A Websense Technology White Paper

SECURITY OVERVIEW

# WEBSENSE® ACE (ADVANCED CLASSIFICATION ENGINE)

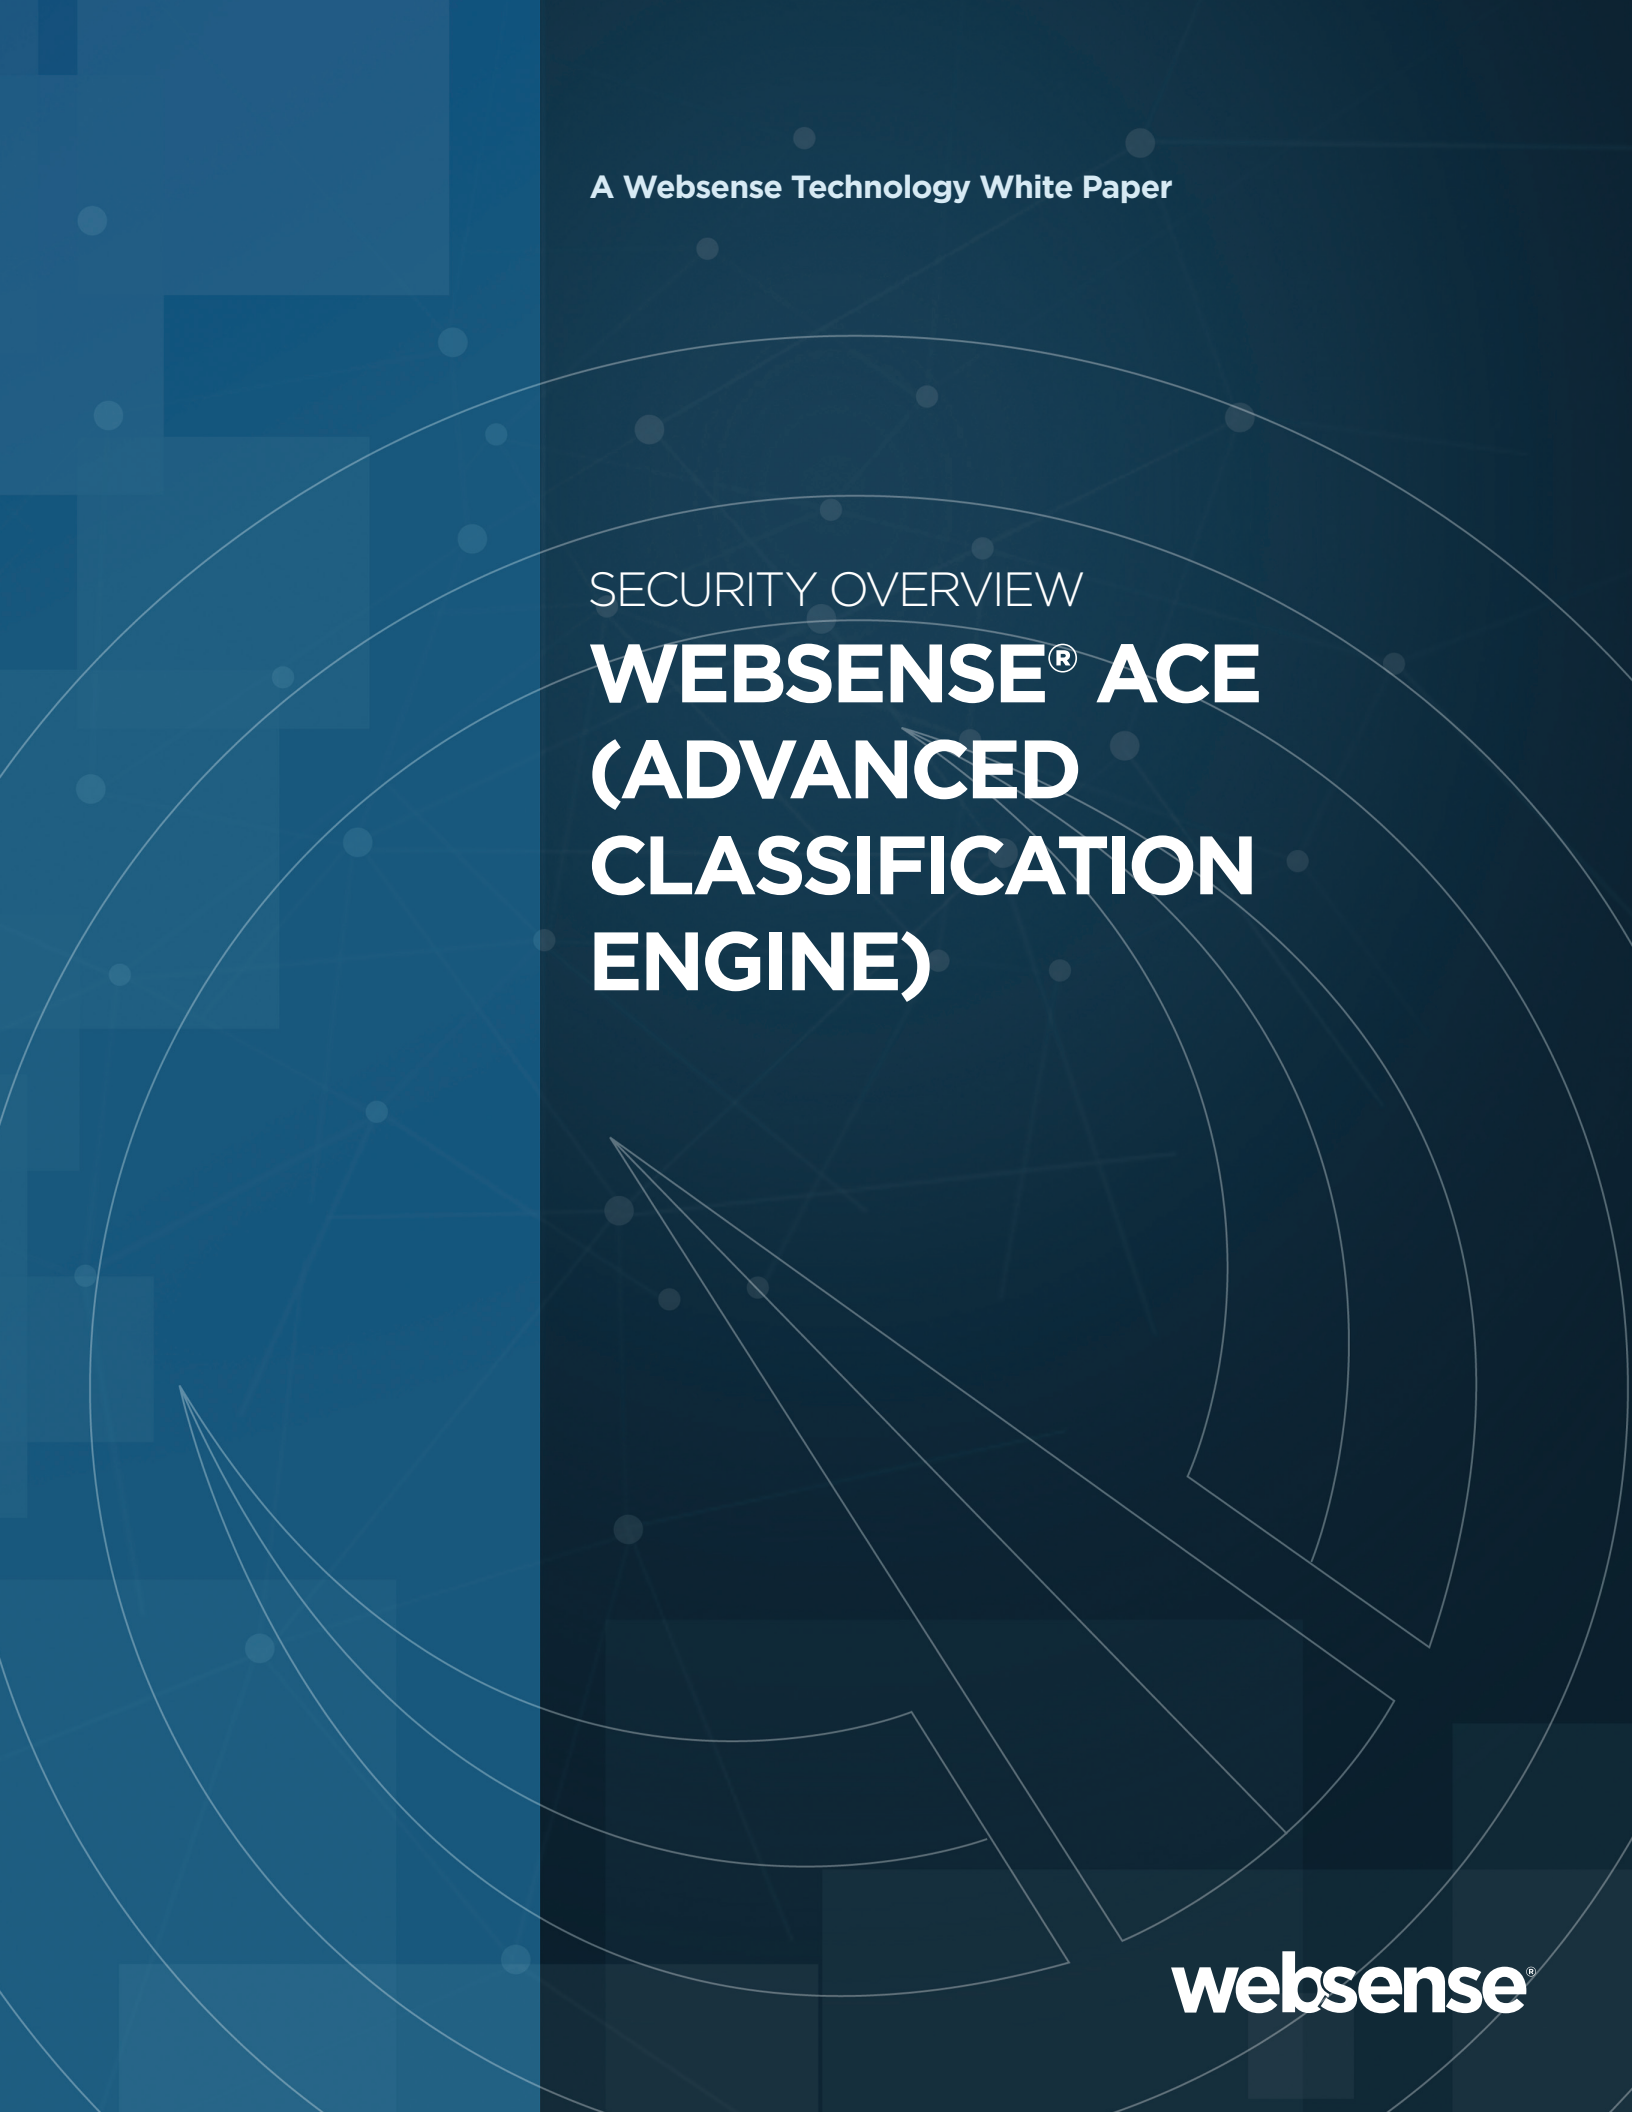

websense®

# CONTENTS

# EXECUTIVE SUMMARY

Those wishing to understand why no one stops more threats than Websense can do no better than begin with a review of the Websense® ACE (Advanced Classification Engine) technology. Our web, email, data, mobile, cloud and forensic security solutions all rely upon ACE to optimize their defensive capabilities. ACE unites seven key defense assessment areas including real-time security classification (e.g., active web page scripts), real-time content classification (e.g., category, language, link analysis), real-time data classification (e.g., password files, criminal encrypted uploads, drip DLP), multiple anti-malware engines, URL classifications, contextual reputation analysis, and anti-spam/spear-phishing. This document outlines these key ACE defense assessment areas and discusses how ACE constantly evolves to keep pace with an expanding, adaptive threat landscape in today's increasingly mobile, in-the-cloud and social world.

Developed and maintained by Websense® Security Labs™, ACE is designed to address today's threats that are more targeted and selective, involve multiple attack stages, and rapidly adapt to circumvent traditional defenses and to exploit your network environment.

- Targeted threats evade many defenses that require attack samples from a variety of regions and industries to respond.
- Multi-staged attacks have become the standard, involving some or all parts of the seven-stage threat model:
