffic.[6] Organizations should also consider running integrity monitoring tools that detect changes to file systems and registries.

Another consideration for user workstations and laptops is to have some sort of monitoring agent that can remotely gather statistics and information from each machine and return that data to a centralized server. Such software would allow security staff to quickly examine any system they believe may have been compromised, which can often help stop an attack while it is at its infancy. It also increases the efficiency of a forensic team in the event of an investigation. The team will be able to scan systems for malicious software indicators such as hashes, mutexes, strings, and registry file changes.

6 Trend Micro Incorporated. (2013). "Trend Micro™ Intrusion Defense Firewall." Last accessed September 13, 2013, http://www.trendmicro.com/cloud-content/us/pdfs/business/datasheets/ds02_osce.pdf.

Data Protection

Keeping the Value of Data

Remember that it is very likely that one of the attackers' main motivations in moving through a network is to look for data to steal. Keeping all data in a centralized, weakly protected location is just as good as boxing it all up and shipping it to attackers by overnight delivery. Attackers are patient enough to keep looking until they find what they want. The longer the attackers search a network, the more time an organization will have to detect attack activities, remove the threat, and patch all of the pathways the threat actors used.

Data Segmentation

Not all data is created equal. While it is true that employees generate a vast amount of work products every day, most of them are not critical to an organization's success and continued well-being. Each business unit needs to take a close look at its data to determine which would be detrimental if it were to fall into the wrong hands. This data should be treated differently than normal day-to-day documentation. Organizations should consider such kind of data "crown jewels" and protect it accordingly. This data, for instance, should not be made available for downloading onto a workstation and must only be accessed on the file server by users with special access privileges. It should be heavily encrypted as well.

Routine and possibly misleading data may be left on individual workstations. It can be considered sacrificial data. The amount of time it may take an adversary to realize the data is not very useful may be the amount of time required to detect and start eradicating a network threat.

Any sensitive document sent via email should be encrypted separately from what the corporate email system offers using tools such as GPG. If an adversary should gain access to an employee's desktop and have the ability to open his/her email client, the threat actor still should not be able to decipher any information he/she sees.

Protect individual pieces of important information as much as possible. In a pharmaceutical manufacturing environment, for instance, you need to have certain pieces of information spread throughout the factory environment. Each piece on its own may not give away the recipe. However, if a patient threat actor had the ability to examine an entire plant's environment, he/she would eventually be able to piece various bits of information together to gain complete formulas and process information.

Data Protection Infrastructure

As previously mentioned, different kinds of data require different levels of protection. Data that requires the highest level of protection can be kept on a disconnected network that requires physical access rights to view. The next tier may be hosted on a secure server that requires a different form of two-factor authentication than ordinary local network access. It should not accept cached credentials nor should it allow data to be removed from a server. Note that the access logs on such a server should be closely monitored as well. Third-tier data can be stored on a regular file server. Note, however, that this kind of data should be something that, if stolen, would not cause great harm to an organization.

Organizations should also consider watermarking critical documents and adding data loss prevention (DLP) protection to their infrastructure. This will also allow them to detect unauthorized document movement.

Response Team

Reasons for Building a Response Team

If an organization's IT security team has reason to believe that it has been breached, time is of the essence. This is not the time to attempt to build a team whose members have all of the skills required to respond to a breach. A response team requires a very specific set of skills that are very difficult to learn while in the middle of an investigation. The organization also runs the risk of having staff missing a required skill. If the adversary has been inside a network for some time, expect a long, expensive remediation effort. A proper team can, however, help shorten this process.

How Incident Response Efforts May Look

First and foremost, responding to an intrusion is a costly and stressful event. People will look for answers to questions that begin with "who," "what," and "why" long before a forensic team is prepared to offer any factual conclusion. If news of a breach reaches external sources, the pressure from media as well as investors will be even greater. Unless properly handled, the initial days of the investigation will be chaotic and filled with fear, misleading information, and incorrect conclusions. One of the first things an organization has to consider is if it has staff qualified to properly handle the forensic investigation. If not, it has to decide who to bring in to handle the investigation. It will also need to decide if it needs to involve local, state, or federal law enforcement agencies or any regulatory agency for that matter.

Once a team has been put in place, its members should be given some time to conduct interviews with people involved to start their investigation. How well an investigation goes often depends on how well an organization's infrastructure was configured. If all of the steps mentioned in the previous sections have been followed, the team will stand a good chance of conducting a reasonable forensic investigation and most likely be able to piece the threat actors' activity together as well as have an idea as to what was exfiltrated. It is important to understand though that even with the proper infrastructure in place, an investigation may still take longer than expected or hoped for. But if the infrastructure is not even in place, the team will flounder and likely not have much success.

