

Realtime
publishers

Advanced Persistent Threats and Real-Time Threat Management

The Essentials Series

sponsored by



TREND
M I C R O™

Dan Sullivan

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers.....	i
Article 1: Beyond the Hype: Advanced Persistent Threats	1
APTs Today	1
The Evolving Threat Landscape.....	2
Elements of APTs.....	3
Changing Business Practices that Compound the Problem	3
Pragmatic Assessment of the Potential to Control APTs.....	4
Summary	5
Article 2: Need for Real-time Management and Responding	6
Limits of Standard Endpoint and Perimeter Security Controls.....	6
Stages of Response to a Breach	8
Ideal and Realistic Assessment of Preventing a Breach	9
Summary	10
Article 3: Planning for Real-time APT Countermeasures.....	11
Business Case for Real-Time Threat Management	12
Assessing the Current State of Readiness for Real-time Threat Management	12
Planning the Deployment of a Real-Time Threat Management System	13
Controls for Blocking	14
Controls for Monitoring.....	14
Containment Mechanisms.....	15
Summary	15

Copyright Statement

© 2011 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

Article 1: Beyond the Hype: Advanced Persistent Threats

Businesses face a constantly evolving threat landscape. One of the greatest challenges is presented by advanced persistent threats (APTs), which are sophisticated, multi-faceted attacks targeting a particular organization. Mitigating the risk of APTs requires advances beyond traditional layered security to include real-time threat management. This Essentials Series describes the nature of APTs, the risks they pose to businesses, and techniques for blocking, detecting, and containing APTs and other emerging threats. We begin with a pragmatic assessment of the nature of APTs, specifically:

- The nature of APTs today
- The continuously evolving threat landscape
- Elements of APTs
- Changing business practices that compound the problem
- Assessment of potential to control and mitigate the risk from APTs

Clearly, the threat landscape continues to become more challenging. The motivation and means for carrying out attacks on information systems is changing. Determined, committed attackers are employing multiple means to breach security controls. Businesses need to respond in kind with multiple security controls, including real-time monitoring and rapid containment measures.

APTs Today

APTs are sophisticated, multi-faceted cyber-attacks targeted at a particular organization. Such attacks are advanced in terms of the techniques that are applied and the insider knowledge the attackers have about their targets. APTs may use multiple vectors, such as malware, vulnerability scanning, targeted hacking, and malicious insiders to compromise security measures. APTs are long-term, multi-phase attacks. Early stages of an APT attack may focus on gathering information about network configuration and server operating system (OS) details; later, efforts may focus on installing rootkits or other malware to gain control or establish communication with a command and control server. Later stages of an attack may focus on stealing intellectual property by copying confidential or sensitive data.

It is important to understand that APTs are not a new means of conducting an attack and are not something that can be blocked or disrupted once and the problem goes away. APTs are better understood to be more like a cyber-attack campaign than a single type of threat; think ongoing processes. An antivirus program may block malware used in an APT attack but that does not mean the attack is stopped. By its very nature, an APT is an ongoing attack. If one tactic does not work, another will be attempted. Realistically, we should not be thinking in terms of a single countermeasure or even adding more layers to a layered security strategy; rather, we should be thinking of processes that together can block when possible and detect and contain breaches in other cases. It's reasonable at this point to ask, How did we get here?

The Evolving Threat Landscape

Businesses and governments face an evolving threat landscape. What began with attempts to gain bragging rights about defacing a major newspaper's Web site or blocking service to a popular site with a Denial of Service (DoS) attack has shifted to attacking for financial gain. Attackers can realize direct financial gains by fraud and intellectual property theft or indirectly by disrupting a competitor's ability to deliver services or conducting a widely publicized data breach that compromises customer private financial information. Besides the changes in motivations, there are changes in the means of implementing attacks.

Changes in application architectures and the decentralization of core operations create opportunities for attackers. In the past, bank tellers and ATM machines were the only ways to conduct transactions with your bank accounts—now you can do it with your phone. It was not that long ago that talk about retailers invoked images of brick-and-mortar stores and malls; now it is just as likely to bring to mind Web sites that sell everything from books to appliances. The Web applications that provide many of the services businesses offer implement workflows that ultimately lead to back-office systems like inventory management and accounts receivables. These can readily become the target for vulnerability scans, injection attacks, and other probes that reveal information about the application architecture and potential vulnerabilities.