  - **Reconnaissance.** Attackers profile potential target user accounts and activities via stolen identities and credentials, calendars, contacts and interests.
  - **Lures.** Attackers create targeted email and web content that preys on human curiosity to trick users into unsafe actions.
  - **Redirection.** A method to lead unknowing users down dynamic cybercriminal paths for potential infection and control.
  - **Exploit Kits.** Attackers detect open vulnerabilities in target systems to deliver malware and advanced threats.
  - **Dropper Files.** Attackers infect systems with malware to establish an inside presence to deliver a scam, demand ransom or collect information.
  - **Call-home.** Infected systems request instructions, replace malware, send home password files, remain cloaked and expand reach within the target.
  - **Data-theft.** Attackers design activities to reap the desired payload and reward. These can be slow-drip data leaks to avoid detection, or heavily obfuscated or encrypted.
- As threats have become more dynamic, ACE applies machine learning and other dynamic capabilities to proactively defeat these attacks through the seven stages of advanced threats.

ACE is a key real-time defense engine used inline within Websense gateways and cloud security services. ACE is also integrated with Websense ThreatSeeker® Network, which unites more than 900 million endpoints that generate up to 5 billion requests of potential threat activity every day. This broad exposure to web activity, including Facebook and other social media, fuels ACE defense assessment areas with threat intelligence and analytics for the predictive composite analysis required for proactive capabilities (e.g., identifying zero-day exploits, new malware, innovative social engineering tactics and other aspects of emerging threats).

Unique to ACE is an embedded enterprise-class DLP engine that enables containment defenses to stop data theft and loss. For example, it can detect and stop password files (e.g., AD/SAM database, text files) from exfiltration, and detect and stop the uploading of suspicious, custom-encrypted files. More advanced uses involve drip (behavioral) DLP that, over a series of events for a defined time period, detects slow data leaks; geo-location destination awareness; and optical character recognition (OCR) of text in image files to detect data theft or loss (the ability of mobile phones to quickly capture images from screens, slides or documents is a growing security concern).

In February 2012, IDC reported that "Signature based tools (anti-virus, firewalls, and intrusion prevention) are only effective against 30-50 percent of current security threats. Moreover, customers expect the effectiveness of signature-based security to continue to decline rapidly." With the declining effectiveness of security solutions previously considered 'core', it is vital to consider what ACE can offer through Websense web, email, data, and mobile security solutions, whether through appliance gateways or cloud security services or a hybrid deployment.

