

CipherTrust IronMail: World-Class Enterprise Spam Filtering

This white paper examines the major components of CipherTrust's anti-spam solution. The reader will gain a better understanding of the spam epidemic, learn specific actions they can take to manage spam and its related threats, and understand the technologies CipherTrust employs to accurately detect and combat spam. Finally, the paper explains the importance of an anti-spam solution as part of a comprehensive e-mail security plan implemented at the network boundary, as opposed to a piecemeal solution installed on the desktop or network.

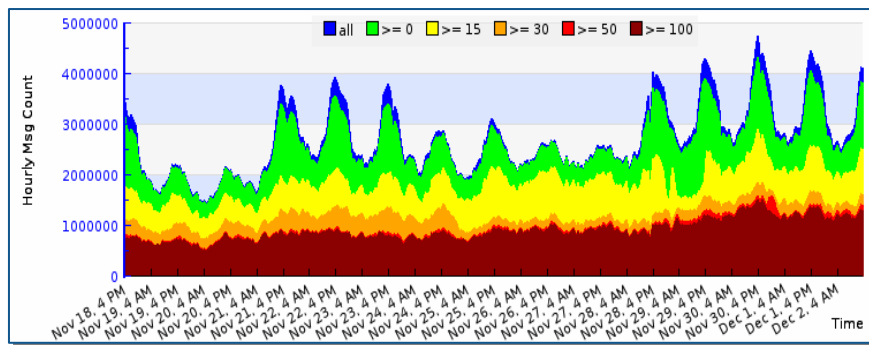
In terms of both volume and costs to enterprises, 2005 has been a record year for e-mail spam. Exponential growth in the amount of unwanted e-mail messages has validated CipherTrust's prediction of the worst spam year on record, and 2006 will undoubtedly build upon the massive volume seen in 2005. Additionally, the costs of spam reaching organizations have risen considerably, to an estimated \$170 per enterprise inbox in 2005¹. As long as sending spam presents a viable income opportunity, the practice will continue to flourish.

Stopping spammers from sending their junk mail is not the ultimate goal, as it is unlikely that the spammers will simply give up their craft; the true goal is to stop their unwanted messages from ever reaching your mail servers.

Once considered only a minor nuisance, spam has emerged as one of the greatest Information Technology (IT) issues for enterprises today. From the minute users log onto their e-mail system, they encounter a deluge of unwanted e-mail that flows into their mailboxes all hours of the day and night. In fact, IDC expects the number of daily spam messages to nearly double over the next few years, from 23 billion in 2004 to 42 billion in 2008.² The exponential growth of spam, combined with increasingly sophisticated security attacks delivered via e-mail, should make deploying an anti-spam solution one of IT's top priorities.

A Growing Commercial Industry

Spam is commonly defined as unsolicited commercial e-mail and is a powerful advertising channel for many products and services. This has resulted in an overwhelming volume of unwanted messages into personal and enterprise e-mail boxes. Spamming is a profitable business, driven by the low cost of sending e-mail compared to other direct marketing techniques.



The graph above displays a sampling of spam volume observed by CipherTrust during a two-week period from late November to early December 2005. The chart indicates multiple 10% compounded increases in volume overnight, highlighting the incredible continuing growth of worldwide spam sending behavior. The areas displaying lower volume indicate weekends, during which time spam volume typically decreases.

Message Volume

Most organizations encounter high volumes of spam on a daily basis. According to CipherTrust research, spam has increased from less than 20 percent of corporate e-mail in 2002 to well over 80 percent in 2005. This presents multiple issues for organizations, as

spam consumes bandwidth, threatens e-mail and network availability, diminishes worker productivity and has the potential to create corporate liability for offensive material that reaches enterprise end users.

A large percentage of the spam sent originates from zombie PCs that have been compromised by hackers via a Trojan horse or other virus. Once infected, these machines are controlled remotely and networked together to send targeted attacks or simply mass quantities of spam to lists of e-mail addresses. As of January 2006, CipherTrust research identifies an average of more than 250,000 new zombie machines every day, an increase of 50% from only a few months earlier.

The recent onset of fraudulent spam variants such as phishing and spoofing pose an even greater risk.

¹ CRN: Spam Costs World Businesses \$50 Billion (<http://www.crn.com/sections/security/security.jhtml?articleId=60403081>)

² IDC: Worldwide Secure Content Management Market Analysis 2005-2009

Phishing

Phishing is a specific type of spam message that solicits personal information from the recipient, such as social security, credit card and bank account numbers. Phishing attacks use social engineering techniques to build trust with the victim before requesting their personal data.