Another factor in the evolving threat landscape is the combination of techniques that may be used. Malware can be used to perform a specific task, such as capture keystrokes, or it may include a communications module that works with a command and control server to download instructions allowing attackers to probe, make discoveries, and adapt their tactics to their findings.

Some of the techniques we see in APTs we have seen in the past with blended threats that used a single attack vector to deliver multiple forms of malicious software. We also see attacks will change in response to countermeasures. When antivirus software successfully detected viruses using pattern-matching techniques, malware developers employed encryption and polymorphic techniques to scramble their code enough to avoid detection. Similarly, if one route of entry in a system is blocked, an APT will look for another. The dynamic nature of APTs is a common characteristic of security threats, but there are characteristics that distinguish APTs from other types of attacks.

Elements of APTs

At the most basic level, there are three characteristics of an attack that make it an APT:

- Motivated by financial gain or competitive advantage
- A long-term, sustained attack
- Targeted at a specific company, organization, or platform

Businesses and governments are the targets of APTs for obvious reasons. Businesses have both financial assets and intellectual property that are highly valued. Governments have faced outside aggression probably for as long as there have been governments—thus, the concept of APTs is in many ways nothing new. What is new is that the means of executing such threats have moved into the realm of networks and applications.

Long-term attacks may continue for days, weeks, months, or even longer. APT attacks can begin with intelligence gathering, which may continue for some time. It may involve both technical and human intelligence gathering. The intelligence gathering efforts can shape later stages of attack, which can be either quick or prolonged. For example, an attempt to steal trade secrets may take months of intelligence gathering about security protocols, application vulnerabilities, and file locations but take only minutes to execute once a plan has been established. In other cases, attacks may continue over longer periods of time. For example, after successfully deploying a rootkit on a server, an attacker may regularly send copies of potentially valuable files to a command and control server for review.

A number of widely publicized APT attacks demonstrate the breadth of means and motivations driving the deployment of APTs:

- The Zeus botnet, for example, started as a platform for attacking financial institutions but was changed to become a framework for other types of APTs.
- The Aurora APT attacked Google and other technology companies seemingly in an attempt to gain access to and possibly modify application code.
- Stuxnet is highly specialized industrial malware that includes a rootkit for a programmable logic controller used in industrial equipment. There has been speculation in the press that Stuxnet was developed by one or more governments.

APTs such as these can take advantage of changes in the way we deliver services.

Changing Business Practices that Compound the Problem

Changes in technology and motivations for attack are only part of the reason APTs have become such a significant threat. The way we architect systems and allow access to business applications is also part of the puzzle.

Consider de-perimeterization. In the past, firewalls would have blocked traffic that was not specifically allowed. As applications advanced, there was more need for more flexible movement of network traffic. Outsiders needed access to internal resources. Developers wrote applications to tunnel blocked traffic over protocols that were allowed through (that is, HTTP). Rather than having a single boundary around all network assets, businesses opened access to more servers and depended on device-based controls and network traffic monitoring.

Another factor that can be exploited by APTs is the increased use of mobile and other unmanaged devices. IT departments do not always dictate the kinds of anti-malware software or access controls that must be in place before a device can be used with internal services. These devices can be used by APTs to stage part of an attack on a business or government network.

Similarly, the increased use of publically available Web applications provides another potential method of attack. For example, an injection attack on a Web application could be used to collect intelligence about the contents of databases as well as the structure of the application.

By expanding employee access to critical information infrastructure, businesses can make it easier and more efficient for employees to perform necessary tasks. However, doing so also increases the potential points of entry for attackers.

Technical and organizational factors are at work with regards to the potential for executing an APT attack. Many of these factors, such as empowering employees and accessing applications from mobile devices, are so beneficial that it is difficult to imagine curtailing them. We can mitigate the risk of APTs without necessarily sacrificing these and other advances.

Pragmatic Assessment of the Potential to Control APTs

From a pragmatic perspective, it is reasonable to assume that APTs will be with us for the foreseeable future. The history of cyber-security is filled with examples of new forms of attacks emerging in response to new types of controls. APTs are long-term process-oriented attacks that are a product of changes in the motivations of attackers and the means available to them to conduct their attacks. Given that APTs are here to stay, what is the appropriate strategy to mitigate the risks associated with them?

We should continue to deploy blocking countermeasures. Anti-malware, encryption, vulnerability scanning, and patching are all good practices. They are not enough, though, to counter APTs, so we should assume there will be a breach. This is not to say there are problems with those countermeasures; this perspective only recognizes the fact that a determined, persistent attacker may find a way to bypass blocking measures.

