



A Websense® White Paper

The Websense® ThreatSeeker™ Network: Leveraging Websense HoneyGrid Computing

Abstract: Many security teams struggle for visibility and control over the data residing **inside** their organizations. Imagine trying to index and classify the safe use of all content *outside* on the Internet as well. There are terabytes of data inside enterprise networks, petabytes of content on the Internet, and staggering growth trends for each. This data is highly volatile, sometimes changing in seconds as needs dictate and Web 2.0 allows.

Not all of these rapid changes are desirable, or even intentional. In particular, Websense Security Labs found that more than half of all Websites hosting malicious content during the second half of 2007 were legitimate sites that had been recently hacked. With Web 2.0, hacks happen in the space of a few keystrokes.

Websense security researchers overcome these challenges using the Websense HoneyGrid, an adaptive classification and research system within the Websense ThreatSeeker™ network. The HoneyGrid active feedback network uses over 50 million systems to monitor and accurately classify the full range of Internet and enterprise content—not just Web URLs, but all types of Web, email, data, and application content—in real-time. These broadly distributed systems automatically track changing content and trends, collect security research material, and instantly adapt to changes through a perpetual stream of probes and updates.

Without intervention by end-users or system administrators, the Websense Internet HoneyGrid™ and Internal Network HoneyGrid provide crucial context about data types and details on changing Internet content and usage. These systems are used together by Websense to safeguard essential information and Internet use in the enterprise. This technical overview of the Websense HoneyGrid system provides an insider's perspective into each tier of this new technique.

Table of Contents:

Introduction.....	3
Websense HoneyGrid	3
• Honeypots Updated for Web 2.0 Threats	3
• Grid Computing For Classification on an Internet Scale	4
Websense Internet HoneyGrid	4
Websense Internet HoneyGrid Components	5
• Data Collectors	6
• Probes	7
• Analytics	8
• Researchers	8
• Real-Time Updates.....	9
Websense Internal Network HoneyGrid	9
• Internal Data Identifiers and Collectors.....	9
• Analytics and Researchers.....	10
• Real-Time Updates.....	11
Putting It All Together	11
• Browse-By Malware.....	11
• How would the HoneyGrid respond?.....	11
• Inadvertent (Or Possibly Malicious) Outbound Data Loss	12
• How would the HoneyGrid respond?.....	12
Conclusion	12

Introduction

Efficient businesses rely more and more on the Internet business platform—through software-as-a-service and web-By any measure, the Internet continues to grow at an explosive rate. With over 1.3 billion users, Internet use has now penetrated 20% of the World's population.¹ Yahoo stopped publicizing the number of unique web pages in 2005, when the number approached 20 billion pages.² This figure only represents the so-called Surface Web—the small fraction of pages that are visible and can be indexed directly with a static URL address. The Deep Web—dynamic content buried behind hidden databases on Internet portals—is estimated to be between 400 to 550 times larger still.³

The huge extent of the Internet made information protection a challenge when Internet content was mostly static. Now Web 2.0 makes the Internet experience dynamic and far more dangerous, as well as making it a mandatory business platform for enterprise users. Enterprise users increasingly create, access, and transmit essential information—intellectual property, business plans, and customer and employee information—over the Internet. This usage only accelerates as email and Web applications converge, enterprise applications become hosted and browser-based, and Web 2.0 programming techniques propagate into ever more applications. Generally, the more enterprises use the Internet, the more there is a chance for employee error or malicious release of confidential data. Specifically, Web 2.0 enables targeted exploits such as browse-by malware, where payloads are invisibly planted to send emails and steal data.

According to a recent IDC study, two-thirds of organizations are currently using at least one Web 2.0 application. But using these emerging technologies without effective security could be disastrous. In a January 2008 IDC report, Program Vice President Chris Christiansen wrote that “Web 2.0 and Business 2.0 applications and communities will become a major source of identity fraud, privacy violations, and corporate data loss.”⁴

Security teams have no choice but to find a reliable way to allow productive use of the Internet, while safeguarding essential enterprise information from loss or theft.