While the target of a phishing attack is the most obvious victim, phishing is unique in that it creates multiple victims per attack. The company that is mimicked in the phishing message also suffers, as their reputation for trustworthiness can be severely damaged, and future legitimate communications from the company may be intentionally disregarded by recipients who no longer trust the source of the message.

Spoofing

Phishing messages and other unwanted mail often incorporate a technique called *spoofing* to trick mail filters into allowing the message through. Spoofing is a deceptive form of spam that hides the domain of the spammer or the spam's origination point. Spammers often hijack the domains of well-known businesses or government entities to enhance the validity to their commercial message or scam.

The combination of spoofing and phishing presents a major threat that can trick most anyone into providing personal information to a spammer.

The Cost of Spam

With the volume and threat of spam on the rise, the business costs of spam have increased dramatically. The sheer volume of spam pouring into enterprise e-mail systems has required enterprises to increase the capacity of their e-mail systems with costly network and infrastructure investments to keep pace. A study from the Radicati Group reported that spam forces enterprises to spend an average of \$49 per e-mail user per year to handle the load, including adding servers, bandwidth and IT staff.

Spam drains employee productivity as workers waste time reading, deleting or even responding to spam e-mails. Additionally, the sexually explicit or otherwise inappropriate nature of many spam messages poses liability issues for enterprises.

The CipherTrust Solution

CipherTrust's IronMail gateway e-mail security appliance protects the e-mail systems of the world's most respected organizations. Key features include:

- The most effective and accurate spam blocking available
- Immediate effectiveness against new spamming techniques and threats, and learning capabilities via

the CipherTrust TrustedSource® behavior-based sender reputation system

- Enterprise readiness
- Integration within an overall e-mail security architecture

Although it takes a person only a moment to process a message and identify it as spam, it is difficult to automate that human process because no single message characteristic consistently identifies spam. In fact, there are hundreds of different message characteristics that may indicate an e-mail is spam, and an effective anti-spam solution must be capable of employing multiple spam detection techniques.

In addition to effectively identifying spam, enterprises must be assured legitimate mail is not blocked in error. Even one false positive (an incorrectly blocked e-mail) can have a significant impact on business. Accurate spam blocking requires a multi-faceted approach that includes examining various message criteria and combining that data with real-time research and intelligence.

TrustedSource™ Reputation System

By aggregating multiple spam detection technologies, IronMail combines the benefits of each individual technique while minimizing the drawbacks. The key to IronMail's effectiveness is the revolutionary TrustedSource behavior-based sender reputation system.

TrustedSource receives and analyzes billions of messages per month from CipherTrust's network of more than 4500 IronMail appliances deployed globally. Like a virtual credit agency, TrustedSource assigns a reputation score and further classifies senders as good, bad or suspicious based on an in-depth analysis by processing more than a dozen of behavior attributes of each sender. TrustedSource is the first and only reputation system to combine traffic data, whitelists, blacklists and network characteristics with the unparalleled strength of global enterprise data. In developing TrustedSource, CipherTrust has succeeded in defining a reputation for *every* sender, not just those that have been encountered in the past. As opposed to other offerings that do not integrate reputation into the spam scoring, TrustedSource data provides the most accurate and effective protection against spam, viruses and other unwanted traffic.

The role of reputation scoring in combating spam cannot be understated. As an example, a CipherTrust customer recently noticed that e-mail from one of their vendors was being blocked by the IronMail appliance protecting them. After a few days, they contacted the vendor to try to gain an understanding of the issue, and found that the vendor had been having a virus outbreak for the past several days, resulting in zombie machines that were spewing spam from within the vendor's network. Because of the real-time TrustedSource data feed, the IronMail appliance was "aware"

of the outbreak and had been blocking connections from the vendor due to the recent dramatic change in their sending behavior.

Social Networks

One of the most effective methods used by TrustedSource to determine reputation scores for senders is the *social network* of the sender. Relationships between all senders are examined by TrustedSource to determine with whom a sender communicates, at what frequency, what volumes of e-mail are generally transferred, and how much mail flow moves in each direction. Based on this social network, TrustedSource is able to instantly detect deviations in behavior, which are typically considered to be suspicious activity.

Real-Time Data Feeds

Information provided by TrustedSource is streamed in real time to CipherTrust products around the globe, ensuring that they are always up to date with the most recent and relevant IP scoring data. As a primary information source for CipherTrust's complete line of messaging security products, TrustedSource plays a large role in ensuring that CipherTrust remains the undisputed leader in messaging security accuracy and performance.