[1] IDC Threat Intelligence Update, 14-Feb-2012

# **WEBSENSE ACE** (ADVANCED CLASSIFICATION ENGINE)

ACE delivers real-time security ratings for protection, productivity and compliance. Protection must address advanced threats, modern malware and data theft. Productivity controls must be accurate and understand social media and other applications. And compliance requires strong outbound content visibility and containment controls including, but not limited to, data loss prevention (DLP) as a defense and risk reduction solution.

ACE provides a number of unique capabilities to support these business requirements including:

- **Predictive security engines** that see several moves ahead.
- **Contextual assessments** to help ensure accuracy.
- **Inline operation** to tackles threats through HTTPS and private social media vectors.
- **Real-time results** for a constantly changing world.
- **More than 10,000 analytics** available to support deep inspections.
- **Composite scoring** to support decision making and the effective integration of a broad set of defense assessment methods.

The more than 10,000 analytics and other methods applied by ACE encompass a broad spectrum of security assessment technologies — from proprietary real-time security, content and data classification processes to traditional security technologies that have declining effectiveness. Understanding the strengths and weaknesses of multiple security techniques enables a system of checks and balances that enable ACE to minimize false positives and improve accurate classification.

ACE can be described as an integrated set of defense assessment capabilities in seven key areas (see Figure 1). A summary of their primary functions and capabilities are listed below; each will be discussed in more detail later.

- **Real-Time Security Classification.** Empowers social web controls, and inspects all web content for malicious or suspicious code such as open or obfuscated scripts, exploit code and iframe tags.
- **Real-Time Content Classification**. Employs advanced machine learning to quickly and accurately classify pages based on content including images, multimedia, and links.


*Figure 1: ACE Defense Assessments*

- **Real-Time Data Classification**. Classifies structured and unstructured data with paring and decoding support to address outbound data theft.
- **Anti-Malware Engines**. Applies multiple anti-malware engines to identify both general and specialized malware.
- **Reputation Analysis**. Considers more than twenty characteristics for detailed assessment and more accurate reputation scoring that encompasses contextual awareness.
- **URL Classification**. Used independently and in conjunction with other defense assessments to apply current classification information for known pages, or to help assess new pages and links.
- **Anti-Spam/Spear-Phishing**. Provides matchless, proactive protection against traditional and emerging threats in email.

Though you may find the depth and effectiveness of the capabilities in each area of defense assessment to be impressive, perhaps the most compelling aspect of ACE is its ability to combine information derived from multiple assessment areas with a unique composite scoring process to improve accuracy and detect things other security solutions simply miss.

## Real-Time Security Classification (RTSC)

RTSC is an advanced, highly dynamic defense assessment focused on emerging threats that are web based, fast moving and use exploit code, browser plug-ins, malicious JavaScript, ActiveX, shell code, exploit kits, cross-site scripts (XSS) and other malicious content. These are the kinds of threats that are often used for drive-by infections and other web page-based attacks.

This area of malware activity is constantly moving and dynamic. They are used by broad, mass-market attacks designed to infect as many people as possible, zero-day threats, and targeted Advanced Persistent Threats (APTs). RTSC was designed for these threat profiles.

Malicious code of this kind is often used in conjunction with obfuscation techniques to further help new, unknown and already hard-to-spot content to evade detection. In order to clearly inspect for potential threats, the RTSC defense assessment incorporates built-in parsing, obfuscation detection and de-obfuscation.

Some threat obfuscation methods mimic techniques used by spam generators to change the spellings of words (e.g., replacing the letter 'o' with a zero) to confuse some content inspection technologies. RTSC faces little challenge here, and works in conjunction with other ACE defense assessments such as anti-spam/spear-phishing, which can help to hone an accurate identification.

Websense social web controls are enabled through RTSC to provision regulatory, productivity or the organizations own risk reduction measures. For example, employee satisfaction may be gained by providing access to social media, while concerns can be mitigated through policies that restrict upload, posting or other activities. As these services continue to expand and evolve, RTSC is charged with identifying these specific activities for customers so that policies are properly applied.

RTSC also provides additional outbound defense services within ACE, including the identification of call-home botnet activity to command and control servers which are then captured in the "Bot Networks" or "Advanced Malware Command and Control" categories. In this way, RTSC supports both proactive defenses as well as network monitoring of activity that can help identify already infected systems that may connect to the network — a valuable capability in a world with an increasingly mobile workforce.

## Real-Time Content Classification (RTCC)

The ultimate challenge for any content analysis application is real-time analysis of dynamically produced web pages with no user-perceptible delay. This is complicated because an accurate assessment must consider the entire page, including the text and images as well as scripts, code, tags, links and HTML artifacts such as fonts and background colors. Each piece needs to be pulled apart, analyzed, classified and correlated before enforcement policies are acted upon. RTCC achieves this in more than 50 natural languages through machine learning and rules-based algorithms along with other special classifiers.

RTCC machine learning is based on a Support Vector Machine (SVM) model and used to analyze the web page at it would be delivered to the user. It extracts and groups all of the prominent features and assigns positive and negative weights. These weights are based on a library of millions of sample pages in more than 50 natural languages, with new samples added daily from Websense ThreatSeeker® Network, which sees up to 5 billion requests every day from more than 900 million unified endpoints.