Enter the Websense HoneyGrid. The Websense HoneyGrid is the research engine inside the Websense ThreatSeeker Network and consists of a classification network in two parts. The Websense *Internet* HoneyGrid sits *outside* of an organization's network boundary on the Internet, providing extensive content coverage and real-time analysis capabilities. The Websense *Internal Network* HoneyGrid sits on the *inside* of an organization's network, providing context on data use within the enterprise. These two systems team to correctly classify and filter content, proactively detect threats, and allow Websense products to mitigate risks to essential information.

Websense HoneyGrid

The Websense HoneyGrid is founded on a combination of two well-known components in computing infrastructure today – *honeypots* and *grid computing*.

Honeypots Updated for Web 2.0 Threats

A honeypot is “a trap set to detect, deflect, or in some manner counteract attempts at unauthorized use of information systems”⁵. Hackers that are lured to penetrate a honeypot expose the nature of their attack methods and, in some cases, the identity of the attacker. As no data of significant value is stored in this traditional honeypot, no sensitive data is at risk.

¹ “Internet Usage Statistics: The Internet Big Picture”, <http://www.Internetworldstats.com/stats.h>, February 2007.

² “Our Blog is Growing Up and So Has Our Index”, <http://www.ysearchblog.com/archives/000172.hl>, August 2005.

³ “The Deep Web: Surfacing Hidden Value”, [http://www.brightplanet.com/images/stories/pdf/deep Webwhitepaper.pdf](http://www.brightplanet.com/images/stories/pdf/deep%20Webwhitepaper.pdf), September 2001.

⁴ Chris Christiansen, Program Vice President, Security Products & Services, IDC, “Security's Troublesome Twins, Crime & Compliance, Ride the Web and Drive 2008 Trends”, January 2008

⁵ “Honeypot (computing)”, http://en.wikipedia.org/wiki/Honeypot_%28computing%29.

Using this concept, all endpoint systems that are connected to the HoneyGrid act as honeypots for observing new attacks and usage trends. There is an important distinction, however. The HoneyGrid threat data collection is actively driven by what all users see and do. Unlike traditional honeypots, it does not wait passively for a direct assault from an attacker. This active model better suits today's Internet, since the vast majority of threats today are not frontal assaults on network infrastructure, but subtle drive-by exploits.

Grid Computing For Classification on an Internet Scale

Grid computing divides processing tasks among a collection of autonomous systems connected by a network to create distributed parallel processing. It provides a cost-effective alternative for complex compute tasks that have historically been tackled only by supercomputers, such as protein folding, financial modeling, and weather modeling. Grids are common in enterprise and service provider networks, and SETI@home and BOINC demonstrate this technique used with distributed networks of independent systems. Leveraging these concepts, the HoneyGrid divides its computational work across millions of desktops and servers worldwide.

The Websense HoneyGrid melds both the honeypot and grid computing concepts. A range of endpoints acting as honeypots provide the volume and diversity of systems necessary to monitor virtually the entire Internet content stream. Distributed data capture and dynamic analysis in small increments enable efficient, real-time adjustments to fine-tune content classification. The HoneyGrid approach gives Websense Security Labs the broadest possible picture of data usage and captures a much larger yield of current threats than traditional honeypot or filtering techniques.

Websense Internet HoneyGrid

The Websense Internet HoneyGrid is a collection of autonomous systems that transmit data to a central location. They characterize the nature of the content that they encounter to improve content classification and security. Each system monitors what it sees on the Web, in applications, or over email; sends feedback on content it encounters; and accepts updates to monitor or refine content classification. As new kinds of content or threats are discovered, an adaptive feedback mechanism queries and evolves the HoneyGrid to classify the new trends within minutes of discovery.

