A Websense® White Paper

ADVANCED PERSISTENT THREATS AND OTHER ADVANCED ATTACKS:

THREAT ANALYSIS AND DEFENSE STRATEGIES FOR SMB, MID-SIZE, AND ENTERPRISE ORGANIZATIONS

REV 2

websense®

Table of Contents

Abstract 2

Overview of APTs. 2

APT Characteristics. 2

The APT Process 3

The Malware Adoption Life Cycle 5

APT Defense Requirements 5

Strategies and Tactical Defenses 6

Websense APT Defense Tactics 9

Abstract

APT: A buzzword or an imminent threat? Advanced Persistent Threats (APTs) have become a major concern for IT security professionals around the world, and for good reason. Recent attacks targeting Canadian government officials, French government officials, RSA, and elements of the European Union have all been linked to APTs. But what exactly is an APT? Too much hype has clouded the facts surrounding a very real danger for organizations *of all sizes*. This paper clarifies the nature of APT risks and provides recommendations on how organizations can better protect themselves. More specifically, it:

- Provides a practical understanding of APTs for security professionals
- Analyzes how APT methods are used to steal confidential business data
- Outlines best-practice APT security strategies and tactics
- Describes Websense's unique defenses against APTs

Overview of APTs

Most security researchers agree that the term "Advanced Persistent Threat" was first coined by the U.S. Air Force, circa 2006, to describe complex (i.e., "advanced") cyber attacks against specific targets over long periods of time (i.e., "persistent"). Originally, the term was used to describe nation-states stealing data or causing damage to other nation-states for strategic gain. Since then, the definition has been expanded (or some might say, hijacked) by security vendors and media to include similar attacks carried out by cybercriminals stealing data from businesses for profit. We have seen attackers go after customer records, blueprints, product roadmaps, source code, and other confidential information.

Whether the term APT applies strictly to nation-state attacks or also to cybercriminal efforts to steal data from corporations is purely a matter of semantics. From a practical perspective, the important thing for security professionals to understand is that the ***same APT techniques*** used by nation-states for strategic gain are now used by cybercriminals to steal data from businesses for financial gain. A recent breach at RSA was one of many recent targeted attacks that employed APT techniques. Although the perpetrator may not have been a nation-state in this case, this was a targeted, complex attack that occurred over a long period of time.

> ***"... the important thing for security professionals to understand is that the same APT techniques used by nation-states for strategic gain are now used by cybercriminals who steal data from businesses for financial gain"***

The bottom line is that whether you work for a government agency or a private business, you need to clearly understand and protect against APT techniques. In the next two sections, we look at typical APT characteristics, the APT process, and the malware adoption life cycle — from well-funded, newly developed methods to their downstream application by common cybercriminals who use these techniques to steal data for other objectives (data collection, property damage, etc.). Once we understand how these threats work, we recommend new strategies and tactics to defend against them.

APT Characteristics

Targeted: APTs target specific organizations with the purpose of stealing specific data or causing specific damage. This stands in direct contrast to most historical malware, which wreaks havoc on any randomly infected system. The Aurora/Google attack targeted source code (with possible political motives). The Sony attack targeted personally identifiable information (PII). The RSA attack targeted intellectual property. These were not opportunistic attacks victimizing just any organization with vulnerability to a given exploit. These were focused campaigns by perpetrators willing to invest time and money to achieve specific objectives. There are two conclusions here. **First, any organization, large or small, with valuable data is subject to APT methods.** Second, the more valuable your data, the more likely you are to be targeted. The cybercrime economy is well organized and funded, with attackers investing more to achieve bigger paybacks.

Persistent: APTs play out in multiple phases over a long period of time. Prior to the actual attack, attackers only know the target organization and objective. They do not know where their target data resides, what security controls are in place, or what vulnerabilities exist that might be exploited. To steal the data, the attacker must identify vulnerabilities, evaluate existing security controls, gain access to privileged hosts within the target network, find target data, and finally, extract data from the network. The entire process may take months or even years. The lesson here is that attack detection cannot rely on any single event, but should look for patterns of events that are characteristic of APT methodologies.

Evasive: APTs are systematically designed to evade the traditional security products that most organizations have relied on for years. For example:

- To gain access to hosts within the target network while avoiding network firewalls, the attacker delivers threats within *content* carried over commonly allowed protocols (http, https, smtp, etc.).
- To install malware on privileged hosts while avoiding antivirus programs, the attacker writes code designed for the specific target environment. This code has never been seen before and therefore, no AV signatures exist to provide protection.
- To send data out of the target network, while again avoiding firewalls, the attacker uses custom encryption and tunnels content within protocols that are allowed outbound by the firewall.

Complex: APTs apply a complex mix of attack methods targeting multiple vulnerabilities identified within the organization. A given APT may involve 1) telephone-based social engineering to identify key individuals within the target organization, 2) phishing emails sent to those key individuals with links to a website that executes custom JavaScript code to install a remote access tool, 3) binary command-and-control code (either custom code or code generated by commonly available malware kits) and, 4) custom encryption technology. Clearly, no single security control provides coverage against all of these vectors. Any successful APT defense strategy must take a multi-layered approach in which multiple detection mechanisms work together to identify complex patterns of evasive behavior.