Many of today's security threats leverage huge numbers of pages that, individually, contain no threats. RTCC examines embedded links to other pages, content with embedded scripts, streaming media, and file downloads to identify potential threats.

The capabilities of RTCC are particularly strong for handling large, highly active web properties full of small, highly variable and possibly unrelated content, such as social networks, where only a single link or embedded link could indicate a threat.

## Real-Time Data Classification (RTDC)

Monitoring outbound traffic is a critical security layer that is often neglected. RTDC provides a number of unique, powerful capabilities to classify outbound content quickly and effectively to mitigate the loss of sensitive information.

Content analysis must see through a variety of data transformation methods often used when criminals attempt to steal data. These include encryption,

[2] Bots, or Botnets, are a network of comprised computers often made available for hire by cybercriminals who use them to send spam, launch Distributed Denial of Service (DDoS), etc. A 'command and control' server acts as a communication hub for this network, and is used to direct the actions of the compromised computers.

formatting changes, screen captures, document format changes and breaking content into smaller components for reassembly later. While many of these are traditional challenges addressed by RTDC, there are a number of unique capabilities:

- **OCR (optical character recognition)** is applied by RTDC to thwart attempts to exfiltrate information (e.g., credit cards and patient information) in the form of images (e.g., screen shots).
- **Files Containing Passwords** is a unique defense to identify, secure and report on events involving files with network passwords including AD/SAM databases.
- **Drip (Behavioral) DLP** is a defense to counter the growing use of data theft techniques such as "low and slow" data theft, where only small amounts of data are taken over time to avoid detection.
- **Custom-Encrypted Uploads** are outbound transmissions that cannot be inspected because they are encrypted in a non-standard way that, in itself, is an indication of suspicious activity. Combined with geo-location destination awareness, the context of the security severity can be assessed.
- **Cascading Fingerprinting** 'fingerprints' overlapping segments of a document to enable identification of even 'partial' versions of sensitive data.

Accuracy is paramount, and the RTDC defense assessment takes into consideration contextual and other aspects of the outbound communication in making a decision. ACE is unique with an embedded enterprise-class DLP engine to provide unmatched data theft defenses.

## Anti-Malware Engines

Cybercriminals still use static code for malware, primarily due to the increased use of 'kits', where they choose their options and simply click a button to turn out new malware with customized content. This makes anti-malware engines very useful in detecting and analyzing this aspect of today's threats.

This defense assessment area applies multiple anti-malware engines to potential malware in real-time. Some are general purpose, publicly available anti-virus (AV) solutions while several others are Websense proprietary security engines designed to identify specific classes of malware or malware designed to exploit a specific document, application, or other file type. Here are a few examples of the anti-malware engines in the ACE arsenal:

- **Customer configurable third-party AV scanner,** either Authentium or Sophos, to provide broad anti-malware detection capabilities to identify mass malware threats.
- **Websense Advanced Detection Scanner**, which employs a combination of packer detection, heuristics and generic detection techniques for certain classes of zero-day malware.
- **Websense Application Recognition**, which is a fast and effective scanner to detect malicious files with very high accuracy.
- **Websense Suspicious PDF Identification Engine (SPIE)**, which is a signature-less scanner designed to secure documents containing content with suspicious characteristics that could lead to the exploitation of a machine, and classifying them as "Potentially Exploited Documents" for control and tracking.

To evade anti-malware engines, malware may use compression techniques known as 'packers' to compress and obfuscate them. The ACE anti-malware engines defense assessment area supports more than one hundred compression formats, which allows it to decompress and extract the content. The identification of the packer itself can serve as an indication of malware, and this is taken into consideration as part of the overall assessment.

The effectiveness of this blend of anti-malware engines and other ACE defense assessment areas, combined with the resources of ThreatSeeker Network, enables Websense to detect hundreds of new malware threats before the leading AV vendors every day. This enables ACE to protect customers even during the window of vulnerability commonly associated with traditional AV — the time from when the malicious content is first seen until the update to the AV signature file is deployed to all systems, which can take from hours to days.

So despite the declining rate of protection provided by exclusively AV solutions, there is still great value in the technology when applied as part of a layered defense, and when AV results are considered in conjunction with other defensive assessment areas. The collective results from the ACE anti-malware engines defense assessments are an important part of the process to identify threats.

[3] See report on "Packer Detection and Generic Unpacking Techniques", http://securitylabs.websense.com/content/Blogs/2927.aspx
[4] http://securitylabs.websense.com/content/AdvancedDetection.aspx gives a list of unknown viruses that are not currently detected by the top AV solutions.

## Reputation Analysis

ACE considers more than twenty different characteristics in its URL/IP reputation analysis by including attributes such as traffic volumes, DNS registration details and the autonomous system number (ASN) , thereby going beyond legacy reputation systems that simply consider the historical malicious use of a URL or IP address. Understanding the deficiencies in legacy security techniques such as reputation analysis enables Websense to expand the information used for reputation scoring and decisions, and use its unique strengths to augment other ACE defense assessments.