Working with the assumption that there will be a breach at some point, we must monitor network traffic and host activities in real time. Once a breach occurs, it is imperative to detect that breach as soon as possible and to contain the impact. Containment can include isolating compromised devices, shutting down services, and collecting data for forensic analysis.

Summary

APTs are a class of security threats that pose particular challenges to IT and security professionals. Motivated by financial or other long-term gain and armed with a wide array of malware and hacking techniques, these attackers are willing to spend the time and effort required to breach an organization's defenses. Many of the best practices used in the past are still required today, but as we shall see in the next article, we need to add real-time monitoring and containment techniques to our set of countermeasures.

Article 2: Need for Real-time Management and Responding

Ideally, we can deploy security controls that would prevent a successful attack by an advanced persistent threat (APT), but we should be pragmatic in our assessment. APTs are multifaceted and although one countermeasure, such as an antivirus system, may block one part of an APT, there can be other elements of the attack that do not depend on detectable malware. Just consider a malicious insider who uses social engineering to discover the password to an administration account of a document management system in order to copy the contents of the repository and mine them for intellectual property. When planning a response to the threat of APTs, we should assume there will be a breach at some time. The overall goal of risk management in this case is to minimize the impact of threats by blocking when possible and detecting and containing when not—to do that, we need real-time monitoring and remediation mechanisms.

This article considers the need for real-time threat management and response, specifically:

- The limits of conventional endpoint and perimeter security controls
- The stages of a response to a breach by an APT
- Ideal and realistic assessments of preventing a breach

As in the first article in this series, a dominant theme is the assumption that we should take the threat of APTs seriously and plan for a breach. This is not to say all businesses will be the victims of an APT attack or that all APT attacks will be successful. From a purely pragmatic perspective, it is better to be prepared for a breach and not suffer one than being unprepared if a breach does occur.

Limits of Standard Endpoint and Perimeter Security Controls

Standard endpoint and perimeter controls can work well to block opportunistic and unsophisticated attacks, but APTs are designed to circumvent these countermeasures. For example, an attack can begin with the identification of employees with access to key information systems followed by spear-phishing and other social engineering techniques. The goal at this stage of the attack is to lure the victim into installing malicious software under the guise of some legitimate operation, such as clicking on a link in an email to access a form or retrieve content.