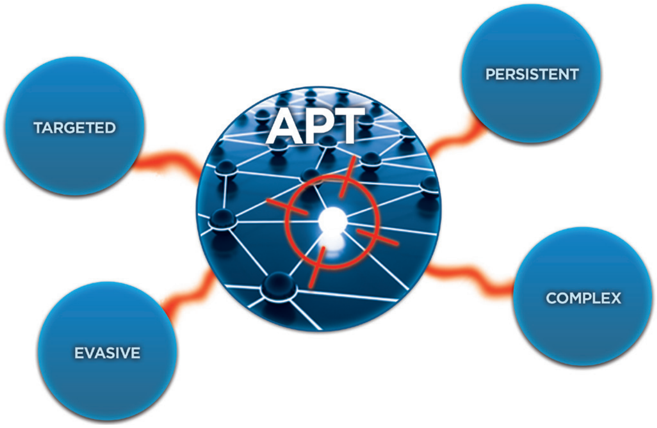

Figure 1

The APT Process

The APT process includes three major phases that occur over a period of months.

- **Phase 1 - Reconnaissance, Launch, and Infect:** The attacker performs reconnaissance, identifies vulnerabilities, launches the attack, and infects target hosts.
- **Phase 2 - Control, Update, Discover, Persist:** The attacker controls infected hosts, updates code, spreads to other machines, and discovers and collects target data.
- **Phase 3 - Extract and Take Action:** The attacker extracts data from the target network and takes action (sells data, etc.).

Attack Phase 1: Reconnaissance, Launch, Infect

The attack phase consists of three sub-phases.

Reconnaissance: Attackers research points of entry, vulnerabilities, key individuals, and key assets. This may include top-ranking executives, IT administrators, and hosts that can provide access to target resources within the organization.

Launch: This phase typically includes one or more methods aimed at gaining access to a privileged host. Targeted attacks and spear phishing keep a low profile to further evade detection. Common methods include the following.

- Email lures with embedded links to websites with zero-day malware downloads
- Emails with file attachments in common formats like Office or PDF. These attachments may include zero-day attack code targeting a previously unknown vulnerability
- Infected web sites of interest to key individuals identified by social media profiles
- Social engineering to gain access to privileged user account credentials

Infect: Custom code is typically installed onto a privileged host. This code reports back to a command-and-control location with network and other data that helps the attackers in Phase 2.

Attack Phase 2: Control, Discover, Persist

This phase can be broken into three sub-phases.

Control: The attacker remotely controls infected hosts with a command-and-control service. Although we have seen cases where this service is located on a compromised host within the target network, it is more typically found on the internet, often on a dynamic DNS host. The C&C allows the attacker to remotely update malware, add new malware (encryption tools, etc.), and send commands to the host. Although the original infection often involves custom (zero-day) attack code, we frequently see commonly available toolkits used for command and control.

Discover: At this stage, infected hosts download additional components with the ability to discover target data on the infected hosts, on mapped network drives, and in other network locations. Key targets may include Active Directory (AD) and certificate PKI servers to establish accounts and gain access privileges to confidential data within the network or cloud-based storage. Monitoring data-in-use once a user accesses it with their credentials is another discovery method, along with breaking into systems where users have administration rights. The attacker may also attempt to gain more control by discovering additional hosts within the target network and using network or other system-level vulnerabilities to infect those hosts. Very often, the tools used to gain more control are standard network tools such as gsecdump, Cain&Abel (to crack passwords), SSH, and RDP.

Persist: A key difference between traditional malware and an APT is the ability to persist. Traditional malware will often remove itself or be removed by an antivirus program once known and identified. An APT is designed to go unnoticed. Additionally, it is designed to persist by calling back to command-and-control centers for updates to download new undetected code to avoid detection by updated antivirus solutions.

Attack Phase 3: Extract, Take Action

At this stage, attackers have taken control of one or more hosts within the target network, may establish access credentials to expand their reach, and have identified target data (assuming data was the goal). The only thing left to do is send the data out of the network to either the command-and-control server or a previously unused server. This server may be located in the same location as the attacker or in a foreign country.

If new target data continues to become available (e.g., new customer records or updated business plans) and holds value for the attacker, this final phase can go on for a long time. Eventually the attack will stop, either because the attacker has achieved their goal or because the victim notices and cuts off the attack. At that point a number of consequences can result.