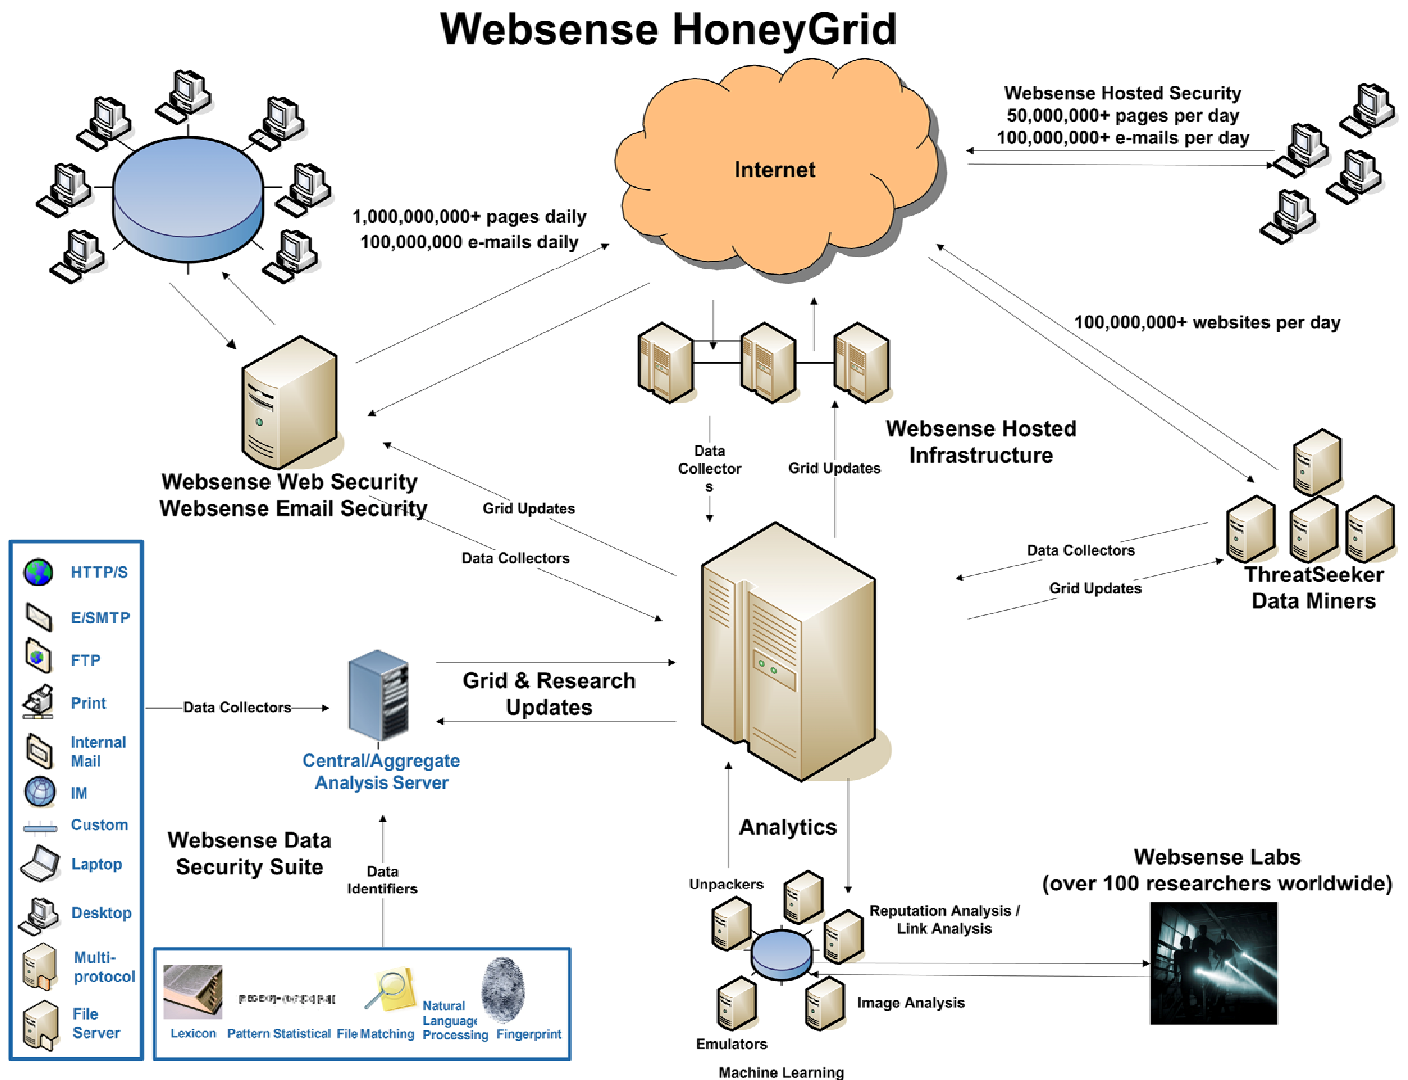


Figure 1: The Websense HoneyGrid dynamically distributes the analysis workload across thousands of systems for accurate classification.

Websense Internet HoneyGrid Components

The HoneyGrid system includes several distributed tiers:

- Data collection endpoints to gather traffic data and usage trends
- Automated analytical tools that examine code and scale overall capabilities for inspection
- Human researchers that perform detailed examinations and drive improvements
- Software probes, the connective tissue that researchers use to drive and refine analysis across the network

Data Collectors

Data Collectors are the individual nodes of the HoneyGrid. Over 50 million systems scan inbound and outbound data, continually transmitting intelligence to Websense Security Labs as they encounter new or flagged content. In total, over one billion Web pages and hundreds of millions of emails are scanned daily. Of these, several terabytes of real-time data are submitted for analysis at Websense Security Labs. While this is a huge amount of data, it represents just a fraction of Internet content: the crucial fraction that has changed in a way that researchers need to know about, or exhibits the flagged attributes that advance research. The data collectors offer a birds-eye view into legitimate content usage trends and malicious activities, all on an enormous scale.

The primary data collectors include:

WebCatcher™

Every Websense Content Filter installation on a customer site includes *WebCatcher*, a configurable component that automatically sends uncategorized URLs and other flagged content to Websense Security Labs for analysis. About 5%, or 50 million, of the billion or so pages viewed through Websense Content Filters ends up being transmitted back throughout each day in WebCatcher submissions. For privacy, static personal information and other identifiers are removed, and other techniques are applied to ensure no personal or confidential information is divulged. This data provides important insight on new sites and changing Internet usage in tens of thousands of organizations.