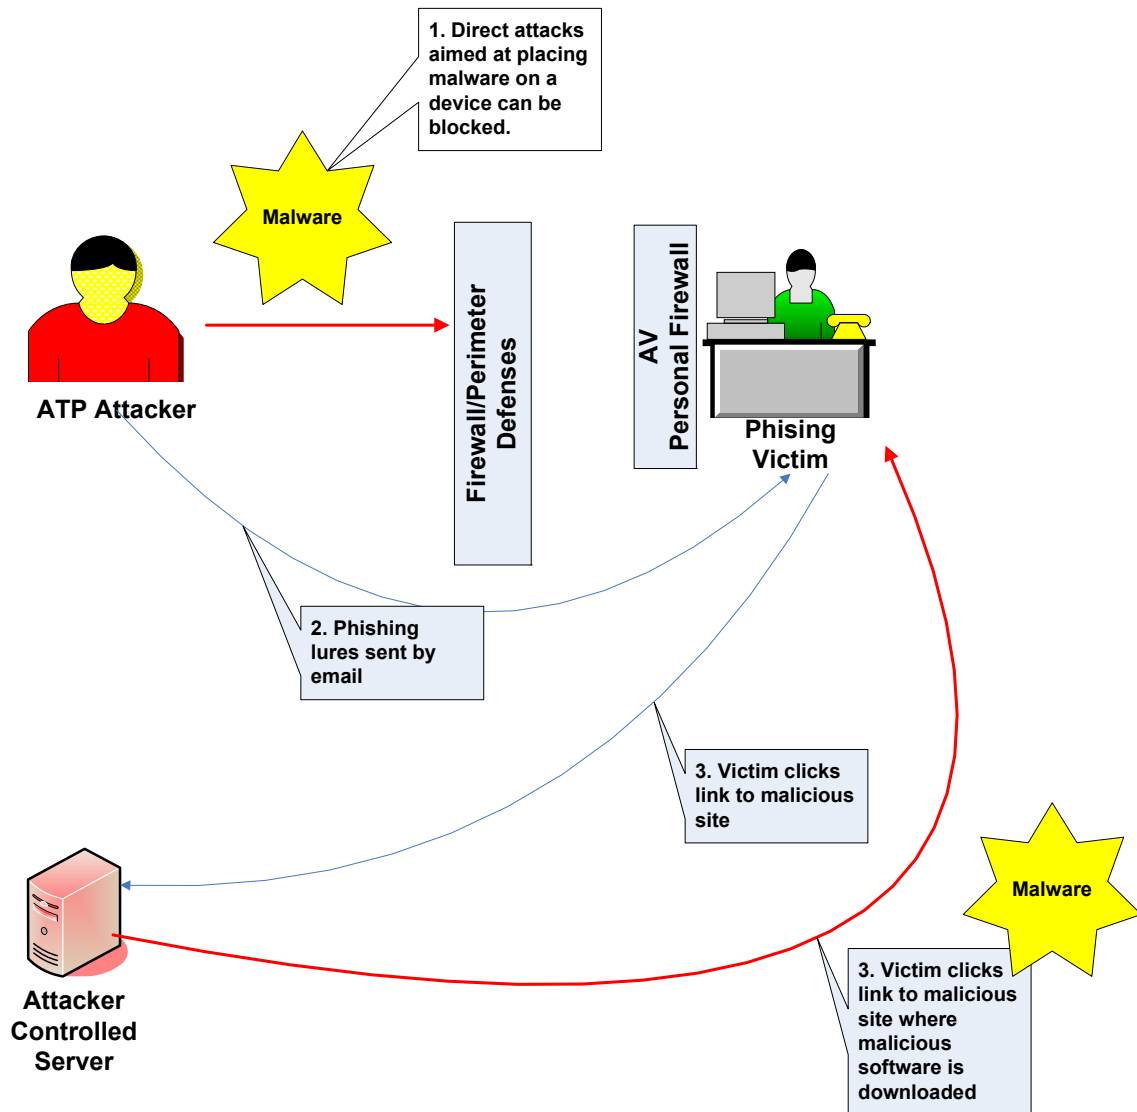


Figure 1: APT attacks use phishing to circumvent perimeter and local device security measures.

Once an attacker has a victim using an attacker-controlled server, the attacker can download malware. Attackers can use encryption and other techniques to avoid detection by pattern-matching-based systems, making it difficult to determine whether the content a user is downloading contains malicious software.

Getting malware into a victim's device is just the first step in an APT attack. The sooner such a breach is detected, the better the chance of containing the damage. This is why we need real-time threat management.

Stages of Response to a Breach

There are four stages to responding to a breach by an APT:

- The initial point of entry
- Compromise of systems and information
- Discovery of a breach
- Containment of a breach

The key stage from a risk management perspective is discovery. Assuming an APT attack has successfully avoided or bypassed perimeter, network, and local defenses, it is then a question of how long the attack continues before it is detected.




Figure 2: Within minutes, a significant proportion of attacks move from breaching security controls to compromising systems or information. Many of these attacks are not discovered for weeks or months (Source: [Verizon 2011 Data Breach Investigations Report](#)).

Two facts about the Verizon data breach study statistics are worth highlighting. First, a significant percentage of attacks lead to a compromise within minutes of a breach. The speed at which APTs operate means that responses that require manual intervention will be too late in many cases. This is why real-time management is required. Often there is no time to waste in initiating a response.

The second fact that we should pay particular attention to is the significant number of attacks that require weeks or months to discover. Until an attack is discovered, it cannot be contained. Some attacks may be point-in-time attacks in which data is stolen or some other malicious act is performed and then the attack terminates. Other attacks could go on as long as they are not detected, for example, streaming customer credit card data to a command and control server 24 hours a day. Active, constant monitoring and analysis is required to discover breaches as soon as possible.

Ideal and Realistic Assessment of Preventing a Breach

As noted earlier, ideally, security controls such as antivirus and perimeter controls would be sufficient to mitigate the risk of a security breach, but it is simply not the case. Attackers understand how perimeter controls and antivirus systems work; and they work well in many cases. The proof of this is the fact that attackers choose to avoid confronting antivirus and perimeter controls by going around them. After all, why bother trying to devise sophisticated malware that can avoid detection when you can use social engineering to trick a legitimate user. Phishing attacks exploit the fact that some users have sufficient access privileges to targeted systems and data. With a sufficiently well-crafted phishing lure, attackers can get these users to unintentionally act as a conduit to reach their target. As humans are sometimes the weakest link in a security strategy, we have to develop strategies that accommodate those weaknesses and mitigate the risks they pose.

A pragmatic approach seeks to prevent a breach and reduce the impact of a breach should one occur. This requires a three-part approach.

