

A Comprehensive, Proactive Approach to Web-Based Threats

TABLE OF CONTENTS

- 1 Executive Summary
- 2 Introduction
- 2 Legitimate Websites, Invisible Threats
- 4 Reactive Filtering Can't Keep Up
- 5 Proactive Protection with Cisco IronPort Web Reputation
- 6 A Comprehensive Approach
- 10 Conclusion

Executive Summary

A sophisticated, dynamic malware economy threatens users—and businesses—with financial loss and brand damage. Cisco® IronPort Web Reputation technology provides advanced protection against these new risks, using unparalleled threat visibility and real-time responsiveness to thwart attacks that evade traditional defenses.

Internet users are under attack. Organized criminals methodically and invisibly exploit vulnerabilities in websites and browsers to infect computers, stealing valuable information (login credentials, credit card numbers, and intellectual property) and turning both corporate and consumer networks into unwilling participants in propagating spam and malware. Simply allowing a user to visit their favorite website, or clicking on a link from their top ten search results, is all it takes for the malware infection process to begin. More and more, malware writers are targeting legitimate, trusted, websites as the starting point for malware distribution. Both BusinessWeek.com and MSNBCsports.com had portions of their websites used for distributing malware. Although no threat is present on these websites today, users became infected simply by visiting trusted sites. Knowing these website are trusted by millions of users makes them easy targets for malware writers.

The sophistication, innovation, and dynamic nature of these attacks often render traditional defenses useless. URL filtering and IP blacklisting are reactive and cannot adequately assess new or previously uncompromised sites in a timely fashion, while signature-based scanning solutions have trouble keeping up with the constant mutation of malware.

A new approach is needed. Protecting users from today's web-based threats requires a layered, holistic, and integrated approach that uses multiple advanced methodologies to assess each threat and type of network traffic. Our best defense as a community of users is to share information about threats in a real-time, automated way so that we can quickly block new threats and shut down the window of opportunity for criminals. Cisco IronPort Web Reputation technology, incorporated into Cisco IronPort S-Series web security appliances, detects and assesses suspicious patterns and websites, as well as vulnerable and compromised elements on individual webpages.

Cisco IronPort Web Reputation technology is based on the extensive knowledge provided by the Cisco Security Intelligence Operations (SIO) framework. Cisco SIO is a cloud-based security service that correlates data received from the Cisco SensorBase® Network—the largest web and email traffic monitoring service in the world—and advanced technologies such as rapid, granular scanning of each object on a requested webpage, rather than just URLs and initial HTML requests. This helps networks significantly reduce their vulnerability—not just to threats from known malicious websites, but also from zero-day and unknown threats from new websites or from sites that are legitimate but invisibly compromised.

Introduction

For most malware creators, recognition for creating a clever piece of malware is no longer the point. With a thriving, maturing malware economy in place, it's more valuable to create malicious code that generates revenues for online criminal networks—for example through click-fraud, massive spam campaigns, or identity and data theft.

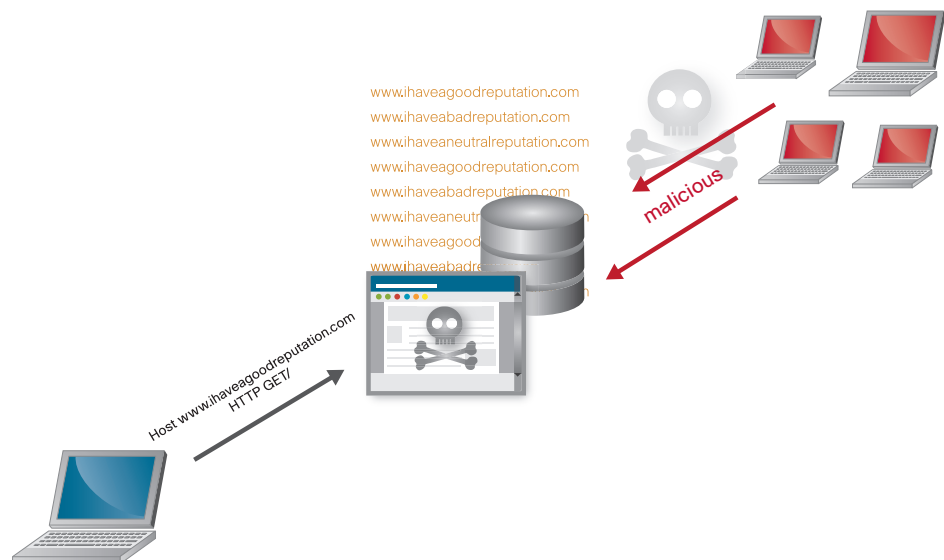
To be successful, the malware must be both easy to distribute to as many victims as possible and difficult to detect. That makes the Internet a very attractive malware delivery mechanism. Originally, malware was delivered directly through email, but the visibility of large attachments and the store-and-forward nature of email made it relatively simple to stop. The near-real-time nature of the web, with threats hidden directly in the content, makes malware exponentially more difficult to stop.