- **Ransom:** The attacker threatens to publicly disclose the theft if the victim does not agree to pay a ransom. The organization may pay the ransom to avoid brand damage, regulatory fines, lost customers, etc. This is a common way for the attacker to convert stolen data into cash.
- **Share or sell attack methods:** If the attack wasn't thwarted by the victim, methodologies are shared with or sold to other attackers who repeat the attack on the same and/or other targets.
- **Sell information:** If PII was stolen (names, credit card numbers, email addresses, etc.), the attacker may sell that information to other criminals who perpetrate downstream crime against those individuals whose data was stolen. An example is using a stolen credit card to make a purchase.
- **Public disclosure:** Eventually, data theft events may be publicly disclosed to the media. Typically, the victim organization chooses to disclose theft once they become aware it, or are required by local compliance regulations. However, the attackers themselves may also announce their achievement before the victim knows anything has happened.

The Malware Adoption Life Cycle

The methods developed for an APT don't always end with one attack. These techniques are often copied and applied by other perpetrators against other targets, including organizations of all sizes. Eventually, these techniques may be commoditized and turned into malware kits that are readily available to common hackers for a nominal cost.

In this respect, the life cycle of an APT may extend for many years beyond its original target and victimize hundreds or thousands of other targets.

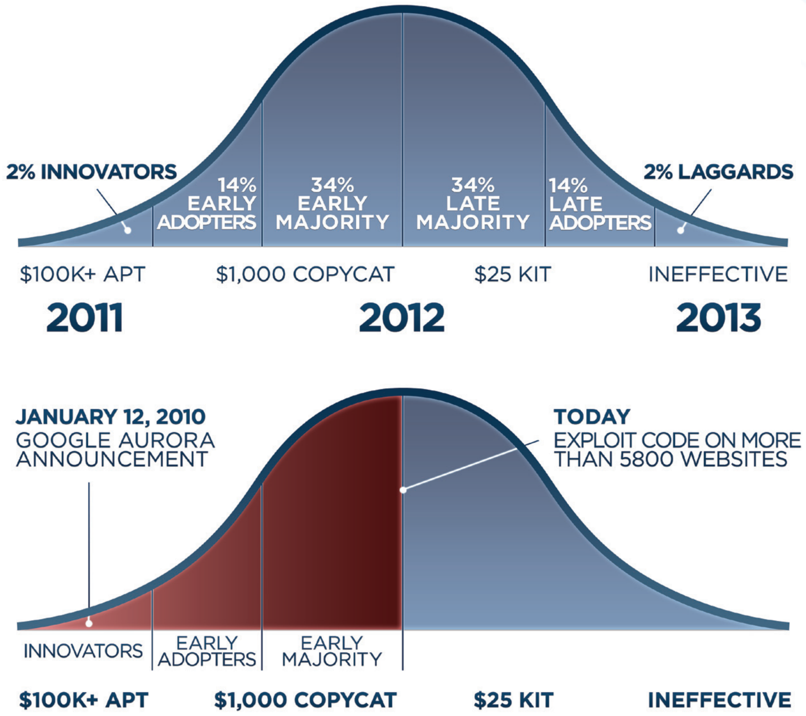

Figure 2

One example of this (Figure 2) is the exploit code from the Aurora APT (announced in 2010), which has since been detected on thousands of other infected sites.

APT Defense Requirements

By analyzing the characteristics of APTs as described above, we can describe the key requirements of an effective security solution.

- **Content-aware** – Since APTs uniformly penetrate network firewall defenses by embedding exploits within *content* carried over commonly allowed protocols, APT defense solutions require deep content awareness.
- **Context-aware** – Since most APTs use custom-developed code and/or target zero-day vulnerabilities, no single IPS or antivirus signature is likely to positively identify the threat. Without definitive attack signatures, we must rely on less definitive indicators. Although a single suspicious indicator is not enough to identify an attack, if we evaluate each suspicious indicator in the *context* of other indicators, we can amass enough evidence to reliably identify malicious activity.
- **Data-aware** – Although target organizations may not know *exactly* what an individual APT looks like (they are all unique), most organizations can identify their own sensitive data. Therefore, data loss prevention (DLP) technology can be applied as a layer of defense to identify sensitive data and prevent outbound transfers of that data. Identifying the use of proprietary encryption on outbound web traffic is also important to an APT defense.

Strategies and Tactical Defenses

Today, most IT security budgets are largely consumed by antivirus, firewall, and IDS/IPS products, yet the news is filled with stories of targeted attacks—including APTs—that elude these defenses. Traditional security measures do not adequately address today's threats. Without a new security posture, many more attacks using APT techniques will succeed in victimizing their targets.

A sound defense against APT techniques needs to monitor inbound and outbound traffic for content, context, and data, preferably for both email and web communications. More specifically, the defense layer should monitor outbound communications for the detection of data-theft behavior. Some examples of malicious outbound behavior are command-and-control traffic, requests to dynamic DNS hosts, requests to known bad web locations, movement of sensitive files that should never be sent outside the organization (e.g., SAM database), and the use of proprietary encryption.

Networks with firewalls, IDS/IPS, and antivirus defenses focus on inbound threat protection using signatures and individual defense analytics, and mostly ignore outbound communications. Behavioral context analysis and threat scoring from multiple defense analytics is missing, as is outbound traffic analysis for data theft as noted above. Traditional defenses such as firewall and antivirus are necessary because they block known threat vectors; however, they are not sufficient and their limitations against APT techniques and targeted attacks must be recognized and fixed.