First, keep security controls in place and up to date. These include antivirus, encryption, access controls, and vulnerability scanning. Zero-day threats will not be detected by vulnerability scanners, so advanced network monitoring is required to detect and block intrusions. This leads to the second requirement.

Networks and servers should be continuously monitored for signs of a breach. This should include:

- Network traffic analysis
- Server log analysis
- Host intrusion prevention
- File integrity monitoring

Comprehensive monitoring can help detect footprints of an attack, such as an unusual amount of traffic between a server and an external IP address in the middle of the night or the creation of a server account with elevated privileges.

The third requirement is to contain the impact of a breach. Techniques such as virtual patching and automated remediation can disrupt an attack and prevent the vulnerability that enabled the attack from being exploited again. The specific steps that should be executed in order to contain a breach should be defined in a set of risk management procedures.

The overall objective here is to reduce the time between the point of entry of an attack and the point of containment. Real-time threat management, which includes both monitoring and response mechanisms, is required to address the threats posed by APTs. The value of real-time threat management lies in the value of data *not* lost or compromised because containment occurs faster than it would have if manual procedures were required to discover and contain the attack.

Summary

Commonly used endpoint and perimeter security controls are insufficient to block APT attacks. Phishing and other forms of social engineering allow attackers to circumvent those controls by luring users with sufficient access controls into inadvertently being used in the attack. APTs can rapidly move from the point of breach to the point of compromise, often within minutes. Manual intervention to detect and contain APT attacks is often too slow to be effective. Real-time threat management is needed to respond as rapidly as the APT attack progresses.

Article 3: Planning for Real-time APT Countermeasures

Advanced persistent threats (APTs) have emerged as a significant threat to businesses, governments, and other organizations. The previous two articles in this series have examined technical aspects of APTs and the challenges to mitigating the risk of an APT attack. APTs are not just malware and they cannot be stopped with just antivirus or perimeter controls. APTs employ social engineering techniques designed to circumvent conventional blocking defenses. Rather than try to outsmart an antivirus program, an attacker gets around the antivirus system. When an employee willingly follows a link in a phishing lure email and downloads what appears to be a legitimate program but is in fact encrypted malware, there is little chance of blocking it. Users have access control rights to download and save applications. Pattern-based detection techniques do not detect encrypted malware. In summary, conventional perimeter and endpoint defenses will not stop an APT.

To be clear, perimeter defenses and endpoint security are necessary to address the risks posed by APTs, but they are not enough. We need real-time threat management. Before deploying such controls, it is advisable to assess the current state of hardware, software, and security controls, prioritize assets, and perform a gap analysis. The results of these efforts will help to plan what proactive controls should be deployed.

This article is organized around basic steps to plan for the deployment of real-time threat management to mitigate the risk of APTs:

- Developing a business case for real-time threat management
- Assessing the current state of readiness for real-time threat management
- Developing a deployment plan

Not surprisingly, some of the recommendations that follow would fit equally well when describing other types of countermeasures. APTs are a collection of well-established techniques used for malicious purposes applied in methodical and comprehensive ways. Countermeasures used in the past can still be useful here. The key distinguishing characteristic of APTs is the speed at which they can progress. This, in turn, drives the need for real-time threat management to complement perimeter and endpoint defenses.

Business Case for Real-Time Threat Management

Executives and IT managers have no shortage of competing demands for resources. Why when a business has invested so much in antivirus, network filtering, identity management, and other security controls should they focus additional resources on real-time threat management? The short answer is because those countermeasures are not enough.

The risk from APTs is well documented. Well-publicized cases, such as Stuxnet, Zeus, and Aurora show that APTs can threaten financial to industrial control systems as well as businesses and governments. The success of these attacks also speaks to the limitations of widely used layered security mechanisms. Again, these mechanisms are essential, but they are not sufficient to mitigate the risk from APTs. APTs are designed to use human and technical resources to collect intelligence, probe for vulnerabilities, and plan multiple-step coordinated actions against a target. The techniques used in APTs are chosen precisely because they can either compromise or avoid such security measures.

The business case justification for real-time threat management is a pragmatic one: APTs exist, organizations with information, financial resources, or intellectual property of sufficient value are potential targets, and commonly used layers security defenses are insufficient to block a sophisticated attack. In addition, once a breach occurs, damaging acts can take place within minutes in many attacks. A well-planned and executed response that requires hours or days to implement may be as effective as no response at all. APTs can operate sufficiently fast enough that automated responses triggered by constant monitoring is required.