Constant Feedback

The more unwanted messages IronMail units encounter, the better they get at detecting and stopping them. TrustedSource provides constant updates on sender status to CipherTrust; these updates are then sent out to other IronMail units in the field in real time, creating a cycle of feedback that benefits all parties involved (except the spammers) and allows IronMail to achieve the highest level of accuracy in distinguishing the good e-mail from the bad. By tracking sender behavior over time, CipherTrust's database of sender reputation is constantly becoming both more robust and more accurate at the same time.

Connection Control™

With the rapid rise in e-mail's popularity and the corresponding increase in the frequency of spam, the new challenge for today's businesses is to effectively handle the flood of unwanted messages effectively, rejecting it quickly

and efficiently. Ideally, enterprises would handle spam in a manner that not only blocks unwanted e-mail, but actually costs the spammers time, money, and resources in the process.

Connection Control is the first offering to combine network-based traffic shaping and reputation services to reduce e-mail volume from senders who consistently send spam. Reputation systems monitor and categorize e-mail senders by IP address according to their sending behavior. Traffic shaping is the concept of controlling incoming connections to a mail server in order to reduce unwanted traffic.

Taking these concepts to heart, CipherTrust married a powerful reputation system, our TrustedSource™ technology, with traffic shaping. The CipherTrust research team determined that rejecting connections from known spammers at a combination of pre-defined intervals could significantly

reduce the amount of spam that must be inspected at the gateway. As a result, Connection Control convinces spammers that your domain is a poor target for their efforts.

Connection Control allows IronMail to work smarter, by focusing only on the messages that have a reasonable likelihood of being legitimate. And it even offers the opportunity to strike back at the spammers:

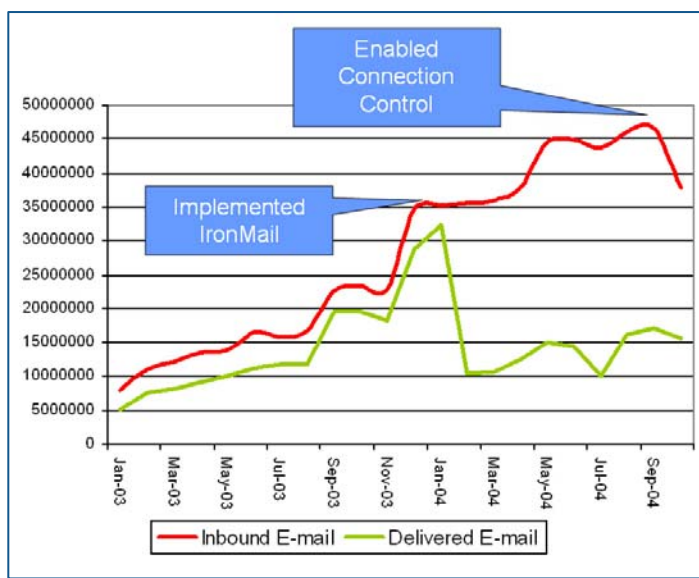
Defensive mode – Connection Control will simply not accept messages from violators for a predetermined time interval. Because of the prevalence of zombie PCs sending unwanted mail, this

technique gives senders the opportunity to repair infected machines and begin sending legitimate mail again once the malicious software has been removed.

Offensive mode – Connection Control turns the tables on spammers by accepting a connection, but slowing the flow of e-mail to a handful of messages per hour. The spammer continues to expend resources but IronMail doesn't. This makes domains protected by IronMail very unprofitable targets for spammers.

The Spam Profiler™

Working alone, each individual spam-blocking technique works with varying degrees of effectiveness and is



Connection Control further increases IronMail's world-class spam-fighting abilities. This graph shows the dramatic effect the technology had at Cox Communications, resulting in a nearly instantaneous drop in inbound messages. Cox reported that nearly 50% of all incoming mail was blocked at the gateway, freeing up valuable bandwidth and mail server storage space.

susceptible to a certain number of false positives. IronMail provides a highly accurate solution by correlating the results of proven first-generation techniques with its industry-leading correlation engine, the Spam Profiler.

The core of IronMail's spam fighting capabilities, the Spam Profiler analyzes, inspects and scores e-mail on over one thousand different message characteristics. Each method is weighted based on historical accuracy rates and analysis by CipherTrust's experienced research team.

Genetic Optimization™

Optimizing the Spam Profiler requires precise calibration and testing thousands of combinations of values associated with various message characteristics. To automate this process, CipherTrust developed Genetic Optimization, an advanced analysis technique that replicates cutting-edge DNA matching models. Genetic Optimization identifies the best possible combination of values for all characteristics examined by the Spam Profiler and automatically tunes the IronMail appliance to the specific needs of each organization, reducing administrator intervention and assuring optimum protection against spam and spam-born threats.

