WHITE PAPER | OCTOBER 2014

Advanced Persistent Threats: Defending from the Inside Out

Russel Miller
CA Technologies, Security Management

Table of Contents

Executive Summary

Challenge

Securing an organization is an increasingly difficult challenge. Attacks are growing in complexity, and the rise of Advanced Persistent Threats (APTs), a type of targeted attack, has made organizations more aware of their vulnerability to attack. Companies from RSA Security to Google to Northrup Grumman have found themselves the target of APTs. Not having been the victim of a significant breach in the past does not guarantee safety moving forward, as organizations specifically targeted by APTs face challenges not normally faced by security administrators, such as spacing out actions over the course of months or years to avoid detection. The damage caused by a breach is also increasing, making this challenge all the more real to senior executives.

Opportunity

There is no "silver bullet" when it comes to defending against APTs. Multiple layers of protection must be employed that combine to reduce both the potential for a breach and mitigate the damage if one were to occur.

The initial approach to defending against targeted attacks was to secure the perimeter using firewalls and intrusion detection systems to detect and block anomalous behavior. This approach can be effective in defending against certain types of attacks, but cannot secure against all attack vectors, such as "spear phishing" and "social engineering."

While no point security product—technology-based or otherwise—can fully protect an organization from APTs, today's availability of cross-domain security solutions can help organizations protect themselves better than ever before. Privileged identity management, information protection and control, and internal infrastructure security are areas that have traditionally been looked at in silos, but may now be combined to enable organizations to secure their IT infrastructure and datacenters in complementary ways. CA Technologies calls this identity and data intelligence.

Benefits

By understanding and protecting against advanced persistent threats, organizations reduce their risk in the event they are specifically targeted for an attack. The risk that is reduced is not only financial, but reputational, operational, legal and regulatory as well.

By taking a holistic view of security that can be used against APTs, an organization is also protecting itself from less-advanced, automated, and even internal attacks. A comprehensive approach to security has many other advantages, including improving compliance, enabling cloud-based services, improving virtualization security and enabling cost savings.

Section 1: Challenge

Advanced persistent threats: not business-as-usual

Advanced persistent threats present challenges that are distinct from traditional security risks. Advanced Persistent Threat refers to a long-term and sophisticated attack on a specifically targeted entity. The attacker is often state-sponsored and seeks to gain high-value intelligence from other governments, but may also be performed by and target private organizations. The term was first used by the United States Air Force in 2006.[1] The National Institute of Standards and Technology (NIST), defines APTs as follows:[2]

"The advanced persistent threat is an adversary with sophisticated levels of expertise/significant resources, allowing it through the use of multiple attack vectors (e.g., cyber, physical, and deception), to generate opportunities to achieve its objectives which are typically to establish and extend footholds within the information technology infrastructure of organizations for purposes of continually exfiltrating information and/or to undermine or impede critical aspects of a mission, program, or organization, or place itself in a position to do so in the future. Moreover, the advanced persistent threat pursues its objectives repeatedly over an extended period of time, adapting to a defender's efforts to resist it, and with determination to maintain the level of interaction needed to execute its objectives."

While other definitions vary, all three words help clarify what an advanced persistent threat is:[3]

- **Advanced:** The attacker has significant technical capabilities to be able to exploit vulnerabilities in the target. This may include access to large vulnerability databases and exploits and coding skills, but also the ability to uncover and take advantage of previously-unknown vulnerabilities.
- **Persistent:** APTs often occur over an extended period of time. Unlike short-term attacks that take advantage of temporary opportunities, APTs may take place over the course of years. Multiple attack vectors can be used, from Internet-based attacks to social engineering. Minor security breaches may be combined over time to gain access to more significant data.
- **Threat:** In order for there to be a threat, there must be an attacker with both the motivation and ability to perform a successful attack.

Purely automated tools are not considered an APT by themselves, although they may be used by an organized and coordinated group as part of a larger attack.

Stages

A typical advanced persistent threat may consist of the following four stages:

Figure A.