The growing significance of the web as a threat delivery mechanism is shown by the fact that more than 80 percent of spam messages include URLs, which can direct a user to a web server where malware is located. That percentage is even higher for malicious emails, such as phishing campaigns. These URLs are intended to lure readers to websites that will engage them in questionable transactions or download malware onto their computers. Typically, both the spam messages and the malicious websites the messages refer to use a combination of social engineering and software vulnerabilities to compromise users.

Malicious websites, specifically created to distribute malware, are not the only sites compromising users. Hackers are now frequently distributing malware through legitimate websites that have been compromised, taking advantage of security flaws in web applications.

Legitimate Websites, Invisible Threats

Trusted, legitimate websites are the perfect vehicle for malware distribution. Unlike botsites, which are websites specifically designed to host malware, legitimate sites are well-known reputable sites, trusted by users. They often see high user volumes on a daily basis and, most of the time, they are allowed under corporate acceptable use policies (AUP), making them prime targets for online criminals looking to infect as many users as possible.



Criminals compromise legitimate websites to infect unsuspecting users.

Security audit provider White Hat Security estimates that more than 79 percent of the websites hosting malicious code are legitimate websites that have been exploited. Worse, White Hat states that nine out of any ten websites may be vulnerable to attack, with seven out of ten websites being vulnerable to cross-site scripting (XSS) exploits, and one in every five websites being vulnerable to SQL injection attacks.

Popular methods of attacking legitimate websites include:

- **Cross-site scripting (XSS)**

With this exploit, a flaw within web applications lets malicious users, vulnerable websites, or owners of malicious websites send malicious code to the browsers of unsuspecting users. These attacks are frequently executed using HTML image and frame elements (, <frame>, <iframe>) and JavaScript.

- **SQL injections**

This technique takes advantage of a security vulnerability in the database layer of widely used web applications and servers. Hackers take advantage of website administrators not properly sanitizing data transmitted in user input fields (such as forms and user logins) on webpages that use the SQL language and gain control of the website, which they then turn into a malware redirect hub.

- **Exploited iFrames**

Both cross-site scripting and SQL injection exploits use the flexibility of the HTML iFrame as a way to deliver malware directly users. An iFrame, or inline frame HTML tag, allows the embedding of one HTML document—often from another web server inside another HTML document (for example, to include a banner ad on a webpage). iFrames are a very powerful enabler for how the Internet works today, but they can also be used maliciously. Hackers simply incorporate a malicious URL within the iFrame, while using CSS and JavaScript to manipulate the properties of the iFrame (such as its position and size). Recently, hackers have been setting the size attributes of malicious iFrames to zero or, by setting its visibility to hidden in the style sheet, making the iFrame invisible to webpage visitors.

- **Buffer overflows**

With this vulnerability, too much data for an application to handle is sent to its temporary storage, which can enable security breaches.

With a growing number of legitimate websites showing signs of delivering malware, it has become clear that those trusted sites are not producing or even serving the malware. Instead, the site is hacked to include iFrames that refer to malicious IP addresses and URLs, as has happened to Business Week, Wired, CNET, Bank of India, and a host of others. Or, through online ad syndication deals, a legitimate site ends up inadvertently showing compromised or intentionally-malicious advertising banners hosted by third (or fourth, or fifth) parties, as has happened with MySpace and Photobucket.