Team Overview

Organizations should consider creating a cross-functional incident response team that is separate from their normal network or IT operations team. This team should be created as soon as possible and not when they have to respond to an incident. Roles, responsibilities, and authorities should be documented and clearly understood by everyone on the team.

Team Composition

An incident response team should have members from the following functional areas:

- **Technical:** People from any existing security group (e.g., computer and physical), network operations, or workstation management who will handle the technical side of the investigation.
- **Threat intelligence:** Employees focusing on the threat landscape who can answer questions such as who might be attacking, why they are attacking, what techniques they are using, and others.
- **Human resources:** People who can answer any question related to the organization's employees and policies.
- **Legal:** Employees who can answer any question about the legality of an investigation or reporting an incident.
- **Public relations:** People who can handle external communications regarding the incident.
- **Executive management:** People who can provide insights from a corporate perspective.

Subject matter experts can be brought into the team as necessary. The team should have a technical leader who is the interface to the broader "nontechnical" team. This leader should have the authority to make decisions regarding all technical details and staffing issues. In addition, an overall incident leader tasked to handle all coordination efforts with internal and external entities should be chosen. The overall team leader should have the decision-making authority about anything related to the incident. Leaders should have the ability to make decisions for several reasons. Trying to explain technical details to a nontechnical manager takes additional time and all of the nuances cannot always be conveyed. The higher up the decisions need to go, the harder it becomes to schedule time to meet. As has been said, time is of the essence. Furthermore, giving leaders decision-making power shows confidence in the team and its members' abilities.

Team Operation

When responding to an incident, the team should meet regularly to track their progress and issue updates to senior management. All interactions should go through the incident leaders and not end-run to individual team members. Going directly to individual team members for updates and questions is very disruptive to the flow of the efforts of the technical team and should be avoided at all costs. During the initial stages of an incident, the team should have frequent but very short and efficient progress-tracking meetings. As the investigation continues, the team can meet less often but should still do so on a regular basis.

Team Training

The technical team should be formed ahead of time and afforded an opportunity to train and work together outside of their normal responsibilities (i.e., if they are not dedicated to the incident response process). Since incident response causes tremendous pressure on staff to get answers and results, it makes sense for them to already understand the process and their individual roles and responsibilities beforehand. Training will make the first few weeks of the response process much more effective. As the team trains and continues to learn, the things that they discover should be fed back to the IT groups to help improve an organization's overall security posture. One way of looking at this would be having a well-versed and trained city fire department versus a small-town volunteer force who never trains. If a facility, for instance, catches fire, which group do you think people would rather have respond?

Threat Intelligence

Defining Threat Intelligence

Simply put, threat intelligence refers to information about potential adversaries and their behavioral patterns. Threat intelligence is created when many pieces of raw data are analyzed to give a more complete image of the big picture. It can help thwart initial intrusion attempts. It can also help determine not only where an attacker has already been in an organization's network but also where he/she is likely to go and how he/she will get there.

Why Current Threat Intelligence Is Necessary

Raw data without intelligence is of limited value. To detect an adversary in a network, one of two things needs to happen. First, an analyst may detect an unusual traffic pattern or a series of events within his/her organization's internal network. This may involve a user's workstation accessing another's, a series of failed log-in attempts at an unusual time of day, a triggered firewall deny rule, or even a user call to report something unusual happening to his/her workstation. Another case may involve an external communication event. This could involve a workstation attempting to access a known C&C server or sending a .RAR file to some unknown remote location.

In both cases, an analyst needs to know what to look for. This is where threat intelligence comes into the picture. There are far too many Internet addresses out there for anyone to know which are legitimate and which are malicious. It gets even worse when one tries to determine the validity of every email that employees receive. Once an attacker infiltrates a network, understanding his/her tactics, techniques, and procedures (TTPs) can spell the difference between quick successful detection and years of undetected data exfiltration.

Attackers' skills continue to advance, new undetected tools created, and different ways to achieve the same goal found. Much of the knowledge on how attackers operate is kept behind closed doors. Once that information becomes public knowledge, attackers would know that they need to find different ways of doing the same thing. Some attackers even modify their TTPs based on their target. As such, what they do for company X may not be what they would for company Y. It becomes a game of cat and mouse. And the more knowledge an organization obtains, the better its chances of detecting and successfully remediating attacks.

External Threat Intelligence Sources

An organization can obtain external threat intelligence in two ways—partnering with a threat intelligence provider and utilizing automated software.

Threat intelligence providers have skilled employees who understand threat actors and TTPs, and can put the puzzle pieces they are presented with together. They typically provide their clients two deliverables—reports and feeds. Reports typically focus on a single subject. They can be short as is the case when something new is breaking or more detailed after there has been time to perform a complete analysis. Some possible subject areas are new campaigns (e.g., attacks targeting several companies or geographic locations at the same time), discussions of a newly discovered TTP or piece of malware, or descriptions of a threat actor. Feeds, meanwhile, are sources of data that can typically be included in automated network defenses. They may list items such as malicious URLs, email senders or subject lines, and hash representations of malicious documents or malware.