Four stages of an advanced persistent threat

1. **Reconnaissance:** Investigation into an organization's vulnerabilities. This can include basic research, including domain queries, to port and vulnerability scans.
2. **Initial entry:** Exploitation of weaknesses to get a foot-hold into the target network. This can be done by using sophisticated technical methods or by techniques such as spear phishing (directed phishing attacks) which result in obtaining a regular user's access to a single system. "Social engineering," or the exploitation of people, is also often a common method of gaining access.
3. **Elevation of privileges and expansion of control:** Once an attacker penetrates the network perimeter, he or she attempts to gain additional privileges and control over critical systems. This step also may involve installation of "back door" tools to make future access to the network simpler.
4. **Continuous exploitation:** Once control is established, an attacker may continually export sensitive data.

The third and fourth stages may take place over the course of years, in order to reduce the risk of detection.

What makes APTs different?

The most critical difference between APTs and "normal" threats is that an organization is specifically targeted. While defending "the perimeter" and using standard security controls may protect an organization from standard attacks, these techniques may not be sufficient when facing APTs. Patient attackers can wait for new vulnerabilities to open up a weakness or can combine seemingly small vulnerabilities into a large-scale and damaging attack.

When facing such a threat, the normal rules do not apply. In the past, many organizations needed to simply have better security than other Internet-connected organizations and businesses, as many attackers would choose easier targets. However, with APTs, organizations need to be able to defeat a motivated enemy who will take the time to look for weaknesses rather than moving on to another target.

The timeframe of APTs can also make detection particularly difficult. In a standard security breach, significant amounts of data may be exported in a short period of time, making discovery of the breach possible by firewalls and intrusion detection devices. An attacker in an APT may take months or even years to export the targeted data, defeating even fully-featured and well-configured systems.

Objectives	Targets
Because of their targeted nature, perpetrators of APTs often have different objectives than standard internet hackers, including an increased focus on the following instead of simple theft and recreational damage: ▪ Political manipulation ▪ Military espionage ▪ Economic espionage ▪ Technical espionage ▪ Financial extortion	Specific types of organizations are more at risk for APTs due to the often political and state-sponsored nature of the threat: ▪ Government agencies ▪ Defense organizations and contractors ▪ Critical infrastructure systems (e.g., public utilities, communications and transportation systems) ▪ Political organizations ▪ Financial institutions ▪ Technology companies

Examples

RSA

In 2011, RSA Security announced that it had been the victim of what it defined as an APT[4]. The attackers gained initial entry by tricking an internal user to open an email that included a spreadsheet attachment that exploited a zero-day vulnerability in Adobe Flash. From there, the attackers escalated privileges, installed backdoors, and gained control of additional systems.

The attackers were able to gain access to RSA systems that held information related to their two-factor authentication tokens, known as SecurID. This information potentially included "seed" values, which RSA uses with their tokens to generate one-time passwords that change every 60 seconds. If the source code itself were stolen, attackers could look for vulnerabilities in the SecurID implementation or even the encryption itself.

Operation Aurora

Operation Aurora was an APT that targeted numerous large companies, including Google, Adobe, Rackspace, and Juniper Networks. Media reports suggest that many other companies were targeted, including Yahoo, Northrup Grumman, Morgan Stanley, Symantec and Dow Chemical.[5] It is believed that China's Politburo directed the attacks as part of a large-scale and coordinated campaign against the U.S. and other western countries.[6]

Section 2: Opportunity

Defense-in-depth

The key to defending against Advanced Persistent Threats is defense-in-depth. Given enough time, a determined attacker will be able to breach most network perimeters. A successful defense will:

1. Make the initial penetration difficult.
2. Reduce the potential for privilege escalation in the event an account is compromised.
3. Limit the damage that can be done by a compromised account, even if it is privileged.
4. Detect compromised accounts and suspicious activity early in the process.
5. Gather information useful to a forensic investigation, to be able to determine what damaged occurred, when, and by whom.

Securing the perimeter with firewalls and intrusion detection systems on the network border can only help with the first and fourth defenses! A more active protection strategy is required.