HoneyClients

A dedicated grid of Websense computers complements WebCatcher data collection by continually downloading content and tens of thousands of executable programs from over 100 million Websites every day. These *HoneyClient* systems mine data from a list of sites whose priority updates dynamically as new or unusual traffic patterns are observed by WebCatcher systems. This interaction ensures examination of the most fresh and potentially risky content, including images, dynamic user-generated content, and niche sites often ignored by reputation-based systems.

Websense Hosted Content Filtering

Through the Websense hosted services, Websense Security Labs also has access to content from hosted customer and non-existent accounts. Customers submit more than 50 million URLs and 100 million emails every day for classification: "does this site contain inappropriate content? Is this attachment malware? Is this email really spam?" The filtering system classifies the content it knows, and transmits anonymized versions of unknown content or flagged content for further analysis. Millions of additional emails, URL links, and executable files that are sent to non-existent accounts are also sent to be analyzed as probable undesirable or malicious content.

Honeypots & Spam Traps

Finally, Websense also hosts and advertises classic honeypots and spam traps that capture more than 10 million unsolicited email and Web-based attacks daily. Details on these unsolicited bulk email, phishing, or exploit campaigns yield highly targeted specifics on the latest attack patterns.

Probes

This diverse set of data collectors allows Websense to scan the entire content stream and collect a great deal of information about the nature of the data. To focus the data collection process and improve the accuracy of data classification, Websense researchers release software probes. These software programs run on data collection endpoints. They identify and retrieve interesting data samples, as well as collections of characteristics that hold specific interest: insight often comes from a cluster of behaviors or attributes rather than a single data point.