Secure web gateways provide an additional defense layer with URL filtering and antivirus scanning, including the ability to analyze SSL traffic. In order to protect your organization in a holistic way, it is recommended that you have a layered defense solution for both inbound protection and outbound data-theft prevention in the event you are compromised.

First, you should have a secure email gateway that has the ability to inspect for malicious web links and attachments to prevent initial infection. Second, you should choose a secure web gateway that has more than just traditional URL filtering and antivirus. To be effective, the solution must have real-time threat analysis to detect zero-day malware and non-binary-based malware (e.g., JavaScripts) to prevent clients from being compromised. Third, the solution should have strong outbound web detection capabilities to detect malicious behavior indicative of a data theft operation in process. To complement this, the gateway should have the ability to see inside encrypted/SSL traffic and attachments so they can be properly inspected for potential sensitive data or malware. Your solution should also have strong DLP capabilities to be able to see when your most valuable data is leaving your organization. While DLP can plug into secure email and web gateways, it often lacks information-sharing about content and context to create a stronger defense.

In short, it is recommended that you employ a unified defense solution that analyzes content, context, and data for both inbound and outbound traffic across web and email egress points to provide the optimal defense for APT techniques and targeted attacks.

Websense Approach to APT Defense

The Websense approach to APT defense is built on three pillars.

- Research is provided by the Websense Security Labs™, which includes a global team of expert analysts leveraging over 16 years of security intelligence and an adaptive feedback network that uses real-time data collecting systems to parse billions of pieces of content daily.
- The Websense Advanced Classification Engine™ (ACE) incorporates several powerful analytics that work in concert to evaluate both content and context for more effective risk detection. ACE is at the heart of every Websense ® TRITON™ product.
- In addition to unified analytics, the TRITON architecture includes unified platforms (enabling Websense products to be deployed on-premise, as a SaaS solution in the cloud, or in effective and economical hybrids) and a unified console for easy and efficient management.

Websense Security Labs

Sophisticated security research is a critical component of any security product positioned to deliver APT defense. APT perpetrators are themselves a sophisticated community of security researchers who constantly evolve their tools and methods to more efficiently achieve their objectives. They constantly work to discover new vulnerabilities and develop malware that takes advantage of those vulnerabilities. Security measures need to evolve as well, with world-class research that stays a step ahead of evolving threats.

Where APTs are concerned, it's not enough for researchers to simply react to individual pieces of malware as they become known. It's not enough to create signatures as malware samples become available. APTs are multi-phased attacks that are unique to each target organization. Although a signature may impede an attacker's path during a single phase of an APT, he will persist and find paths to achieve his objective. This is a fundamental weakness of independently operating signature technologies like antivirus, IPS/IDS, and firewall/UTM gateways with open ports to the web.

To effectively characterize APTs, research teams need to go beyond dissecting attack samples and toolkits. They also need to understand motives, *modus operandi*, and structure. With a holistic approach they can design systems that detect larger patterns of behavior without relying on a single signature.

Advanced Classification Engine

The Websense Advanced Classification Engine (ACE) is a unified content security threat detection engine that powers Websense web security, email security, and data loss prevention (DLP) solutions. It combines a broad range of advanced analytics (Figure 3) including URL Filtering, Reputation, Antivirus, Real-time Security Classification, Real-time Content Classification, PreciseID DLP, and Anti-spam.

ACE analytics are the cumulative result of Websense Security Labs analyses of billions of content elements over more than 16 years. Each analytic on its own provides unique detection capabilities that differentiates it in the market.

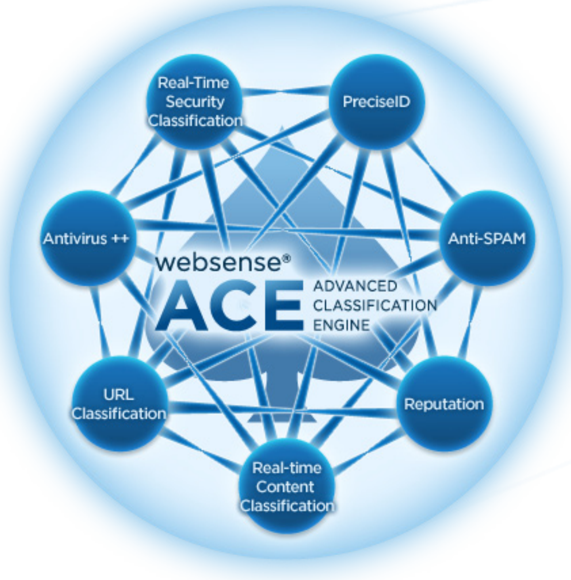
Real-Time Security Classification
PreciseID
Antivirus ++
Anti-SPAM
websense®
ACE
ADVANCED CLASSIFICATION ENGINE
URL Classification
Reputation
Real-time Content Classification

Figure 3: Advanced Classification Engine (ACE) Analytics

But it's not just the *individual* analytics that set ACE apart. The most compelling aspect of ACE is its ability to *combine information* derived from multiple analytics with unique Composite Risk Scoring technology. Each individual analytic assigns a risk score and provides contextual information as input to proprietary risk scoring algorithms. These algorithms then calculate overall risk and detect patterns that indicate the presence of an attack. By combining information from multiple analytics to make more informed decisions, composite risk scoring enables ACE to detect APTs and other complex attacks that evade *independently operating* analytics.