Early detection

Breaches are often detected after the attacker has gained access to an internal network and caused damage or stolen large amounts of data. At this point, APT "defense" involves a costly damage control, clean-up and continuous monitoring process. The key to affordable and manageable protection against APTs lies in detecting the threats as early as possible. In the initial phase of an attack, when an attacker first gains a toe-hold in the network, an organization can use various techniques to detect a breach, including decoupling and externalizing system security from system administration, preventing and detecting attempts at privilege escalation and unauthorized use of privilege, and auditing and recording of user activities outside the operating system logs (such auditing and recording may be unknown to the attacker).

Privileged identity management, information protection and control and internal infrastructure security form the core of an in-depth APT defense, along with early detection. These techniques are detailed in the sections below.

Privileged Identity Management

Privileged Identity Management (PIM) tools manage and monitor administrative accounts, such as "Administrator" on Windows and "root" for UNIX and Linux. PIM systems:

- Implement the principle of "least privilege," even for administrative accounts.
- Manage access to shared accounts via privileged user password management capabilities.
- Track user activities both to help ensure accountability and to assist in a security breach investigation.

Least privilege access

All people should have the minimum privileges necessary to do their jobs. While this concept is understood by many organizations, they often fail when implementing it in practice, particularly for administrative accounts. Individuals who require some level of privileged access are typically given the password to the relevant administrative account, which is shared by multiple people.

What organizations must realize with the rising prevalence of APTs is that privileged access does not need to be an "all or nothing" decision. Individuals can be granted elevated privileges to enable them to accomplish only a very specific task. In the past, this has been accomplished on UNIX and Linux systems using the "sudo" tool, but modern access control tools can grant and deny access centrally for both UNIX and Windows® systems.

Security model: decoupling security from system administration

A typical operating system has a two-layer security model: privileged users and regular users. In order to protect against APTs, however, a more sophisticated model is needed. This model is based on the standard security principles of "least privilege" and "segregation of duties." At a minimum, three primary administrative roles should be defined:

- **System administrator:** The administrator of the system itself should have the privileges required to make necessary server software updates, configuration changes, and install software. System administrators should not be able to change critical security settings or view security-related logs.
- **Security administrator:** These administrators should be able to update and change security settings and configurations and view security-related log files. Security administrators should not be able to install software or access sensitive data on a system.
- **Auditor:** Auditors need to be able to check security settings and view log files, but should not have the ability to make any changes to a system. While access to sensitive files may be required, all access should be read-only.

Additional administrator types should be created wherever appropriate, such as database administrators or for other particularly sensitive applications.

Using a multi-tiered security model accomplishes two goals simultaneously: it protects against insider threats from internal administrators by limiting what each individual can do, and also makes APTs significantly more difficult for external attackers. Instead of needing to compromise one "superuser" account, attackers will now need to gain access to multiple accounts in order to have full access to a system.

Fine-grained controls

Fine-grained controls, aside from being good security practice, are particularly helpful in mitigating the damage caused by an APT. Once attackers gain administrative privileges, they usually install backdoor "rootkits" and begin to export sensitive data. With proper access controls, an attacker with even privileged access is limited in what he or she can do, and may be prevented from accessing sensitive files, executing malicious commands, installing programs, stopping or starting services, or changing log files. On a system where fine-grained controls are implemented, an attacker may be forced to compromise multiple accounts in order to do what was previously possible with a single account.

Implementing fine-grained access controls can also mitigate the risk of one of the greatest security weaknesses in an organization: its people. Using what are called "social engineering" techniques, attackers often trick employees and other insiders into providing information that can be used to gain access to their accounts or reveal other security weaknesses. By limiting access to critical systems and data of employees, the damage that can be done by an attacker who gains access to accounts through social engineering is diminished.

Shared account management

Shared account management (or "privileged user password management") is a key defense against APTs. Gaining access to privileged identities (often through privilege escalation) is a key intermediate step in almost all successful attacks. Privileged user password management tools should be able to:

- Securely store encrypted passwords.
- Manage password complexity and regular automated changes according to policy.
- Restrict access to administrative accounts by requiring all accesses to go through a centralized portal.
- Use "automatic logon" functionality to prevent even authorized users from knowing the passwords to privileged accounts.
- Provide emergency account access, which has additional controls and required approvals.
- Eliminate the use of hard-coded passwords in scripts (which are often stored in clear-text and can be stolen by a malicious user).