As data comes in, researchers refine each probe to eliminate false positives and aggregate those data attributes that contribute to an accurate classification. For example, researchers classifying general content might use a probe to examine the keywords and features of a Web page or email, score them for a particular category, and seek examples that exceed a pre-determined threshold for analysis. For security threats, a probe would look for the presence of a collection of specific characteristics (e.g. vulnerability exploit code, script constructs, intentional format malformations) that, when taken together, are indicative of an attack. In both cases, the content of interest is stripped of any private data and transmitted to Websense for analysis.

The probes that simply gather information are called visibility probes. Those probes that evolve are called classification probes. Eventually, classification probes that are considered accurate are promoted to official detections and populated back out across the ThreatSeeker Network to ensure that risky content is blocked appropriately.

Visibility probes

Visibility probes monitor content to collect statistical trends and help determine which parts of the Internet need further exploration. Large spikes in activity can be very instructive in targeting research efforts, and tracking of content and activity can provide a bellwether of significant usage shifts.

Visibility probes help the HoneyGrid system adapt quickly to change. Probe reports of unusual or suspicious traffic patterns result in new high-priority requests for the HoneyClient data miners. In some cases, a new social phenomenon or current event, such as a sports match, could trigger activity surges, raising awareness of a potentially new classification requirement. In other cases, an Internet attack could drive upload or download activity to a new domain, as with a cross-site scripting attack. This event would trigger a high-priority alert that allows the Websense ThreatSeeker Network to put immediate remediation efforts into action.

Classification probes

Classification probes address one of the greatest difficulties in securing unknown content: the risk of misclassification. Historically, trying to research and classify unknown content required a researcher to download a sample of the content, run some analysis tools, define common characteristics, and deploy detection for these characteristics in the field. The main problems with this approach are, first, that the sample being researched is usually a tiny subset of the content, lacking real-world visibility, and, second, that the data used is no longer real-time and quickly becomes outdated. The absence of real-world real-time data renders it impossible to know or tune the true effectiveness of the heuristic, forcing a risky guess. If deployed aggressively, it risks a body of false positives on an Internet-wide scale, needlessly blocking access to sites and pages. If deployed conservatively, its potency as a classifier is largely negated, and risky content remains active. Most technologies adopt the latter approach: safer, but ineffective.

Websense classification probes leverage the HoneyGrid to get to the root of this accuracy problem. Initially, new classification probes do not attempt to block the content. Their accuracy is monitored in real time as they interact with real-world data on millions of systems. Since the probes are not performing blocking, improvements to each probe can be made without affecting the endpoints. Once a probe has achieved a

high degree of accuracy, it will be upgraded to an official detection and immediately be used to enforce policy on all Websense customer systems.

The HoneyGrid's real-time visibility into a massive body of content mitigates the risk of misclassification, so Websense Security Labs can address new kinds of content aggressively, but safely.

Analytics

A solid research infrastructure must be built and maintained to analyze a huge volume of incoming data, correlate it, and present it for investigation. Probe data delivered to Websense Security Labs is examined first by a set of *analytics*, automated tools that can deeply inspect the nature of the content.

Analytics include:

- Emulators that run executable content in a virtual world and collect behaviors known to be associated with malicious activity.
- Machine learning algorithms that can cluster the content and re-train real-time classifiers, such as support vector machines and Bayesian filters.
- Reputation and link analysis that examine the ecosystem in which the content resides to associate content with its peers and identify potentially risky sites and malware sources.
- Image analyzers that categorize digital images by identifying objects and colors within them. They reveal malware embedded in images and detect inappropriate content, such as pornography.
- Unpackers and deobfuscators that remove wrappers intentionally placed around content to mask its intent.

The key concept for the HoneyGrid is that these automated analytics are able to distill raw data into insight very quickly and reliably. This insight can automatically spawn activity in the HoneyGrid or the ThreatSeeker Network, or it can be channeled to the human research staff at Websense Security Labs to aid deeper examination.

There is a lot more to say about analytics, and more techniques are being automated as this field evolves. To gain a deeper understanding of the many technologies applied by Websense Security Labs, refer to <http://www.websense.com/securitylabs>.