Enterprise-quality products from security vendors that are kept updated with the latest threat indicators can also help protect networks. Products such as Trend Micro™ Deep Discovery, for instance, provide the advanced threat detection and real-time intelligence an organization needs to discover and respond to targeted network attacks and APTs. Deep Discovery's specialized network detection uniquely discovers and identifies evasive threats and provides in-depth analysis and actionable intelligence about the nature of an attack. Integrating Deep Discovery with an organization's security infrastructure creates a "Custom Defense" solution, a complete network security strategy to detect, analyze, adapt, and respond to attackers.[7]

Internal Threat Intelligence Group

Whether an organization contracted a vendor to provide threat intelligence services or not, if it has the opportunity, it should still set up its own internal threat intelligence group (ITIG). Having members of an organization's own forensic team's full-time responsibility be gathering threat intelligence can provide tremendous value. An organization's ITIG will be responsible for two areas. First, they should monitor the Web for any reference to their company. Second, they should research any group or actor they believe may be a threat.

An ITIG should monitor the Internet to look for any reference to the organization. Its members should check blogs, Pastebin, and underground forums for anything that may be related to not just the organization but also the industry or group it belongs to. They should also look at sites where employees are mentioned such as conference sites that list attendees, as this can be used as a list of spear-phishing attack targets.

7 Trend Micro Incorporated. (2013). "Deep Discovery Advanced Network Security." Last accessed September 5, 2013, http://www.trendmicro.com/us/enterprise/security-risk-management/deep-discovery/index.html.

An ITIG should also spend as much time as possible learning about exploits, malware, and TTPs that can be used against the organization and its networks. Plenty of resources (e.g., books and sites) talk about exploits and malware. Note, however, that very little information about the TTPs adversaries use once inside a target network is available. Organizations need to evaluate their risk posture but one of the greatest places to learn what threat actors are up to is to have a secure area separate from the remainder of the corporate infrastructure. They should consider this a playground that can enhance their threat knowledge.

If an organization, for instance, received a spear-phishing email but caught it before anyone opened it, the IT staff can open the email and trigger the exploit, allowing the adversary to create a backdoor to the "playground." A properly configured playground would allow the organization to monitor all of the adversary's actions, capture all of the tools that he/she uses, and observe the techniques he/she uses to laterally move throughout the network. If the playground was large and realistic enough, the organization can gain a tremendous amount of knowledge that it can then feed back to the network in order to enhance its security posture. Since it is monitoring all activities, the organization can "detect" the adversary at any point when it became uncomfortable and clean up the environment in a way that the threat actor would believe he/she was caught. Setting up this playground is not a trivial task though and, if incorrectly done, will cause the adversary to immediately leave or, worse, become aggressive and hostile.

Penetration Testing

If an organization is not part of an industry that is required to conduct regular penetration tests, it should seriously consider doing so. Penetration testing can serve several purposes in an infrastructure. First, it can help identify areas in the network that need to be improved and patched. Second, the testing can also help identify areas where your network monitoring has gaps. Finally, it gives security and monitoring teams an opportunity to gain some experience with a realistic scenario.

When considering rules and operating procedures for penetration tests, give the penetration team as much leeway as possible. An important thing to remember is that the penetration test is meant to be a learning experience and not to make the security staff look good. It is okay to tell the penetration-testing team that it can do no damage and certain sensitive systems are off limits to anything considered risky but allowing the penetration testers to leave evidence of their presence behind should be acceptable.

It is a corporate decision on whether most of the security staff is informed of the penetration test. Unless it is management's specific intention, the penetration testing should not turn into a red/blue team exercise. The organization's defensive team can monitor but should not thwart the activities of the penetration testers. When penetration testing, there are many ways to achieve the same goal. If the network security team is to shut down an attempt, this may prevent your organization from learning about a bigger weakness further down in your network. One of the things that is artificial about a penetration test is the timeline. The penetration-testing team only has a week or two to complete all of its activities, whereas a determined and patient attacker may take weeks to months to do the same. So, any effort to shut down attackers' current avenue may thwart them due to time constraints. Since they have limited time, they may not have the chance to pursue other avenues to achieve their attempted goal.

The penetration test can be used by an organization's forensic team as a training opportunity as well. Since the penetration-testing team will provide complete documentation on where they went and what they did, the organization will know all of the answers before the start of a forensic exercise. It could let the forensic team work with no knowledge and monitor their progress but, at any point, reveal any additional data that may help the team's members learn.