The breadth of information used for ACE reputation analysis helps overcome botnet strategies increasingly employed by cybercriminals to mask their use of an Internet location to send spam, host malware, provide botnet 'command and control' and so forth. Victims of a bot attack, unaware of any infection, continue to use the infected system for legitimate purposes, which generate completely legitimate traffic that can confuse less sophisticated reputation analysis systems. Consider also that bot compromised systems are not limited to individual PCs, but often include servers inside the trusted domain of a legitimate organization, which also generate activity that may further contribute to a 'positive' reputation score and mask the anomalous usage taking place. ACE reputation analysis leverages the finer granularity of information to identify the anomalous behavior within the legitimate activity allowing a more accurate reputation score to be calculated.

It is important to note that Websense is also concerned with quickly identifying when a previously compromised site has corrected the situation and is now safe to visit. As an example, consider locations found to host malware, where the unique ACE approach is to monitor such websites and URLs, removing them from the bad reputation database once they have been cleaned, typically within 30 minutes.

## URL Classification

If a URL has been previously classified, this defense assessment will quickly return this rating information. However, URL classification also plays a significant role to assist other ACE defense assessment areas with their contextual evaluations. This is one of many interactions between the various defense assessments so they can optimize their classification recommendations and compensate for any weaknesses in their own capabilities.

Consider that cybercriminals use many methods to bypass security defenses, including the creation of many harmless, legitimate looking pages that are more likely to be presented in search results or otherwise seen by potential victims. And while other ACE defense assessments may consider some aspect of such pages "suspicious", their confidence level in a "suspicious" classification may be low without the additional information the URL classification defense assessment can provide. In this example, the URL classification defense assessment may identify some of the links on the page as known to be malicious in nature, or perhaps that most of the links lead to previously unseen pages — both of which would be suspicious for a finance institution, and contribute to an accurate classification.

## Anti-Spam/Spear-Phishing

This ACE defense assessment area provides a matchless, proactive protection against traditional and emerging email threats. While the majority of today's threats use the web for much of their effort, email still plays a vital role in multi-faceted attacks which have become increasingly targeted. Spam and spear-phishing continue to prey on human curiosity. Spam presents a continued drain on network resources, but spear-phishing attacks are increasingly the result of more detailed profiling of intended victims by cybercriminals. ACE continues to evolve to deal with these shifts in the threat landscape.

Another emerging email tactic is of particular concern because of the simple method used to circumvent most email defenses which examine emails for threats only at the time it is received by the email server. The cybercriminal merely sends an email containing a link to what is currently a legitimate website which they are prepared to compromise, but have not yet done so. Consider the risk if such an email is received over the weekend — the security solution completes its scan, finds no threat and accepts the email. Then, just prior to beginning the work week, the cybercriminal executes their changes to the legitimate website. When the user opens the email and follows the link, they become compromised, often following the advanced threat stages highlighted earlier in this paper.

To prepare for such threats, Websense has developed URL Sandboxing, which is used by Websense email security solutions. When an email containing a web link is first received, ACE may not detect a threat, but it considers the threat potential at which point the web link will be modified within

[5] http://en.wikipedia.org/wiki/Autonomous_System_(Internet)

the email. When a user open his email, the modified link points to cloud-based ACE security defenses to examine the original web link at the point-of-click moment to protect users in real-time. If the link is safe at point-of-click, the user accesses the desired web content. Because it operates within the cloud the URL Sandbox protects users at any time, at any location.

Multiple techniques including sender reputation, content scanning, heuristics and advanced fingerprinting analysis are correlated with input from other ACE defense assessment areas to help ensure that Websense catches more than 99.5 percent of spam (fewer than one in a million emails is misclassified).

## Composite Scoring

Each of the seven ACE defense assessment areas is powerful on its own, and indeed many competitive products depend primarily on technologies in only one or two of these seven areas. But the value of ACE comes from its ability to correlate the results from all seven defense assessment areas with unique, proprietary composite scoring algorithms.

Each ACE defense assessment contributes a risk score and contextual information to the composite scoring algorithms, which then calculate overall risk and consider patterns that may indicate the presence of a threat. By combining information from multiple defense assessment areas, each with specialized capabilities, ACE enjoys superior accuracy. Composite scoring also enables ACE to detect complex attacks such as Advanced Persistent Threats (APTs) that can evade independent or limited sets of assessment technologies.

# WEBSENSE THREATSEEKER NETWORK

Websense ThreatSeeker Network is among the largest security intelligence networks. It unites more than 900 million endpoints to provide greater awareness of web, blog, social network and email content and threats. It operates in conjunction with ACE defense assessment areas to analyze the content of up to 5 billion requests per day.

According to a November 2011 Library of Congress report, the median lifespan of malware distributing domains was only two hours. ThreatSeeker Network responds by driving continuous awareness, particularly in regard to emerging threats, as it collects and identifies content.

## Content Collection

ThreatSeeker Network processes up to 5 billion requests each day to provide content analysis through more than 900 million endpoints. This content takes the form of web pages, documents, executables, streaming media, Emails, mobile applications and other Internet traffic.