Researchers

Security traffic requires constant, evolving vigilance. In many cases, analytics are sufficient to classify content. However, no amount of automation can achieve the level of quality needed for commercial grade filtering for all content, especially malicious content designed by calculating adversaries. Human researchers are critical to any adaptive feedback mechanism. Only through this combination of analytics and examination by researchers does an accurate picture emerge about the effectiveness of current classification strategies.

Websense researchers review all the collected analytic data and resolve any complex classification issues that arise. Researchers also determine when probes need to be released, tuned, and upgraded. Major changes, such as newly discovered content trends or the surge in Web 2.0 techniques, may require new probes and analysis tools, which are defined and built by Websense researchers.

Real-Time Updates

To close the feedback loop, the results of the analytics and other research from Websense Security Labs are repopulated back to the HoneyGrid and all Websense products, either as probes or official detections.

Probe Refinement

Many probes need to be refined in order to increase their accuracy to the point where Websense researchers are satisfied. Often, probes will be refined and redefined through multiple iterations, each time sent out across the HoneyGrid to collect a fresh set of sample data for evaluation. Through automated distribution, these updates can happen in minutes, ensuring content remains fresh. Eventually, refined probes may be released as “official detections”.

Official Detections

When probes classify content with a high degree of accuracy, they can be promoted to “official detections” and distributed in minutes to the HoneyGrid. Subsequent content that matches an “official detection” will then be blocked or allowed as individual policy on the endpoint system dictates. These updates can take effect simultaneously on millions of systems worldwide to begin the blocking of risky or inappropriate content.

Websense Internal Network HoneyGrid

The tools described so far work together to detect and block inappropriate and malicious external content from entering an enterprise, where this content can wreak havoc on productivity or steal sensitive data. With regulatory requirements increasing, it is also critical to identify proprietary, confidential, or regulated information and prevent its inappropriate transmission outside the enterprise, a process called data loss prevention (DLP).

Data loss prevention provides for the discovery, monitoring, and protection of sensitive content within an enterprise, whether that information is stored (at rest), in use, or being transmitted (in motion) via such channels as the web and email. In order to effectively discover, monitor, and protect data, DLP systems first need to identify it. Data identification requires accurate analysis engines to interpret what is and what is not sensitive information. Unlike Web or email content, internal data is private to each organization, requiring an internal network of analyzers and agents to discern sensitive information.

The Internal HoneyGrid applies many of the general principles used by the external Internet HoneyGrid in order to meet DLP requirements. The Internal HoneyGrid is directly integrated into Websense Data Security. Like the Internet HoneyGrid, the Internal HoneyGrid is a network of data identification and collection agents that works together to identify structured and unstructured data. Unlike the Internet HoneyGrid, the Internal Network HoneyGrid especially targets confidential or regulated data.

Internal Data Identifiers and Collectors

The Internal HoneyGrid uses a similar set of grid-like computing principles to distribute analysis agents across thousands of systems in an organization's network. There are two categories of agents, data identification agents and data collection agents.

Data Identification Agents

Internal data identifiers are data-smart agents that identify the data type (confidential, regulated, etc.) through real-time analysis of data in motion, at rest, or in use. These identifiers can also connect to data repositories such as file storage systems and databases in order to explicitly identify data designated as

sensitive. This identification is accomplished by creating an “information fingerprint,” a mathematical representation of a group of characters, words, sentences, or data fields within a document, message, or database that precisely identifies the designated data together with its extended metadata. Data identifier agents include lexical analysis, pattern matching, statistical analysis, file matching, fingerprinting, and natural language processing.

Data Collection Agents

Internal data collection agents run transparently in the network to provide comprehensive enterprise visibility. They look across protocols, at key business systems—email and print servers, gateways, and file storage, and at local hosts operating in and outside the network perimeter. These agents perform a detailed analysis on all internal and outbound communications to audit and, according to policy, block distribution of any data that is identified as regulatory or confidential.