For example, a CipherTrust customer in the retail industry had corporate attorneys who used the website classmates.com to keep in touch with colleagues. Through genetic optimization, the CipherTrust IronMail appliance learned to accept mail from classmates.com that had previously been quarantined.

Techniques for Identifying Spam

IronMail's anti-spam technologies are grouped into six categories:

- Connection analysis
- Lexical analysis
- Protocol analysis
- Authentication protocols
- Traffic pattern analysis
- Auto learning.

Connection Analysis

Connection analysis identifies where a message is going and where it came from. IronMail's transit analysis includes blacklists, whitelists and Domain Name System (DNS) interrogation.

- IronMail queries a variety of blacklists containing domain names and IP addresses from known spammers.
- Whitelists specifically exempt senders and recipients from spam filtering. Particularly useful in

the case of trusted business partners and contacts, whitelists use e-mail addresses, domain names and IP addresses of users who are exempt from filtering, which reduces false positives.

- IronMail's automated whitelisting functionality provides a maintenance-free approach where legitimate e-mail addresses are automatically whitelisted once the recipient has received e-mail from that sender a designated number times.
- DNS interrogation authenticates incoming connections to identify spam from hijacked e-mail servers. IronMail can be configured to deny connections for spammers if reverse DNS lookup fails to authenticate the domain of an incoming connection.

Lexical Analysis

Lexical analysis processes message content to identify spam. IronMail's lexical analysis is based on a combination of URL filtering, content filtering and Bayesian filtering.

- CipherTrust Research has identified thousands of URLs used in spam. These URLs are critical to spammers as they lead to a site where end-users purchase the spammer's product.
- IronMail allows administrators to easily configure words and phrases and manage them in dictionaries, which are regularly updated by CipherTrust Research via CipherTrust Threat Response. IronMail contains a default anti-spam dictionary, as well as dictionaries targeted at confidential, malicious and pornographic content, and administrators may add, delete or edit this list as desired.
- Bayesian filtering creates evolving dictionaries which rate hundreds of thousands of words by their probability of being in a spam message.

Protocol Analysis

Protocol analysis identifies spam by recognizing abuse of or deviation from e-mail protocols. IronMail's protocol analysis is based on forgery detection, header analysis and domain spoofing detection.

IronMail identifies message forgeries by analyzing the connection with a set of heuristic tests to mail headers, including:

- Signatures of spam-generating software
- Violation of e-mail protocols
- Reverse DNS lookups
- Invalid dates

- Forged e-mail addresses.
- Header analysis analyzes custom fields generated by mail servers, particularly bulk mail engines, and gives administrators the ability to monitor and control custom fields used within an organization.
- Domain spoofing allows an enterprise to block messages that originate externally but report the messages originate from an internal domain. The accuracy of the technique is supplemented with the identification of approved relay servers.

Authentication Protocols

Authentication protocols enhance IronMail's effectiveness by identifying legitimate e-mail by applying cutting-edge techniques, including CipherTrust's TrustedSource™ reputation lists and the Sender Policy Framework (SPF) sender identification standard.

IronMail is the first product to offer support for the SPF protocol for legitimate, non-spamming e-mailers to validate their e-mail senders and prevent forgery. SPF, which has now merged with Microsoft's CallerID protocol, protects end users from phishing, spam and viruses. With legitimate e-mailers designating a whitelist of their domains and IP addresses, IronMail's SPF/CallerID analyzes each e-mail on the correlation of the sender's IP address and claimed domain. IronMail recognizes the forged spam when these two essential elements do not match up.

IronMail now supports additional evolving standards for identifying valid e-mails, including the Domain Keys system utilized by Yahoo!

CipherTrust Research & the IronMail Global Network

Over 1800 enterprise customers, including more than 30 percent of the Fortune 100, rely on IronMail's anti-spam and e-mail security solution. With over 7 million enterprise e-mail users, IronMail protects more inboxes than any other enterprise solution on the market.

By pooling the latest e-mail trends and threats from its diverse customer base, CipherTrust has a world-view of enterprise e-mail traffic and behavior that benefits all IronMail users. As a new threat is identified on one network, all other systems are updated in real-time to protect against the threat.

CipherTrust also monitors other sources for information on emerging threats, ranging from open source projects that identify spam to leading industry groups such as the Anti-Spam Research Group of the Internet Research Task Force (IRTF).

This combination of an enterprise network effect and outside sources allows IronMail to identify new spam threats and e-mail variations the instant they hit the Internet.