Traditional defenses are often unable to identify compromised legitimate sites, leaving visitors exposed to infection. Not only are legitimate sites being hacked (often invisibly) but many times, visitors to compromised sites won't be able to detect that they have been compromised and their system has downloaded malware.

In what is known as a “drive-by download,” the malware download happens in the background. In those cases, browser vulnerabilities allow downloads without the user ever clicking on a link, so the user may never know they downloaded malware. When the system is fully patched or a specific browser helper object isn't present, the exploit may more obviously redirect the user to a site that asks them to download an update to legitimate-sounding media-playing software or (ironically) something pretending to be anti-malware software.

Google's 2008 study of drive-by downloads—“All Your iFrames Point to Us”—indicated that, of the billions of URLs the researchers surveyed, more than three million initiated drive-by downloads. Additionally, 1.3 percent of incoming search queries to Google brought up at least one malicious URL in the search results.

Today, downloaded malware casts a very large range of threats. These include adware, keyloggers, botnet participation, data theft and more.

Reactive Filtering Can't Keep Up

Traditional methods of protection are usually not fast, accurate or comprehensive enough to assess and protect users from these new, dynamic web-based threats, which are growing in record numbers.

IP blacklists and URL filtering solutions typically cover only a very small percentage of all URLs and IP addresses—and only the known bad ones. They are also normally binary, offering only “block/malicious” or “allow/safe” options for the URLs and IP addresses they do cover, instead of providing detailed, granular information about any possibly suspicious URL, IP address or object—even those that haven't been known offenders before.

Even with security categories enabled, these URL filtering solutions can't help when a legitimate, normally trustworthy website has been turned into a redirection hub for malware distribution. The website's URL is trusted and not on any blacklist. Consequently, acceptable-use policies designed to protect a network by preventing access to certain sites can't prevent users from getting infected on acceptable websites. Since traditional URL filtering technologies are only concerned with the initial domain request, they don't examine the additional objects needed to load the webpage correctly or their origins, and thus don't observe the malicious redirection. When a webpage has an average of 150 objects, traditional URL filtering technologies simply can't keep up.

This was the case on September 13, 2009 for visitors to NYTimes.com; a trusted source often categorized by URL filtering lists as “news”. A seemingly legitimate advertisement (inserted via a single object on the site) began presenting a pop-up, alerting visitors that a virus had infected their system. Victims were then redirected to a malware site that offered legitimate-looking anti-virus software, which was actually a malicious Trojan.

Meanwhile, web-threat protection systems that depend solely on behavior-based, heuristics solutions tend to rely heavily on the expertise of the administrator. Many of the suspicious or malicious behaviors monitored by these solutions—such as modifying registry settings and accessing system resources—are also exhibited by legitimate software. In that case, it is up to the network's administrator to decide whether such an alert is really dangerous. This is a heavy burden for most networks—particularly when some of the behaviors monitored are highly technical or resemble the actions of legitimate software. Deciding whether to allow or deny largely becomes a guessing game, resulting in a high amount of false positives.

To be useful in today's environment of rapidly evolving, dynamic web-based threats, network protection methods must be proactive rather than reactive, and should offer a layered, integrated approach to security.

They should use multiple means to assess the trustworthiness of websites and the content on individual webpages, so they don't stop users from accessing legitimate content they need, while effectively blocking malicious content. Looking at a limited number of types of threats or depending on historical information is no longer enough. Today's environment requires using comprehensive, deep knowledge based on a large body of data, and the ability to recognize and flag threat patterns in real time, before they can harm users.

This means closely tracking all traffic moving across monitored networks—web, email, IM, downloads and all attempts to access Internet ports. Then building on that information and additional parameters, including rapidly scanning every object on a webpage and its origins—to assess the reputation of websites, objects, URLs, and IP addresses.