Trends identified within this massive request stream are identified and used to direct proactive content collection efforts. To stay ahead of emerging threats, given attackers' preference for compromising legitimate websites, it is important to identify the locations they are likely to target. Popular websites are scheduled for regular assessment, viral sites and content are monitored, geographical hot spots are tracked and hot topics in the media or social networks are used to identify and locate websites hosting related content for frequent, proactive content assessment.

Scanning and collecting content from huge segments of the Internet is no simple matter. As a distributed cloud technology, ThreatSeeker Network scales as needed to efficiently collect and identify massive quantities of content.

To complement the massive content collection efforts of ThreatSeeker Network, many Websense customers share samples of new and potentially malicious content with Websense Security Labs for further analysis. This can be particularly useful for social network sites.

ThreatSeeker Network is further enhanced by the relationship Websense pioneered with Facebook in 2011, which provides awareness of new threats and viral trends, and also provides valuable trend information to help predict potential areas of risk for users. For

[6] Library of Congress, "The Average Lifespan of a Web Page", November 8, 2011

example, in the Websense 2012 Threat Report it was noted that more than 42 percent of Facebook activity is streaming media, which puts Facebook users in a high risk profile to video "lures", often used in the early stages of emerging threats.

The global nature of the Internet, improved connectivity and the explosion of Internet-enabled mobile devices mean malware can be located anywhere and accessed from everywhere. ThreatSeeker Network is designed to be there with you — often before you get there.

## Identification

ThreatSeeker Network applies all seven ACE defense assessment areas as well as a series of out-of-band trending analyses. New analytics in development or adjustments to existing ACE defense assessments may also be performed out-of-band as part of the ongoing evolution of ACE. All in-band and out-of-band assessments are monitored closely by Websense Security Labs to maintain current levels of efficiency and effectiveness while continuously pursuing future enhancements.

Threat intelligence and other services performed by ThreatSeeker Network compliment ACE, often resulting in improved performance of ACE defense assessments.

- **Big Data Analysis**. Proprietary big data analysis tools have evolved over the years at Websense to support up to 5 billion daily content requests and enable automated assessments for key trends and indicators. Security Labs researchers investigate anomalous activity that may lead to the identification of new threat activity, or enhance understanding of emerging threats.
- **Sandboxing**. Sandbox analysis is difficult to perform in real-time, as many malware samples typically require up to two minutes. As an example, the recent Flame/Skywiper malware requires more than 5 minutes before malicious activity begins, which enabled it to evade popular behavior based/activity monitoring solutions. Operating in the cloud, ThreatSeeker Network is able to generate many different types of sandbox environments to simulate samples of various target platforms.
- **Mobile App Profiling**. This sandbox performs traditional malware tests as well as monitoring permission related activities of a mobile app. Specific permission use and abuse can be a strong indicator of malicious intent. Results are used to maintain the Websense "Mobile Malware" and "Unauthorized Mobile Marketplaces" security categories.

# CONCLUSION

Independently, each of the seven ACE defense assessment areas provides unique capabilities and multiple analytics, making them highly effective and accurate in their own right. Some of this may be attributed to use for security assessments within web, email, data, and mobile security solutions providing a unique broad perspective few other security solutions obtain. Used collectively, through Websense composite scoring, ACE enables Websense to stop more threats on inbound and outbound communications over the seven stages of advanced threats.

ACE continues to improve each day as it receives more samples of known safe and unsafe sites with passive lures and active threats from the ThreatSeeker Network. ACE is constantly being primed to secure new devices and new applications as a predictive inline real-time security defense engine. Daily testing against traditional AV defenses confirms ACE is well ahead of the competition in detecting attacks before they break.


[7] Websense 2012 Threat Report
http://www.websense.com/content/websense-2012-threat-report-download.aspx
[8] Websense Security Effectiveness Center: Number of Unknown Viruses Detected:
http://securitylabs.websense.com/content/EffectivenessvsAV.aspx

# APPENDIX A: ACE COMPETITIVE PERSPECTIVE

Comparisons can provide important context when learning about the value of alternative technologies, such as the advanced capabilities offered through ACE. This appendix provides a basis for comparison for several popular, widely available security technologies. Let's begin with a brief review of why the more traditional defense technologies have seen declining effectiveness in recent years.

## Top Reasons Traditional Defenses Fail

There are four reasons that traditional defenses are failing to effectively detect modern malware, advanced threats and prevent data theft.

1. **Primarily based on signatures and/or reputation.** Unfortunately the majority of today's defensive technologies are heavily, if not exclusively, based on the assumption that you will encounter the same threats repeatedly, launched from the same places. Threats today are more dynamic through mechanisms such as self-modifying/encrypting malicious code and frequently changing domain hosts with a median lifespan of two hours. Reputation and signature based solutions simply cannot keep up.

2. **Lack of real-time, inline content analysis**. Most competitive defenses are reliant on sample threat collection by their security labs for analysis by background processes and validation by their researchers to generate updates necessary to detect new threats. Even though updates may be frequent, such updates fail to support a defense until after a lengthy background process.