Assessing the Current State of Readiness for Real-time Threat Management

Once the business case for deploying real-time threat management has been made, the next step is to assess the current state of readiness. This involves three steps:

- Inventory IT infrastructure
- Prioritize assets
- Perform a gap analysis

The final product of this stage is a description of the potential weak spots in current security controls. Real-time threat management does not replace perimeter or endpoint defenses, it complements them. When endpoint and perimeter defenses are up to date and deployed throughout a network, the attackers have to go to greater lengths to successfully breach the infrastructure.

An inventory of IT infrastructure includes:

- Hardware and network infrastructure
- Software, especially enterprise applications
- Database, content management systems, and other repositories
- Security controls

The purpose of the inventory is to understand what can be a target of an attack or exploited in an attack. Network management and asset management tools are available that can discover assets on a network and produce an inventory of both hardware and software on those systems.

With an inventory in hand, the next step is to prioritize assets. Not all applications, servers, or other infrastructure are created equal. The object is to group assets according to their relative importance so that resources can be allocated to the most important assets first.

We should also understand where there are gaps in the current configuration of layered security controls. In particular, what security controls are missing with respect to blocking, detecting, and containing attacks? Do any of the controls in place support real-time threat management? For example, are log analysis tools capable of operating in a real-time manner? What is the delay between an event being logged and an alert being triggered?

Also consider whether governing policies and procedures are adequate for real-time threat management. They should include specifications for how to respond to a suspicious event as well as who (and what automated controls) should be involved with a response. At the conclusion of these steps, you will be in a position to plan the deployment of a real-time threat management system.

Planning the Deployment of a Real-Time Threat Management System

As you plan your real-time threat management system and evaluate candidate systems, consider three key requirement areas:

- Controls for blocking
- Controls for monitoring
- Containment mechanisms

Controls for Blocking

Blocking network attacks is a complex operation and requires a number of types of controls. Network-level malware detection should be deployed even when antivirus is deployed on endpoints. This type of redundancy is helpful when one of the instances of the control is bypassed or compromised. Vulnerability scanning will help to detect weakness in applications. There are different types of vulnerability scanning. For commercial or open source applications, vulnerability scanning can help to maintain appropriate patch levels and mitigate the risk of attacks using known vulnerabilities. For custom applications, vulnerability scanning can help identify potential points of injection attacks, especially SQL injection attacks. As helpful as vulnerability scanning can be, it does not address the problem of zero-day attacks, which exploit as-yet-publically-unknown vulnerabilities in applications.

Compliance verification procedures should also be implemented. Such procedures can help detect configurations that do not meet minimal security control standards.

Controls for Monitoring

Real-time threat management requires a number of types of monitoring mechanisms:

- Network-level analysis
- Log analysis
- Host intrusion prevention
- Blacklisting of known command and control servers

Network-level analysis demands advanced techniques to adequately identify anomalous patterns without generating too many false alarms. A combination of heuristic rules and statistical pattern recognition techniques may improve overall performance by leveraging the strengths of both while compensating for each technique's weaknesses.

Like network analysis, log analysis must be sufficiently accurate and precise to minimize both false positives and false negatives. It must also scale to meet the volume of logs that are generated in your site, so consider performance and throughput when evaluating this and other analysis tools.

In addition to monitoring network traffic and logs, critical servers should be monitored. By establishing a baseline of activity on a server, host intrusion prevention can help detect anomalous activity on a server, such as unusually high volumes of I/O or changes to application libraries. File integrity checks should also be included in this type of monitoring.

Do not forget to monitor higher levels of network traffic and, in particular, block access to known malicious servers. A real-time threat management application should ideally provide access to up-to-date blacklists on known command and control servers that could be used to direct parts of an advanced attack on your network.

Containment Mechanisms

In the event of a breach, a real-time threat management system should be able to automatically remedy the situation. This can include isolating compromised devices on the network and patching known vulnerabilities. Containment mechanisms should also support risk management procedures, such as generating alerts and escalating notifications according to the severity of events.

Summary

APTs present a new set of challenges from a security perspective. APTs are designed to circumvent commonly deployed security controls. They are also noteworthy for the time that attackers are willing to invest in collecting intelligence and probing for vulnerabilities. Conventional perimeter and endpoint security controls are necessary but not sufficient to prevent the full range of threats posed by APTs. Real-time threat management that entails blocking, detection, and containment can help mitigate the damage that can be done by fast-moving APTs that can progress from breaching controls to compromising systems and data in a matter of minutes.