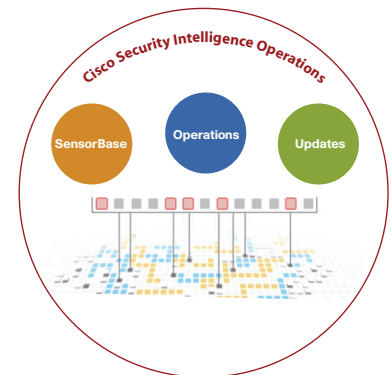
Proactive Protection with Cisco IronPort Web Reputation

Cisco IronPort Web Reputation technology uses pattern-based assessment techniques and fine-grained per-object scanning capabilities to provide its users with timely, accurate threat warnings. Cisco IronPort Web Reputation Filters also take advantage of the Cisco Security Intelligence Operations (SIO) framework, a cloud-based security service that is maintained by the Cisco Threat Operations Center. Cisco SIO correlates data received from Cisco SensorBase—the world's largest threat monitoring network—and provides the Cisco Threat Operations Center with more than 500 GB of threat data daily, with 30 billion individual device queries made directly to SensorBase.

To measure the trustworthiness or reputation of every active web server on the Internet, the Cisco SensorBase Network tracks more than 200 different network-level parameters related to web, IPS, firewall and email traffic.

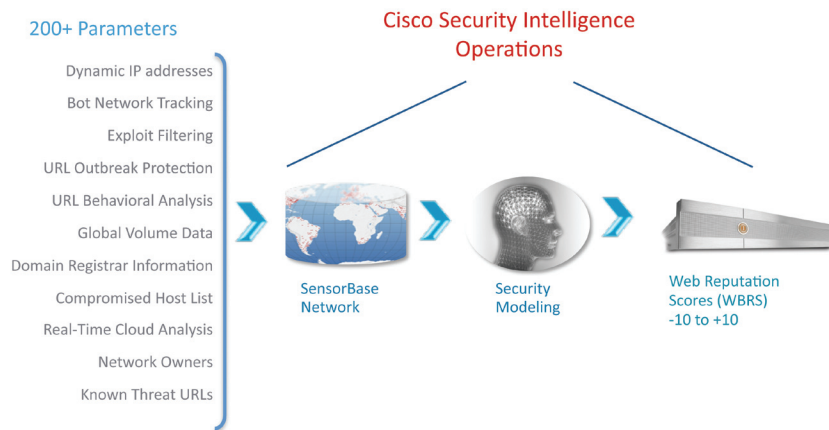
Integrated web, IPS, firewall and email traffic monitoring capabilities enable Cisco security solutions to rapidly analyze the activities, reputation, and potentially malicious affiliations of websites—even if they have never previously been associated with malware. In addition, Cisco IronPort Web Reputation Filters benefit from the data seen by Cisco's global network security services and its unparalleled insight into Internet traffic and trends.

As part of its layered, comprehensive approach to web-based threats, Cisco Security Intelligence Operations also monitors the web for newly created or modified URLs, and receives URL feeds from thoroughly evaluated sources that identify URLs affiliated with malware, spyware, pharming, phishing, and spam. On the request side, Cisco IronPort Web Reputation Filters include Cisco IronPort Virus Outbreak Filtering and Exploit Filtering, both powered by SensorBase. On the response side, Cisco IronPort solutions use multivendor, signature-based anti-malware scanning.



Cisco Security Intelligence Operations provides the highest level of threat correlation—enabling users to collaborate with confidence.

Proactive Protection



All Reputation Filters are Not Created Equal

Cisco IronPort Web Reputation Filters are the world's premier reputation system. Powered by the Cisco Security Intelligence Operations and the SensorBase network, Cisco IronPort Web Reputation Filters have visibility into over 100,000 global networks—including Cisco IPS, 30 percent of the world's email and real-time traffic insights from customer participation.