3. **Forward-facing focus, limited outbound protection.** Emerging attacks often gain a small foothold through targeted email or web activity, after which they regularly call home for instructions or other pieces of malicious code to complete the breach and prepare to steal data. A primarily forward-facing defense only sees part of the attack, is focused on identifying the malicious code, and typically monitors only one entry vector such as the web. This limited perspective is blind to the full attack, cannot assess threat actions in context, lacks valuable data-aware protection, and provides limited if any forensic visibility required by IT to constantly assess and fine-tune defense strategies and policies.

4. **More of the same with new deployment options.** The promise of intrusion detection/prevention (IDS), next-generation firewalls (NGFW), unified threat management (UTM) and other systems have yet to materialize as they suffer from many of the above limitations relying on signatures and reputations. And though they are typically deployed inline, they lack any real-time defense assessments, and the growing use of SSL traffic presents a challenge. Some solutions simply lack the ability to see inside SSL communications, while others offer the support — but at a significant cost in performance which requires additional purchases of capacity for a lower overall return on investment (ROI).

## Comparative Technologies Graph

This graph reflects the effectiveness of comparable detection technologies based on the many elements discussed in this document. Here it becomes clearer why inexpensive solutions based in only one or two defensive assessment capabilities fall far short of addressing todays emerging threats. Those that combine multiple defense assessment capabilities perform measurably better, even more so when inline and using real-time defense assessments for point-of-click protection. And with the seven areas of defensive assessment within ACE, it is more understandable why Websense is able to identify more threats than anyone else.

# APPENDIX B: THREAT MAPPING TO ACE DEFENSE ASSESSMENTS

The power of ACE comes from the combined assessment strength of all seven areas. Nevertheless, certain ACE defenses are more finely tuned to address particular classes of threats. Here is a table of common threat classes and the ACE defense assessment area that takes a lead role in their detection.

| Threat Type | Websense Defense Assessment |
|---|---|
| Advanced Persistent Threat (APT) | RTSC/Anti-Malware Engines |
| Adware | Anti-Malware Engines |
| Anonymity services | RTCC |
| Auto-generated domains | Reputation Analysis |
| Backdoor | Anti-Malware Engines |
| Blackhat SEO | RTSC/URL Classification |
| Botnet Command & Control | RTSC/RTDC/URL Classification |
| Custom packed files | Anti-Malware Engines |
| Data stealers | Anti-Malware Engines |
| Dialers | Anti-Malware Engines |
| Drive-by downloads | RTSC |
| Dynamic DNS | URL Classification |
| Elevated risk profiles | RTSC |
| Embedded code (e.g., inside PDF, SWF files) | RTSC/Anti-Malware Engines |
| Emerging exploits | RTSC |
| Exploit code | RTSC |
| Exploit kits | RTSC |
| Fast flux | Reputation Analysis |
| File infectors | Anti-Malware Engines |
| Hacking tools | RTSC/Anti-Malware Engines/URL Classification |
| Hijacked websites | RTSC/URL Classification |
| Illegal content | RTCC |
| Keyloggers | Anti-Malware Engines |
| Low reputation domains | Reputation Analysis |
| Malicious Active X | RTSC |
| Malicious Applet | RTSC/Anti-Malware Engines |
| Malicious binaries (e.g., Windows Executables) | Anti-Malware Engines |
| Malicious browser plug-in | RTSC/Anti-Malware Engines |
| Malicious flash files | Anti-Malware Engines |
| Malicious insider threat | RTDC |
| Malicious JavaScript | RTSC |
| Malicious Obfuscated code | RTSC |
| Malicious packers | Anti-Malware Engines |
| Malicious PDF | Anti-Malware Engines |
| Malicious RIA | RTSC/Anti-Malware Engines |
| Malicious URL Redirection | RTSC/URL Classification |
| Malicious Visual Basic scripts | RTSC |
| Malicious/suspicious embedded iframes | RTSC |
| Man-in-the-middle | RTSC |
| Packed files | Anti-Malware Engines |
| Password stealers | Anti-Malware Engines |
| Phishing | RTSC/URL Classification |
| Polymorphic binaries | Anti-Malware Engines |
| Porn/gambling/illegal drugs | RTCC |

| Threat Type | Websense Defense Assessment |
|---|---|
| Potentially unwanted software | Anti-Malware Engines |
| Productivity Loss | RTCC |
| Remote Access Trojans (RATs) | Anti-Malware Engines |
| Rogue/Fake AV | RTSC/Anti-Malware Engines |
| Rootkits | Anti-Malware Engines |
| Shell code | RTSC |
| Social engineering | RTSC |
| Social web threats | RTCC/RTSC |
| Spam links | RTSC/URL Classification |
| Spyware | Anti-Malware Engines |
| Targeted attacks | RTSC/Anti-Malware Engines |
| Trojan downloader | Anti-Malware Engines |
| Trojan dropper | Anti-Malware Engines |
| Viruses | Anti-Malware Engines |
| Web spam | RTSC/URL Classification |
| Website Defacements | RTSC/URL Classification |
| Worms | Anti-Malware Engines |
| XSS (cross-site scripts) | RTSC |
| Zero-day exploit code | RTSC/Anti-Malware Engines/URL Classification |