For an example of how Composite Risk Scoring works, consider an obfuscated JavaScript program, downloaded from an unknown URL (no URL database match), at an IP address with a suspicious reputation. In this example, no individual indicator when evaluated in isolation justifies a block.

A JavaScript scanner, URL database, and reputation analytics operating independently would not block this program. However, when the information from all three analytics is combined and evaluated as a whole, it's obvious that the script should be blocked. This is the power of Composite Risk Scoring.

One way to think about ACE is as a *content- and context-aware* attack detection engine.

- **Content awareness** - Each individual content security analytic, from URL scanning to data loss prevention, provides ACE with multi-dimensional "content awareness." Equivalent content awareness would require the use of four or five separate products if purchased from alternative vendors. This baseline content awareness is critical to APT defense given that APTs focus on content-layer (as opposed to network-layer) attacks to evade perimeter firewall controls.
- **Context awareness** - Rather than trying to reach security decisions based on analytics working independently, ACE makes *contextual* decisions based on information derived from multiple analytics. This context awareness is the key to detecting complex APTs designed to evade standalone content analytics like antivirus.

In the Websense APT Defense Tactic section below, we explain how ACE functions during each phase of an APT.

TRITON Architecture

The TRITON architecture is a set of shared security analytics, deployment platforms, and management services that serve as the foundation for Websense products. Websense web, email, and data security products, including cloud, on-premise, and hybrid deployment options, are based on the TRITON architecture. By integrating content security at every level, TRITON solutions lower cost of ownership and protect against complex attacks (like APTs) that go unnoticed by point solutions.

websense® TRITON™

The architecture consists of three main components: unified analytics, unified platforms, and unified management.

Unified analytics: The Advanced Classification Engine (see above) combines multiple web, email, and data security analytics to detect APTs and other complex attacks that evade *independently* operating analytics.

Unified platforms: Websense products can be deployed as software on general-purpose servers, on preconfigured appliances, as a SaaS solution in the cloud, or in powerfully effective and economical hybrids. Regardless of the mix of deployment platforms chosen, the entire system is managed from a single console (see Unified Management below). This approach allows allows organizations to fill gaps in security coverage by using cloud to cost-effectively secure remote offices, mobile users, and mobile devices while preserving on-premise performance and granularity at larger offices.

Unified Management: A single policy management, monitoring, reporting, and forensics console interface extends across web security, email security, data loss prevention, software, appliances, and cloud system components. This consolidated interface reduces management costs by reducing training time, repetitive tasks, and human error. In addition, having a single centralized policy management and reporting server for products and platforms can lower hardware deployment and management costs. Finally, by providing a consolidated view of events across products, the TRITON console enables easier correlation of web, email, and data events that are all tied to a single APT occurring over a period of time.

Websense APT Defense Tactics

In this section, we discuss specific APT defense tactics that can be implemented with the TRITON architecture. These tactics include:

- Identify key data assets and employees
- Prevent infection in Phase 1
- Sever command and control in Phase 2
- Contain data in Phase 3
- Identify infected hosts and data extrusion attempts
- Measure the impact of the event
- Response and adaption of defenses

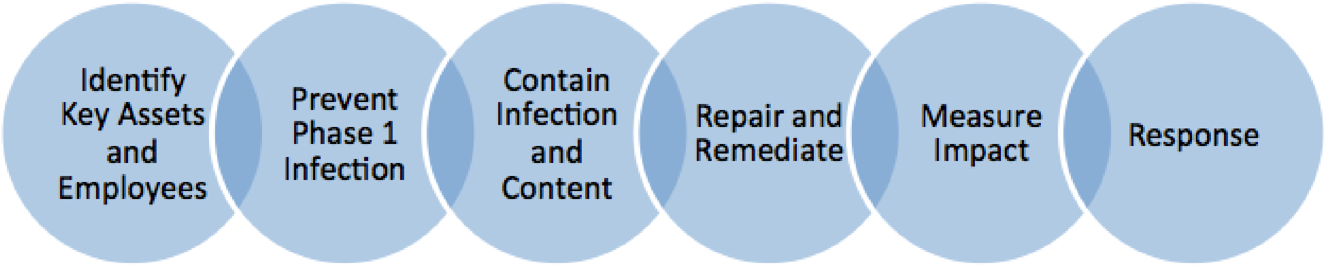
Identify Key Assets and Employees
Prevent Phase 1 Infection
Contain Infection and Content
Repair and Remediate
Measure Impact
Response

Figure 4

Identify Sensitive Data

Before designing a security control system to defend against APTs, you need to identify sensitive data and where that data resides. Websense data loss prevention solutions can aid in this effort by scanning your networks to identify and classify sensitive data. Once sensitive data locations are known, you can reduce risk in several ways.