Data collection agents send suspicious requests to a central server for deep analysis. The agents are customizable to analyze virtually any data type or format. The data collection agents are integrated with the Internet HoneyGrid to discern contextual information in real-time, including the destination, destination category (e.g., Destination: Google Mail; Category: Webmail), and relative safety of that destination.

Analytics and Researchers

The Internal HoneyGrid leverages many of the same analytics for content classification described elsewhere in this whitepaper, with some notable additions for data identification and internal monitoring:

- Similarity measures to identify data that has been altered and might otherwise evade detection
- Lexical analysis, pattern matching, statistical analysis, file matching, optical character recognition (OCR), fingerprinting, and natural language processing, to identify structured and unstructured data, whether confidential, regulated, or proprietary.

The Internal HoneyGrid also benefits from the researchers, tools, and techniques applied to the Internet as a whole.

Some examples:

- Researchers use Internet HoneyGrid systems to test new DLP data rules on millions of websites and emails. These tests help researchers understand the frequency of data occurrence, as well as the effectiveness of the rule in detecting and classifying sensitive data.
- Researchers investigate regulations, standards, data types, and taxonomies on a global scale to create, refine, and deploy new detection and collection mechanisms.
- The Internet HoneyGrid's comprehensive identification of malicious sites complements the data loss prevention of the Internal Network HoneyGrid. It creates a second-tier barrier for data leakage by removing access to malicious Web sites collecting sensitive data.
- The same techniques can be used to parse file formats or deobfuscate/decrypt content both inside and outside the corporate network, supporting a universal approach to content classification.

Real-Time Updates

All internal data identifier and collector agents are connected on the internal network. As new fingerprints, rules, and content are added to the central repository, or real-time updates arrive from Websense Security Labs with new probes, tweaks to existing data, or brand new classifiers, these updates are immediately distributed to all the nodes to close the adaptive feedback loop.

Putting It All Together

We can get a sense for how this comes together using two commonplace scenarios that span both external and internal risks in computing today.

Browse-By Malware

Usually, when a new browser or programming vulnerability is discovered, attackers immediately begin hacking into legitimate Web pages and altering them to host malicious code. When innocent users visit the page, they unwittingly launch attack code on their computers.

1. The exploit payload sends an email with a link to the hacked Website and an enticing subject, such as "Hey, check this out," and
2. Scans the hard drive for any files with credit card numbers, social security numbers, and other account information, sending it to an attacker-controlled Website.

How would the HoneyGrid respond?

Web security view

Based on the announcement of the vulnerability, Websense researchers would immediately define and distribute classification and visibility probes for the new vulnerability and its exploits, updating users of Websense Web Security with real-time content scanning as well as users of Hosted Web Security products. All the pages that are visited by these users would now be scanned for this attack code and blocked. When newly hacked sites are discovered, the HoneyGrid data collectors return feedback about the hacked sites, directing the attention of HoneyClient systems, which begin to mine all the pages on the compromised domain and its links to discover the full extent of the hacked site. New Web page fingerprints are added for all security breaches and updated to customers within minutes.

In the meantime, unprotected users will access the corrupted page and run the exploit code. As the payload launches, a large spike in emails from each user will begin to emerge with the malicious embedded link.

Email security view

Users with Websense Email Security and Hosted Email Security solutions, already updated by the HoneyGrid, will have these systems scan the emails, identify the links as malicious, and block the email. The HoneyGrid returns feedback about the malicious emails to Websense Security Labs, leading to the automatic generation and distribution of email fingerprints for the attack (including a broad set of variations) to all email products in minutes. As attackers alter the emails to bypass traditional filters, Websense Email security products continue to catch the variants and submit both the emails and any newly discovered links back to the HoneyGrid for immediate analysis and Web product updates.

Data security view

The payload will also begin attempting to transmit confidential information. Websense Data Security will identify the type of content being transmitted—in this case, credit card and social security numbers—and block it.

Inadvertent (Or Possibly Malicious) Outbound Data Loss

The second scenario takes us from the externally generated threat to a common internal one. In this instance, an organization's user subverts better business process by attempting a Web post of a sensitive and regulated customer list containing tax identification numbers, credit card information, and customer names.