A Comprehensive Approach

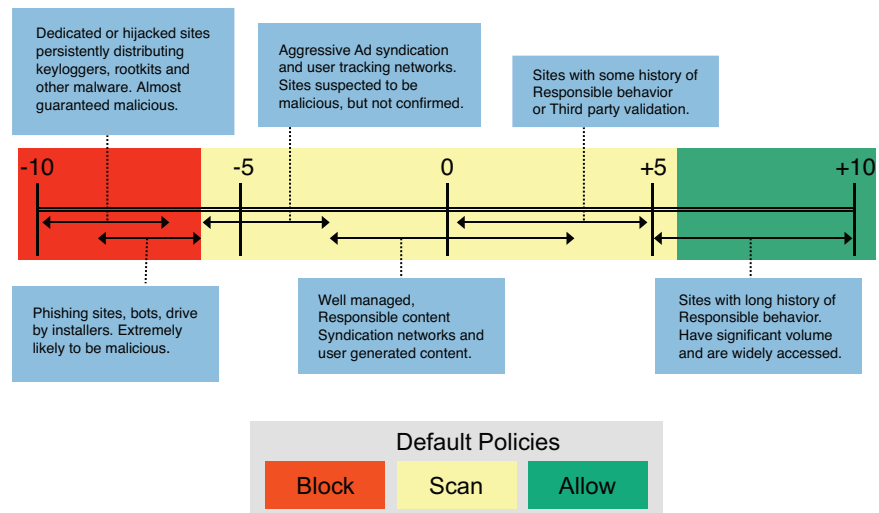
Using Network-Level Parameters to Determine Web Reputation

Analyzing data, even the most difficult to manipulate elements, can reveal much about the trustworthiness of a URL. Data analysis can determine how long a domain has been registered, whether it was registered by machine or manually, who owns it, whether it is associated with an IP address that has previously been associated with a web-based threat, whether the IP address is dynamic or static, what country the website is hosted in, and more.

Cisco IronPort Web Reputation Filters apply data from the Cisco SensorBase Network. SensorBase tracks more than 200 distinct parameters that are excellent indicators of a URL, IP address or web object's reputation. Using sophisticated security modeling and malware detection agents, Cisco IronPort technology evaluates the elements of a webpage and can then create an accurate picture about their trustworthiness. Some of the parameters include:

- Content-based behavioral analysis
- Presence of downloadable code
- Presence of long, obfuscated End-User License Agreements (EULAs)
- Global volume and changes in volume
- Network owner information
- History of a URL
- Age of a URL
- Presence on virus/spam/spyware/phishing/pharming blacklists
- Presence on virus/spam/spyware/phishing/pharming whitelists
- URLs that are typographical errors of popular domains
- Domain registrar information
- IP address information

Web Reputation Score Examples



Sophisticated algorithms analyze and correlate threats with over 200 different web traffic- and network-related parameters to accurately evaluate a web object's malware risk. Using this data, a dynamic score (ranging from +10 to -10) is generated.

Cisco IronPort Web Reputation technology uses the following process to calculate the web reputation of a URL:

1. Through global correlation, Cisco Security Intelligence Operations (SIO) uses pattern-based assessment techniques and security modeling, which associate the above attributes to determine how likely URLs and web objects associated with that particular attribute are affiliated with malware. Depending on that likelihood, a corresponding weight is placed on each of these attributes.
2. Using over 200 network-related attributes, Cisco SIO evaluates the elements of a webpage to determine the overall probability that it contains malware.
3. This aggregate likelihood of containing malware is mapped to a web reputation score between -10 and +10, with -10 being most likely to contain malware and +10 being least likely to contain malware.

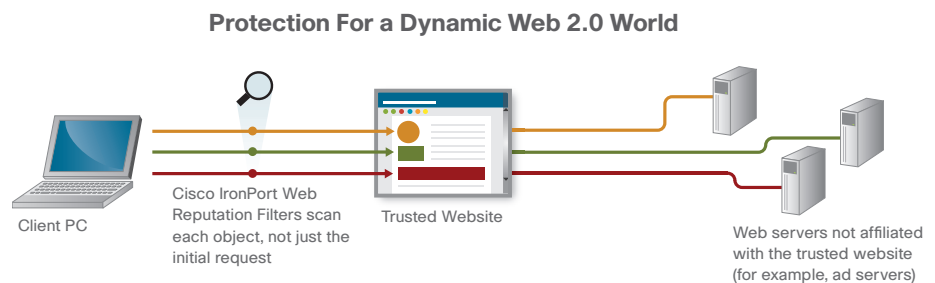
Evaluating Multiple Parameters Offers In-Depth Assessment

Most malware detection applications (including URL filtering solutions) depend on manual assessment and evaluation, and can only offer a binary "good" or "bad" categorization. However, because Cisco IronPort Web Reputation Filters analyze a broad set of data, they can produce a highly-granular reputation score of -10 to +10.