These capabilities not only prevent passwords from being shared, but also prevent password theft from personal password files or through keystroke logging. By requiring all privileged account logins to go through a central proxy, an organization can track all logins and activities in the event of a breach, helping investigative efforts and potentially mitigating damage.

User activity reporting

Understanding what actions are being performed by privileged accounts is a key component to detecting APTs and mitigating the damage in the event of a successful initial attack. By their nature, APTs usually involve the exporting of significant amounts of data, which may be detected by the right tools. User activity logs prove what system and user activities are taking place on a system or a network device and can be used to identify policy violations and investigate security breaches.

Regulations like HIPAA, CA SB 1386 and the numerous state breach notification laws require an organization to disclose the security breach to the person or organization affected. User activity logs can be used to investigate the security breach to find not only who did what but also how it happened so that internal controls can be fixed and processes can be improved.

User activity reporting tools should be able to do the following:

- Track all:
  - Logins, particularly for privileged and shared accounts, including the source IP, the original User ID that accesses a shared account, time and date of both login and logout
  - Shared account activities back to the original User ID
  - Commands, whether entered via a command line or GUI

- Detect anomalous behavior:
    - Identify suspicious activities and generate alerts.
    - Provide log correlation capability, focusing on connecting user activity with the individual who performed it via analysis of complex patterns of audit logs.
- Investigate breaches:
    - Prove "who did what" in a shared account environment.
    - Deliver visual log analysis tools with drill down capabilities that can expedite the investigation of user and resource activities and the identification of policy violations.

In the event of a breach, these abilities will help an organization understand:

- How an attacker was able to gain access to an account
- What they did while using that account and what damage was done
- How to prevent future attacks using the same or similar methods
- Potentially who the attacker was and where he or she came from
- What information to report to regulatory agencies

It is critical to remember that logs must themselves be protected from administrators. Privileged users can determine where logs are stored locally on systems and can discover the auditing policies used within the organization. They can cover their own tracks by deleting records within local log files since they have complete access to the systems (if proper fine-grained controls are not implemented). Organizations should store logs in a remote location which cannot be accessed by those privileged users and also monitor if attempts are made to delete the local log files across systems.

Information protection and control

In an APT, the end goal of the attack is to steal sensitive information, so having control over data is an essential component to a successful defense. To protect sensitive data from an APT, an organization should protect and control data in four states:

- **Data at-access.** Sensitive information attempting to be accessed by an inappropriate role
- **Data in-use.** Sensitive information handled on the local workstation or laptop
- **Data in-motion.** Sensitive information communicated over the network
- **Data at-rest**. Sensitive information stored in repositories such as databases, fileservers or collaboration systems

To achieve this, organizations must define policies to enforce control if inappropriate access or usage of the data is detected. Once a policy violation occurs (such as attempting to access intellectual property, copying the information to a USB drive or attempting to email it) the solution should mitigate the compromise while generating an alert.

Information classification is at the heart of any data security initiative. Without understanding what the information is and where it is located, it is impossible to implement a comprehensive data protection program. An organization must accurately discover and classify sensitive information based on its level of sensitivity to the organization. This includes intellectual property, but also personally identifiable information, private health information, and other non-public information.

Once information has been properly classified, policies have been defined, and controls have been deployed, an organization can then monitor and control the access and handling of all sensitive information. This includes user actions from simply attempting to access and read sensitive data, to copying to a removable device or printing, to emailing outside the network, to discovering data stored in a repository such as SharePoint.

Internal infrastructure security

While securing the network perimeter and privileged identities and data are essential components of an in-depth APT defense, it is also important to secure the internal IT infrastructure. In addition to appropriate network architecture and segmentation, this includes properly configuring and securing individual servers and devices, and their environments.

Unexpected and externalized security

Attackers strategize and employ tactics against known security defenses. They also use common operating system commands, functions, and utilities to gather information, monitor the system, and take action to expand their control. Security professionals can use attackers' basic assumptions against them by adding unexpected elements to a system. For example, files and commands that appear not to be protected or monitored by system logs can be both protected and monitored by an external tool. In effect, the permissions that an attacker sees are not necessarily the permissions that are being enforced. This enables an organization to detect an attacker checking operating system permissions and violating external policies when he or she tests permissions boundaries.