1. Remove data from insecure or unnecessary locations.
2. Ensure that appropriate access control and attack prevention systems are deployed in all areas where sensitive data is stored.
3. Monitor and prevent data theft at email and web gateways plus end-points.

Identify Key Employees

As previously outlined, attackers often target high-profile employees and members of IT staff who have escalated privileges to data. Once these individuals or groups are identified, TRITON policy management capabilities allow for the creation of custom, user-specific policies that are specifically matched to the risks associated with these high-profile employees.

Attack Prevention in Phase 1

One of the more common techniques applied to gain access to hosts within a target organization is to blend web and email methods. Most blended web/email threats start with an email to a target individual. In most cases, this email is carefully crafted (or spoofed) to appear to originate from another employee or other trusted source like a Facebook friend, and it includes a link to a website with a browser or file-based exploit. Let's look at how ACE analytics work individually and in combination to handle this scenario.

ACE Analytics Working Individually

Initially, individual ACE analytics are applied sequentially and employ increasingly advanced detection methods.

1. First, *Reputation* and *Anti-spam* examine email sending location, user, and content for a range of malicious indicators. It is uncommon for APTs to be blocked at this stage, although general attack heuristics are occasionally effective.

2 *URL Classification* then examines embedded web links to determine if they point to known malicious destinations, or destinations with poor reputations.

3. Next, *Antivirus* analysis of email attachments is performed with multiple third-party antivirus engines. It's rare for third-party antivirus to detect an APT. Once the file passes third-party checks, Websense's own antivirus technology examines the file for advanced malware code, exploit attributes, and other anomalies. Websense antivirus often detects attacks at this point. For example, the attachment may be a PDF with embedded SWF (Flash) actions, a Word document, or a PDF with Active Scripts.

4. Finally, *Real-Time Security Classification* performs deep content analysis for known exploits and more general suspicious indicators.

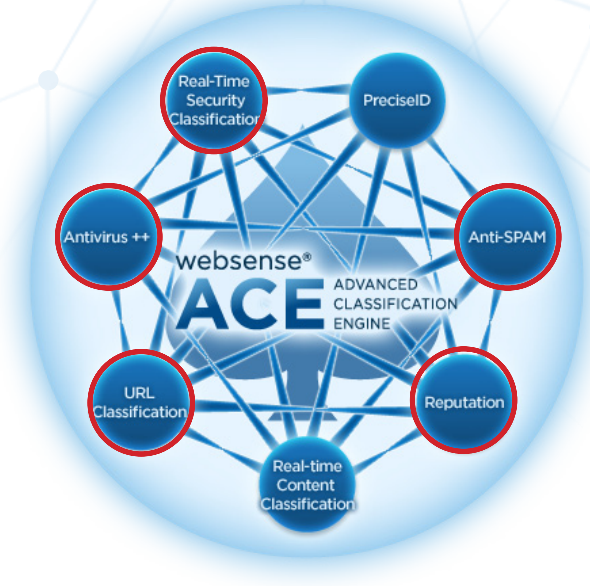
Real-Time Security Classification
PreciseID
Antivirus ++
Anti-SPAM
websense®
ACE
ADVANCED CLASSIFICATION ENGINE
URL Classification
Reputation
Real-time Content Classification

Figure 5

ACE Analytics Working Together

In many cases, APTs are too carefully crafted to trigger an accurate block by any single analytic. In this case, ACE combines all analytics with Composite Risk Scoring algorithms. All analytics provide input to these algorithms with different weights and thresholds applied to each input. By combining analytics in this fashion, ACE is able to detect attacks like APTs, designed to evade individual analytics. Even today's advanced attackers do not design threats with composite analytics in mind. Composite Risk Scoring is therefore the number one method within ACE to prevent Phase 1 APTs.

Severing Command and Control in Phase 2

Once the attacker infects a target host, ACE defense tactics shift to severing command-

and-control connections. ACE applies multiple techniques to this task.

- *Reputation* examines destination address information for malicious indicators such as known malicious IPs, and recently registered domains.
- *URL Classification* identifies known command-and-control web URLs.
- *Real-Time Security Classification* looks for command-and-control indicators within the content of outbound traffic. For example, signatures are used to identify known C&C toolkits, which (as described above) are commonly used in APT attacks.