Enterprise Readiness

A complete anti-spam solution, IronMail is built for enterprise-class networks. Key enterprise-specific features include:

- Built-in redundancy and fail-over for operational persistence
- Low administration
- Adaptive learning via spam traps and "honey pots"
- Continuous updates via CipherTrust Threat Response
- International language GUI, including Chinese, Japanese, Korean, Spanish, German and more
- Fully integrated with LDAP
- Enterprise spam-blocking actions, such as end-user quarantine and administrator quarantine

Zero Administration

IronMail is a zero-administration anti-spam solution. Through continuous automatic tuning and adaptive learning from user-accepted behavior, IronMail eliminates the need for e-mail administrators to become anti-spam experts. When combined with the automatic updates through the CipherTrust Threat Response service, IronMail provides maintenance-free spam protection.

Adaptive Learning

IronMail's adaptive learning systems allow an administrator to create fake e-mail addresses, or "honey pots," that are not associated with an actual employee, for the specific purpose of spam collection. Spam sent to this address is automatically added to the Bayesian spam pool. Honey pots allow IronMail to maintain and improve effectiveness over time without administrator intervention.

Enterprise Spam-Blocking Actions

After identifying spam messages, IronMail provides nine available actions for spam disposal. Available actions include:

- Dropping the message
- Sending a blind copy (for example, to the HR or legal department)
- Rerouting the message
- Attaching a spam prefix to the subject of the message
- Five different forms of quarantine

IronMail can label spam by modifying the message subject line, allowing end users to create their own personal policies

for spam in their e-mail clients with rules that delete the message or send spam to specific folders.

IronMail's end-user quarantine allows designated users to receive regular updates of quarantined messages. If a message has been incorrectly quarantined, the end user can release the message and IronMail will utilize the localized Bayesian systems to learn the marked message was legitimate.

For administrators, end-user quarantine eliminates the need to review quarantine queues for possible lost mail and to construct rules and policies in response to user feedback. They can select which users or groups of users need access to the user quarantine and which do not. All of this is accomplished without allowing suspect e-mail to pass through the gateway until it has been released by the user. Administrators can access the quarantine queue through a secure browser-based interface.

IronMail enables end-user spam reporting to proactively protect against spam while minimizing or eliminating administrative overhead. Many ISPs and enterprises provide a designated e-mail address for users to report spam messages; mail administrators are typically responsible for examining the messages and developing a rule to respond. This overhead usually results in ineffective spam management or hiring a dedicated staff.

IronMail's automated spam-abuse management technology eliminates this burden. IronMail accepts messages users forward to the designated "send spam" e-mail address and automatically process the messages to extract key identification characteristics. IronMail provides feedback to the Bayesian engine as it creates and enforces a new policy rule that blocks, quarantines or labels any future messages from this spammer based on the e-mail address, subject or IP address.

Comprehensive E-Mail Security

Anti-spam is only one component of complete enterprise e-mail security. In keeping with that concept, IronMail not only identifies and blocks spam, including phishing and spoofing messages, but also provides a complete solution to protect enterprise e-mail systems from other threats including denial-of-service attacks, intrusions and Web mail attacks. By providing this protection at the e-mail gateway, IronMail ensures that all traffic is thoroughly screened prior to entering or leaving the enterprise, providing protection that is unavailable using traditional desktop-based software solutions.

"There are very few scenarios where we would recommend desktop filtering unless it was a 'belt

and braces' approach with multiple layers at the gateway, inside the network, and on the desktop."

— *Forrester Research*

Web Mail Attacks

Web mail provides another common point of intrusion to e-mail systems. Many organizations today allow their mobile workers to access corporate e-mail through a Web browser by using Outlook Web Access (OWA) or iNotes. Web mail requires a Web server such as Microsoft's Internet Information Services (IIS), which is subject to numerous vulnerabilities, blended threats, viruses and worms.

IronMail provides a secure platform to protect Web mail. As a hardened e-mail appliance, IronMail acts as an application-specific firewall and allows only valid and safe connections to mail servers. IronMail is designed to block all manner of e-mail attacks, including buffer overflow, denial of service and exploits such as malformed MIME headers directed at internal servers. By acting as a secure gateway, IronMail protects the entire e-mail system from all e-mail threats.

Summary

Controlling spam is a critical requirement for enterprises today. Working in tandem with the TrustedSource reputation system, IronMail's Spam Profiler correlation engine ensures highly effective and accurate spam blocking. IronMail provides a comprehensive solution for e-mail security, encompassing hackers, intruders, policy and compliance, and viruses and worms, as well as spam. CipherTrust's research team and Threat Response Updates ensure that IronMail continues to protect against spam and other threats to ensure continuous effectiveness over time. IronMail's robust administration and Genetic Optimization capabilities ensure administrators win the battle against spam.