This gives administrators far greater flexibility, enabling them to implement different security policies (like scan further or decrypt) based on web reputation scoring ranges. Instead of trying to create acceptable-use policies based on restrictive binary "good/bad" assessments, administrators can use fine-grained reputation scoring. This allows network users to maintain access to productivity-enhancing web content (without being overly restrictive), while still protecting them from emerging threats.

Scanning All Objects on Webpages Protects from Dynamic Threats

Unlike traditional URL filtering solutions, Cisco IronPort Web Reputation Filters examine every request made by the browser. Instead of just looking at the initial HTML request, they also analyze all subsequent data requests, considering each element on a webpage and its origins—including live data (such as JavaScript, ads, and widgets), which may be fed from different domains. This enables Cisco IronPort Web Reputation Filters to give users a much more precise and accurate assessment and block web content in a far more fine-grained way than URL filtering and IP blacklisting solutions.



Cisco IronPort Web Reputation Filters provide visibility far beyond the initial threat.

Industry-Leading Protection Technologies

Cisco IronPort Web Reputation Filters block up to 70 percent of malware at the connection level, prior to signature scanning. Utilizing a holistic, multi-layer approach—combining comprehensive reputation assessment with in-depth scanning—Cisco delivers a malware catch rate that is 60 percent higher than signature scanners alone.

In addition, the Cisco IronPort Web Reputation system is the industry's only reputation solution to include Botsite Defense, URL Outbreak Detection, and Web 2.0 Exploit Filtering.

Botsite Defense uses heuristics and behavior-based algorithms to accurately identify websites hosted on bot networks. Since many new malware attacks (for example, fake spyware scanners, e-card recruitment spam, and phishing attacks) are orchestrated by botnets, this dedicated detection system that isolates botsite malware helps Cisco IronPort Web Reputation Filters protect users before an attack occurs. When Botsite Defense detects active code, it uses back-end sandboxing to execute the code in a safely walled-off environment and determine if malware is being obfuscated.

The March 18, 2008, iFrame attack on the MSNBC sports website illustrates the effectiveness of Botsite Defense. The iFrame, which linked to a malicious JavaScript file, redirected to an IP address belonging to a web server that had previously hosted malware from Interpace and Russian Business Network. Thanks to Botsite Defense, Cisco IronPort Web Reputation Filters were blocking this botsite five days before the attack occurred.

URL Outbreak Detection leverages Cisco IronPort Virus Outbreak Filters to identify and block URL-propagated malware that has no reputation or signature. This type of malware is typically hosted on a botsite and controlled by a botnet.

The outbreak URLs link directly to malicious files. The user is never taken to a website—instead, with just one click where the user thinks they'll be visiting a website, a malware file automatically installs. The Cisco Threat Operations Center monitors for these outbreaks 24x7x365, and is able to deliver rule sets to Cisco IronPort Web Reputation Filters an average of 13 hours ahead of traditional signature vendors.

Web 2.0 Exploit Filtering zeroes in on the latest network security threat: trusted websites that have been compromised to deliver Trojans or phishing attacks through techniques such as cross-site scripting (XSS), SQL injections, and invisible iFrames.

Using real-time cloud scanning, powered by Cisco SensorBase, Exploit Filters proactively examine content and group compromised websites into three risk levels:

Dangerous. These websites have malicious scripts present but have not been made active by the bot network's command and control server, which is responsible for distributing malware. These sites are also automatically blocked.

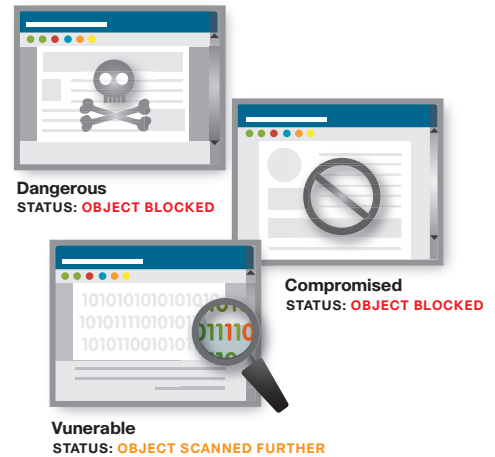
Compromised. These websites are actively serving malware or have malicious scripts injected into them. They are immediately blocked.