Conclusion

The unfortunate reality for many of today's corporations is that it is not a question of whether they will be the target of an attack but when. The ideas and suggestions featured in this research paper are not revolutionary. Unfortunately, most organizations do not implement them. Many of the suggestions will not result in happy end users or IT support staff but they are critical if an organization wants to enhance its security posture and make it possible to detect and remediate an intrusion. Remember that time is on the side of the attacker. All he/she needs to do is get one user or administrator to make a mistake to get in to a target network.

Once the attacker gets in to a network, the organization will race against time. The longer the attackers stay in a network, the more footholds they gains and information they can steal. Note, however, that an organization may not be the actual end target but is being used to gain access elsewhere. Although this scenario seems to pose less damage, it can still have very negative effects, ranging from loss of contracts and investor confidence to lawsuits.

All of the defenses an organization puts in place against targeted attacks can also be used to help detect and mitigate insider attacks. It is, after all, an unfortunate reality that insider information theft can be just as detrimental to an organization as an external attack.

Implementing the suggestions in this paper will not contribute to greater efficiency, streamlined operations, or better user experience. The upfront costs in implementing them should be considered "costs of doing business." In fact, these are very similar to the cost of buying insurance. Insurance provides no benefits to an organization until something bad happens. At that point, the benefits far outweigh the cost. In today's environment, it will be very costly to not heed these suggestions.

References

- Bashar Ewaida. (January 21, 2010). "Pass-the-Hash Attacks: Tools and Mitigation." Last accessed September 11, 2013, http://www.sans.org/reading-room/whitepapers/testing/pass-the-hash-attacks-tools-mitigation-33283?show=pass-the-hash-attacks-tools-mitigation-33283&cat=testing.

- Jesper M. Johansson. (January 2008). *TechNet Magazine*. "Island Hopping: The Infectious Allure of Vendor Swag." Last accessed September 11, 2013, http://technet.microsoft.com/en-us/magazine/2008.01.securitywatch.aspx.

- Karl Dominguez. (November 2, 2011). *TrendLabs Security Intelligence Blog*. "Zero-Day Exploit Used for DUQU." Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/zero-day-exploit-used-for-duqu/.

- Kelly Jackson Higgins. (August 22, 2012). *Dark Reading*. "Shamoon, Saudi Aramco, and Targeted Destruction." Last accessed September 11, 2013, http://www.darkreading.com/attacks-breaches/shamoon-saudi-aramco-and-targeted-destru/240006049.

- Nart Villeneuve. (January 26, 2012). *TrendLabs Security Intelligence Blog*. "Top APT Research of 2011 (That You Probably Haven't Heard About)." Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/top-apt-research-of-2011-that-you-probably-havent-heard-about/.

- TrendLabs APT Research Team. (2012). "Spear-Phishing Email: Most Favored APT Attack Bait." Last accessed September 2, 2013, http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/white-papers/wp-spear-phishing-email-most-favored-apt-attack-bait.pdf.

- Trend Micro Incorporated. (2013). "Deep Discovery Advanced Network Security." Last accessed September 5, 2013, http://www.trendmicro.com/us/enterprise/security-risk-management/deep-discovery/index.html.

- Trend Micro Incorporated. (2013). *Threat Encyclopedia*. "Ransomware." Last accessed August 29, 2013, http://about-threats.trendmicro.com/us/definition/ransomware/index.html.

- Trend Micro Incorporated. (2013). "Trend Micro Intrusion Defense Firewall." Last accessed September 13, 2013, http://www.trendmicro.com/cloud-content/us/pdfs/business/datasheets/ds02_osce.pdf.
- Trend Micro Incorporated. (May 31, 2012). *TrendLabs Security Intelligence Blog*. "Update on FLAME." Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/update-on-flame/.
- Valerie Boquiron. (January 19, 2010). *TrendLabs Security Intelligence Blog*. "Cyber Attacks on Google and Others—Who Is Really at Risk?" Last accessed September 2, 2013, http://blog.trendmicro.com/trendlabs-security-intelligence/cyber-attacks-on-google-and-others-who-is-really-at-risk/.

Trend Micro Incorporated, a global leader in security software, strives to make the world safe for exchanging digital information. Our innovative solutions for consumers, businesses and governments provide layered content security to protect information on mobile devices, endpoints, gateways, servers and the cloud. All of our solutions are powered by cloud-based global threat intelligence, the Trend Micro™ Smart Protection Network™, and are supported by over 1,200 threat experts around the globe. For more information, visit www.trendmicro.com.

Securing Your Journey to the Cloud

10101 N. De Anza Blvd.
Cupertino, CA 95014

U.S. toll free: 1 +800.228.5651
Phone: 1 +408.257.1500
Fax: 1 +408.257.2003