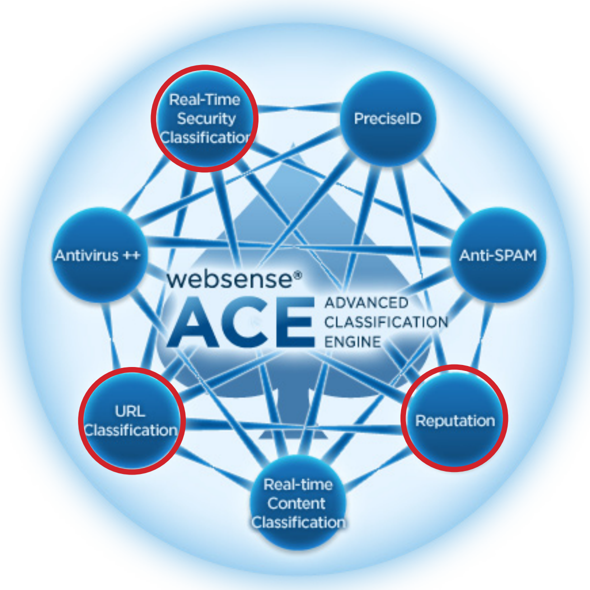
Real-Time Security Classification
PreciseID
Antivirus ++
Anti-SPAM
websense®
ACE
ADVANCED CLASSIFICATION ENGINE
URL Classification
Reputation
Real-time Content Classification

Figure 6: Reputation, URL Classification, Real-Time Security Classification, and Protocol Inspection (not shown) prevent Phase 2 command and control communications

- *Protocol Inspection* validates that outbound traffic on port 80 and 443 traffic is legitimate http and https. APT attacks sometimes attempt to evade content security controls by tunneling non-standard protocols within these commonly open ports. Protocol inspection thwarts this evasion technique, thus forcing attackers to use standard protocols which are subject to deeper content security inspection (such as ACE Real-Time Security Classification).

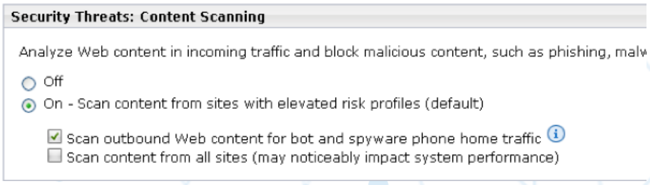

Security Threats: Content Scanning

Analyze Web content in incoming traffic and block malicious content, such as phishing, malw

○ Off

● On - Scan content from sites with elevated risk profiles (default)

☑ Scan outbound Web content for bot and spyware phone home traffic ⓘ

□ Scan content from all sites (may noticeably impact system performance)

Figure 7: Outbound content scanning from the Websense TRITON Console

Containing Data Extraction in Phase 3

Many of the same signals that indicate outbound command and control also indicate outbound data extraction. Therefore, as a first level of defense, the same ACE analytics that identify command and control in Phase 2 also prevent data extraction in phase 3. *Reputation, URL Classification, Real-Time Security Classification,* and *Protocol Inspection* identify a range of extraction indicators such as known malicious URL, poor reputation, suspect content-type (or lack thereof), non-standard encryption, non-standard protocol, and file-type anomalies.

ACE also includes enterprise-class data loss prevention (DLP) to identify and block outbound transfer of the specifically targeted sensitive data. This DLP analytic is referred to as *PreciseID™* (see Figure 8). *PreciseID* includes thousands of predefined classifiers to identify standard data types such as healthcare and financial records. In addition, *PreciseID* includes advanced fingerprinting of data that is unique to each business such as customer database records and business plans.

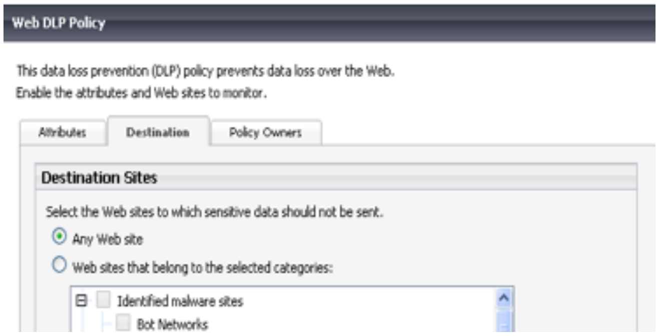

Web DLP Policy

This data loss prevention (DLP) policy prevents data loss over the Web.
Enable the attributes and Web sites to monitor.

Destination Sites

Select the Web sites to which sensitive data should not be sent.

● Any Web site

○ Web sites that belong to the selected categories:

⊟ □ Identified malware sites

□ Bot Networks

Figure 8

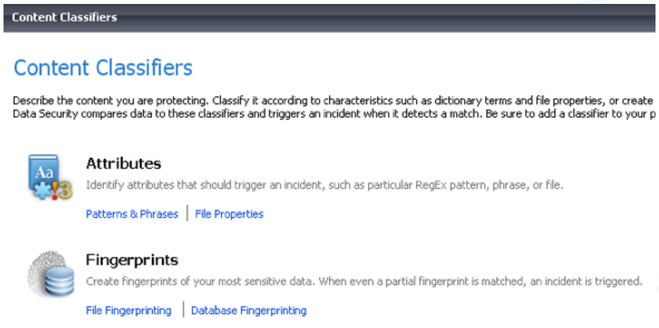

Content Classifiers

Content Classifiers

Describe the content you are protecting. Classify it according to characteristics such as dictionary terms and file properties, or create
Data Security compares data to these classifiers and triggers an incident when it detects a match. Be sure to add a classifier to your p