This is the critical reason that security administration should be externalized and separated from the operating system administration. After gaining initial access to a system, a typical attacker will attempt to escalate privileges in order to bypass operating system controls. With this access, they assume that they will be able to override security mechanisms and effectively "hide their tracks." With an external security function, it is often possible to detect and contain attackers much earlier in the APT process, when an attacker attempts to escalate his or her privileges, change systems security controls or exercise privileges that have not been granted. While an attacker may successfully bypass traditional OS-level controls and logs, external detection processes can catch them unawares. In essence, an organization can implement an access control policy behind the scenes—in a powerful and unexpected way.

In addition, standard system commands can be changed and modified. If administrators rename functions such as "sudo," all attempts to use the original sudo command can trigger an alert and lead to early detection of a breach.

Server hardening

All servers hosting sensitive information should be configured in a manner that minimizes the potential for compromise and dissemination of data in the event a compromise does occur. This includes:

- Using a software firewall to control both inbound and outbound communications, restricting packets by source IP, protocol (e.g., SSH, TELNET, etc.), and TCP port; block insecure protocols (e.g., unencrypted services, such as FTP)
- Blocking all application executions and installations except when explicitly specified ("application whitelisting"), preventing code execution exploits and the installation of "backdoor" software
- "Jailing" applications. Define and allow accepted actions for high-risk applications and restrict any behavior that exceeds these bounds. For example, an ACL can be built based on a logical ID which owns Oracle® processes and services so its jailed behavior prohibits it from any actions besides starting Oracle DBMS services.
- Preventing changes to log files
- Enabling file integrity monitoring to detect changes to key files, such as those made by "root kits"
- Controlling access to sensitive application directory files (e.g., only the payroll application can open payroll files)
- Detecting changes to sensitive files in real-time

Uniform security

A common issue in distributed computing is the variance of capabilities and availability of security controls across platforms (e.g., UNIX file/directory controls are significantly different than Windows). This can lead to multiple exploitable issues:

- Security policies that cater to a system model instead of a business security model
- Security policies must accommodate systems limitations
- Errors and omissions caused by added complexity of security management

In order to provide a comprehensive APT defense, security configurations must be applied as equally as possible to all platforms. Any limitations and inconsistencies must be understood and tracked.

This is another reason that organizations should not solely rely on operating system security. External tools can provide a universal platform to apply a security paradigm across environments, allowing for a centralized, streamlined and business specific approach to security.

Virtualization security

The number of virtualized systems has skyrocketed, making virtual environments a key target for attackers in an APT. The hypervisor is also a critical target because of the level of access it can yield. If an attacker compromises the hypervisor, he or she can gain nearly complete access to all of the virtual machines running on that hypervisor. While operating system security can prevent direct logins and encryption can protect sensitive data, these measures will not stand up to a determined attacker. Someone with

administrative control over a hypervisor can copy entire virtual machines to an external environment as well as bypass host-based security using either brute-force methods or by overwriting key files.

In order to secure virtual environments, organizations must again focus on administrators and applying the principle of least privilege. First, access to privileged hypervisor accounts should be strictly controlled, with all actions monitored and logged. Second, in the same manner as physical environments, privileged hypervisor identities should be restricted to perform only required actions. For example, a finance administrator should be able to only access virtual machines owned by the finance department and not HR systems.

Putting it all together

No one security tool is going to protect an organization from an APT by a determined, capable, persistent, and well-resourced attacker. The goal of any APT defense strategy lies in making it as difficult as possible to penetrate the network, limiting the amount of damage that can be done and the amount of information that can be stolen in the event of a successful breach, and detecting a breach as quickly as possible.

While perimeter security is a required component to preventing the initial breach, it is by no means sufficient and it does little to reduce the damage after a breach has occurred. The key to mitigation lies in an intelligent combination of privileged identity management, data classification and control, and infrastructure security.