Vulnerable. These popular, high-traffic sites show vulnerability to common exploits or have previously been linked to malware distribution. They are put on a "high-risk watch" and actively monitored 24x7x365 by the Cisco Threat Operations Center to ensure continuous protection for all Cisco IronPort Web Reputation customers.

The Cisco IronPort S-Series web security appliance combines traditional URL filtering with advanced Cisco IronPort Web Reputation Filtering and malware scanning on a single platform to proactively address dynamic and sophisticated web-based threats. After web reputation filtering, the Cisco IronPort S-Series uses an advanced scanning engine—the Cisco IronPort Dynamic Vectoring and Streaming (DVS) engine—and multiple signature databases as the second layer of defense. Along with on-box heuristic scanning powered by the McAfee scanning engine, Cisco IronPort S-Series also provides Webroot and McAfee anti-malware signature scanners. Signature-based scanning performs full content inspection to detect and block hard-to-find malware.

The Cisco IronPort Dynamic Vectoring and Streaming (DVS) engine is able to use parallel scanning capabilities to simultaneously scan for malware with both engines, increasing the hit rate on known malware by approximately 35 percent. Although the performance of legacy technology previously made multivendor spyware scanning unrealistic by requiring multiple point solutions, the Cisco IronPort DVS engine was able to remove this performance limitation by employing performance-enhancing techniques such as stream scanning, early exit algorithms, reputation-based verdict caching, and rapid object parsing and scanning. This strategy has allowed Cisco to offer a single web security appliance that integrates multiple malware signature databases that can be used, either in isolation or together, to provide the highest level of threat coverage.

Exploit Filtering



With Exploit Filtering, Cisco offers uncompromised protection against one of the biggest invisible threats on the web: the transparent passing of malware through legitimate websites.

Conclusion

Over the past decade, the threat landscape has changed dramatically. In the past, malware writers were looking for fame, not fortune; unfortunately that is no longer the case. Malware writers are now targeting users with the intent of obtaining their personal information. Items such as credit card numbers, passwords and bank account information are all being exploited for financial gain.

As the motives behind malware have changed, so has the method of the attack. Attack origins were once somewhat predictable and less challenging to stop. But today's malware threats are often times unknown and without precedent. Simply clicking on a search result or viewing a well-known website can be enough to infect a machine with malware.

Malware writers are creating more credible-looking websites themselves, while also increasingly distributing malware by cleverly compromising legitimate websites. These sites are trusted, or had never previously been known as offenders, traditional URL filtering won't suffice to keep users from being compromised. Nor will malware scanning alone offer adequate protection when new malware has not yet been identified with a signature from security vendors.

Allowing users to access valuable productivity resources on the Internet, while protecting them from constantly evolving web-based threats, requires a comprehensive, integrated approach. A solution that looks at more than just URL blacklists or whitelists, but checks every element of the webpage (rather than just the requested URL) from the initial HTML page to all subsequent data requests, considering each element on a webpage and its origins individually—including live data (such as JavaScript, ads and widgets), which may be fed from different domains.

Cisco IronPort offers advanced pattern-based assessment techniques and granular per-object scanning capabilities to provide users with timely, accurate threat warnings. Cisco IronPort Web Reputation Filters also take advantage of the Cisco Security Intelligence Operations (SIO) framework, a cloud-based security service that is maintained by the Cisco Threat Operations Center. Cisco SIO correlates data received from Cisco SensorBase—the world's largest threat monitoring network.

Using this deep, broad collection of threat data on web, IPS, firewall and email traffic generated by this winning combination, Cisco IronPort Web Reputation technology can rapidly detect malicious patterns and attacks. This holistic, layered, integrated approach helps Cisco keep customers safe from the web-based threats of today and tomorrow.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Pulse, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco.Financed (Stylized), Cisco Store, and Flip Gift Card are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Fast Step, Follow Me Browsing, FormShare, GainMaker, GigaDrive, HomeLink, iLynx, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0908R) P/N 434-0206-3 11/09