How would the HoneyGrid respond?**Data security view**

Websense Data Security uses the Internal Network HoneyGrid to automatically detect and identify the data type, regulation(s) it violates, the transmission mechanism of Web posting, and the user performing the Web post. At the same time, the Internal Network HoneyGrid leverages the Internet HoneyGrid by querying it to identify the actual intended post destination site on the Internet and its classification. This destination awareness allows the organization to determine in real-time if the violation is potentially an inadvertent circumvention of the business process for that data, like sending it to a personal Webmail account, or if it is an intentional and potentially malicious transgression, such as posting the data to an auction site as a sample for sale.

The HoneyGrid's real-time detection capability of internal and external content allows an appropriate real-time response appropriate to the situation's context, depending on the result of the detection, the implied intent, and the policy enforced. The inadvertent data leak would be stopped and could trigger notifications to the end-user that the data can only be sent via corporate email, along with a manager notification to alert them of the transgression. The intentional violation would be prevented and could trigger a block, plus a notification to security, human resources, or compliance staff to investigate the end-user and the machine.

Conclusion

Content classification in today's computing environment is an extremely challenging enterprise. The amount of data both inside and outside the network is enormous, and analyzing it requires not only tremendous computational horsepower, but also state-of-the-art analysis tools to make intelligent classification and policy decisions.

The Internet HoneyGrid reaches its Internet-wide visibility in real-time by distributing the data collection and analysis into manageable chunks across millions of systems worldwide. The Internal HoneyGrid leverages data identifiers and collectors within the enterprise for visibility in real-time into key types of sensitive or regulated data and its usage. These research tools and technologies, built from over 15 years of content classification and security experience, are embedded at all tiers. The Websense HoneyGrid is not merely collating data of known attack signatures or packets. Probes identify collections of characteristics that generalize to specific kinds of content such as phishing, proxy avoidance, adult sites, or other forms of irregular content, as well as identifying data types and their associated risk of loss. HoneyGrid systems regularly discover unknown content, and the adaptive network helps ensure risky content gets extra attention. Detailed analysis by researchers also frequently reveals techniques that have never been observed, often resulting in the release of new probes to monitor them as well.

Most heuristic security technology today adopts either a highly conservative route to avoid the risk (yielding a large body of missed detections) or an overly aggressive one (yielding unacceptable false positives), because they lack the visibility into how these heuristics affect users and information assets. By monitoring probes and data collectors in real-time, the HoneyGrid provides a precise understanding of the effect that a new classification strategy will have on millions of systems instantly.

This integrated, multi-tiered, massively distributed approach to attack detection supplies automatic discovery and protection to all Websense and Websense customers in real time.

Updates are transmitted every few minutes as new phenomena or threats are encountered. This ongoing classification process is transparent to Websense users through data identification and collection agents, automatic classifiers, and researchers. The Websense HoneyGrid powers the ThreatSeeker Network, which provides a seamless, dynamic, and adaptable identification and classification infrastructure for Websense web, email, and data security solutions. These solutions provide essential protection for today's real-time inbound and outbound threats to sensitive information.

About Websense, Inc.

Websense, Inc. (NASDAQ: WBSN), a global leader in integrated Web, messaging and data protection technologies, provides Essential Information Protection for more than 42 million employees at more than 50,000 organizations worldwide. Distributed through its global network of channel partners, Websense software and hosted security solutions help organizations block malicious code, prevent the loss of confidential information and enforce Internet use and security policies. For more information, visit www.websense.com and:

- Sign up for security alerts and threat reports <http://www.websense.com/securitylabs/>
- View solution information and supporting educational materials <http://www.websense.com/global/en/ProductsServices/>
- Download whitepapers and case studies or join webcast sessions <http://www.websense.com/global/en/ResourceCenter/>
- Evaluate solutions <http://www.websense.com/global/en/Downloads/>
- Locate and contact a Channel Partner <http://www.websense.com/global/en/Partners/Channel/FindPartner/>