Standard privileged identity management tools can restrict or grant access based on a set of rules. While this capability can provide appropriate segregation of duties, it is an inherently rigid solution. Privileges may be modified over time as roles change, but this is an essentially passive solution.

"Content-awareness" is what is needed to usher in a new generation of active APT defense. This means integrating data intelligence into every decision made when determining whether to grant a request. This should be done by recognizing and understanding patterns of data accesses and use. For example, the following should be noticed:

- **Changes in data type access.** An administrator consistently accesses data of a specific type (e.g., operational records) then requests access to confidential financial information or customer data
- **Changes in data use.** An administrator typically accesses sensitive data via a specific application with read-only access requests to export data to an external hard drive, USB stick, or via email
- **Changes in data quantity.** An administrator accesses 100MB of sensitive data weekly and requests access to 500GB in that same time period
- **Changes in data access frequency.** An administrator accesses highly-confidential data once a month and suddenly accesses the same data daily

None of these changes may in and of themselves indicate that a breach has occurred; however, they do represent a change in behavior. A system that intelligently controls privileged user access should take all of these factors into account when reviewing an access request. This data intelligence can either be used to deny access to resources in real-time or to allow access but create an alert indicating suspicious activity.

Section 3: Benefits

Reduce your risk!

Organizations targeted by an advanced persistent threat face multiple types of damage. Attackers may steal intellectual property and strategy documents, potentially affecting competitiveness. The theft of customer data can lead to customer backlash, reputational harm, and legal action. Stolen private health information or financial records may lead to regulatory compliance issues.

A secondary benefit of a holistic program to defend against advanced persistent threats is that it helps protect an organization from other threats, from automated external attacks to insider threats. Many of the techniques employed to mitigate the damage of APTs also limit the access given to internal accounts, including administrators. By limiting access and segregating duties for even privileged users, an organization is protecting itself against a rogue administrator or other malicious internal user.

What is unique about this approach is that it does not require specific knowledge of vulnerabilities and new exploits and does not rely on perimeter defense. Using these techniques, organizations can apply a security model, and allow or deny actions based on business rules, data sensitivity and anomalous behavior. Because this model can be applied uniformly across platforms and can be separated from operating system security, it can provide an effective means of defending against APTs and detecting attacks early in the process.

Section 4:

Conclusions

Targeted attacks are growing in prevalence. The breaches of companies such as RSA have been highly-publicized and will have far-reaching consequences, both to reputation and profits.

The idea of defense-in-depth is not new. It is a fundamental aspect to any security program. What is new is the focus on securing internal privileged identities in order to prevent damage done by outsiders. With the network perimeter no longer the bastion of security that it once was, identity has become even more critical. In essence, “identity is the new perimeter.”

When using identity to secure against both internal and external threats, such as APTs, “content-awareness” should be a key requirement. By using data intelligence as part of every access decision, today’s organizations can better understand the risks associated with every action a user takes. Access requests to sensitive data can be analyzed and understood with far more context than ever before. Instead of relying on fixed rules to allow or block certain actions, data can be used to create a clearer picture of user activity.

To help your organization get ahead of the game when it comes to defense against targeted attacks, embrace privileged identity management and content-awareness as the cornerstones of your security program.

Section 5:

References

1 http://taosecurity.blogspot.com/2010/01/what-is-apt-and-what-does-it-want.html

2 NIST Special Publication 800-30 Revision 1, Guide for Conducting Risk Assessments, http://csrc.nist.gov/publications/drafts/800-30-rev1/SP800-30-Rev1-ipd.pdf

3 "Advanced Persistent Threat", Wikipedia, http://en.wikipedia.org/wiki/Advanced_persistent_threat

4 http://www.rsa.com/node.aspx?id=3872

5 http://en.wikipedia.org/wiki/Operation_Aurora

6 http://www.nytimes.com/2010/11/29/world/29cables.html?_r=2&hp

Section 6:

About the author

Russell Miller has spent over eight years in network security in various roles from ethical hacking to product marketing. He is currently a Director of Product Marketing at CA Technologies, focused on privileged identity management and data protection. Russell has a B.A. in Computer Science from Middlebury College and an M.B.A. from the MIT Sloan School of Management.