Attributes
Identify attributes that should trigger an incident, such as particular RegEx pattern, phrase, or file.
Patterns & Phrases | File Properties

Fingerprints
Create fingerprints of your most sensitive data. When even a partial fingerprint is matched, an incident is triggered.
File Fingerprinting | Database Fingerprinting

Figure 9

Threat Monitoring and Reporting

The TRITON architecture provides three management views of infected host activity that may indicate APT presence. These provide a real-time window into what is happening in your organization, with the ability to look at a point in time, and the ability to dig deep into the analysis and flow of content.

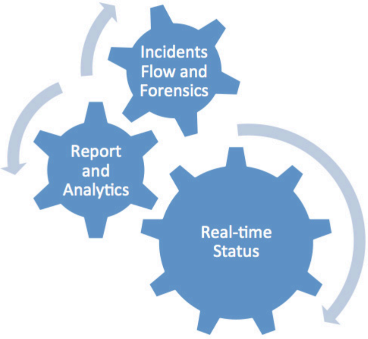

Figure 10

Real-Time Monitoring

The TRITON Real-Time Status Monitor enables monitoring of traffic patterns in real time. Results can be drilled into specifically by user, group, IP address, network, URL, URL category, and policy action (permitted or blocked).

Reporting and Forensics

TRITON reporting capabilities enable tracking of long-term event trends while also allowing security teams to drill interactively into forensic event logs. For example, an investigation into potentially infected hosts can be completed in minutes by first running a report on all traffic to "Botnets," then drilling from that report into hosts that visited botnets, then drilling into outbound content-type, and finally drilling into specific botnet URLs. Even the actual sensitive content carried to malicious destinations is logged and viewable using TRITON forensic tools. If the content is custom-encrypted or proprietary in some way, that will also show up in the event logs. In addition, "outlier" reporting analytics can be used to find infected hosts by automatically identifying anomalous web requests, URL categories, and users.

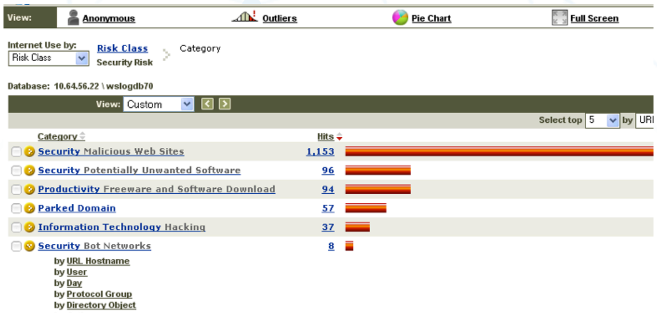

Internet Use by: Risk Class
Risk Class > Category
Security Risk

Database: 10.64.56.22 \ wslogdb70

View: Custom

Select top 5 by URI

Category	Hits
Security Malicious Web Sites	1,153
Security Potentially Unwanted Software	96
Productivity Freeware and Software Download	94
Parked Domain	57
Information Technology Hacking	37
Security Bot Networks	8

by URL Hostname
by User
by Day
by Protocol Group
by Directory Object

Figure 11: Interactive reporting enables easy investigation into infected hosts by drilling into traffic to known Bot Networks.

Response

Once an organization is victimized by an APT, intelligent response is critical. IT needs to work with the executive team and legal counsel to develop a response plan, and may hire outside security experts for assistance. Multiple actions must be taken immediately and in parallel.

- Alert law enforcement authorities
- Analyze forensic logs from infected hosts, Websense, firewalls, IPS, email servers, and other systems. If an attack is ongoing, the smart way to respond may not be to immediately unplug infected computers, but instead to understand how attackers are controlling them before disconnecting.
- Create images of hard drives from infected hosts to ensure that no evidence is lost.
- Reverse-engineer attack binaries to help identify attack methods, communication protocols, and attack servers.

- Identify stolen data. It may be possible to identify data by directly investigating command-and-control servers — if they are still online. This highlights a key point: Taking immediate action without alerting the attacker is crucial to investigating attack servers before they can be taken offline. If a DLP solution was deployed in monitor-only mode (block mode would have likely prevented the theft), then DLP logs should also reveal stolen data.
- If customer data is stolen, develop a plan to notify customers and follow local compliance regulations as required.
- Finally, based on a clear understanding of the attack, develop and implement a plan to remediate existing security controls to mitigate risk of future attacks.

About Websense

Today's productivity tools are increasingly mobile, social, and in the cloud. But so are advanced data-stealing attacks, which antivirus and firewall can't prevent. You can stay a step ahead with Websense® TRITONTM security, which combines best-of-breed web security, email security, and DLP modules (available together or separately) into one powerful solution. With shared analytics, flexible deployment options, and a unified management console, it's the effective and economical solution for today's security challenges.