

# **Spam, Spammers, and Spam Control**

**A White Paper by Ferris Research**

**March 2009. Report #810**

Ferris Research, Inc.  
One San Antonio Place  
San Francisco, Calif. 94133, USA  
Phone: +1 (650) 452-6215  
Fax: +1 (408) 228-8067  
[www.ferris.com](http://www.ferris.com)

# Table of Contents

<b>Spam, Spammers, and Spam Control.....</b>	<b>3</b>
Defining Spam.....	3
Spammer Tactics .....	3
Sending Mechanisms.....	4
Spammer Tricks.....	4
Techniques for Identifying Spam .....	5
Connection Analysis.....	5
Behavioral Analysis.....	6
Content Scanning.....	6
Controlling Spam: How and Where .....	7
The Key Role of Reputation Services .....	7
Conclusion: An Arms Race .....	8
<b>Trend Micro Interview.....</b>	<b>9</b>

# Spam, Spammers, and Spam Control

Just about anybody with an email address knows about spam and spam filters—or at least they think they do. However, it is a complex subject, with far more to it than meets the eye.

This Ferris Research white paper looks at spam: what it is, how it is sent, and how it is filtered. It unravels the complexities and gives an accessible, factual overview into the scourge of spam.

## Defining Spam

Most spam falls into the following categories, listed in roughly descending order of received volume:

- Fake pharmaceuticals
- Fake fashion items (for example, watches)
- Pornography and prostitution
- Stock kiting—that is, spammers driving up the price of stocks by inciting victims to buy them (also known as “pump and dump”)
- Phishing and other fraud, such as “Nigerian 419” and “Spanish Prisoner”
- Trojan horses attempting to infect your PC with malware
- Misdirected nondelivery reports and autoreplies sent by badly configured mail servers replying to forged email (“backscatter”)
- Spam from other types of senders, such as ignorant marketers, rogue affiliates, and misguided politicians or charities

The word “spam” is sometimes used narrowly to describe just commercial advertisements for products and services. This definition is dubbed UCE: *unsolicited commercial email*.

However, spam is more frequently used to describe *any* unsolicited email that is sent in bulk. Therefore, we prefer the definition UBE: *unsolicited bulk email*. This broader definition includes viruses by email, phishing, and other email-borne threats.

## Spammer Tactics

Let’s look at some of the tactics that spammers use to send their email and how they evade detection by spam filters. As this is a brief report, we do not provide a complete list but rather focus on the major issues that spam filter technologists need to deal with at present.

## ***Sending Mechanisms***

A typical spam campaign sends millions of email messages. The economics of spam and ISPs' terms of service usually make it impractical for a spammer to send such a large volume using a conventional mail server.

Here are the main ways that spammers get their messages out, listed in order of pervasiveness:

- *Botnets*. Networks of “zombie” malware-infected PCs send email on behalf of the spammer, without the knowledge of their legitimate owners. Botnets are controlled by a botmaster, who sells a spamming service to those who wish to send spam.
- *Free email services*. Public Webmail (for example, Yahoo! Mail) is misused to send spam.
- *Other free services*. The “send to a friend” functionality of social Web sites (for example, Picasa Web) can also be misused.
- *Open proxies*. Compromised or misconfigured servers can be used to redirect spam. Known in spammer slang as “peas,” open proxies are also sold as a service to spammers in a similar way as botnets.
- *Stolen netblocks*. Spammers set up in business as an ISP, typically by taking over portions of Internet address space—often through illegal means.

## ***Spammer Tricks***

In no particular order, here are the main ways that spammers try to fool spam filters:

- *HTML tricks*. A spam message is manipulated to appear in one way to an email user but in another way to a spam filter—for example, mixing the spam body text with fragments of invisible but legitimate-seeming text. The invisible fragments are hidden using various HTML tricks. Unsophisticated filters are confused by the invisible text and fail to identify such messages as spam.
- *Bayesian poisoning*. The email includes large blocks of legitimate-looking text in an effort to fool statistical content filters.
- *Content morphing*. The sender varies the spam text and headers to fool simple filters that look for known text fragments.
- *Images and other attachments*. Instead of sending spam text in the message, the spammers send graphical files that display the text; they vary the images to try and fool filters that look for known images sent in bulk.
- *Forcing secondary MX*. Many receiving domains specify a backup mail server in case the main server is unavailable. Spammers often expect these secondary “Mail eXchangers” to have poorer spam filtering, so they force their mail to be delivered to the backup.

- *Countering IP reputation.* Zombies normally send spam directly to the victim's mail server, but IP reputation can be harder to implement if they send it via a legitimate mail server. Some bots now submit email via the mail server of the user who owns the zombie PC, as if the user had sent it.
- *Hiding the call to action.* For example, the recipient is asked to manually type a URL rather than clicking a link. Or spammers use short-lived domains for the click target.

## Techniques for Identifying Spam

We discuss three broad categories of spam-control techniques. As before, this is not a complete list, but it covers the main ways the industry currently filters spam.

Usually, the techniques are employed in a cocktail approach. A technique is embodied in one or more tests, each of which calculates a probability that a particular message is spam. The probabilities are intelligently combined to maximize the filter's accuracy.

Do not think of the tests simply as ways to find spam; tests may also conclude that a particular message is legitimate.

In this report, we do *not* discuss discredited techniques, such as challenge/response blocking, e-postage, or "proof of work" schemes. (For more on such techniques, see [www.ferris.com?p=316898](http://www.ferris.com?p=316898).)

### *Connection Analysis*

These techniques act on knowledge of the incoming connection in the early stages of receiving the message—when the connection is first made. Ideally, a spam filter would conclusively decide at this stage whether a message is spam. Techniques that analyze the connection are the most efficient of the three categories. State-of-the-art spam filters can identify and reject most spam at this early stage, without needing to receive any of the message.

- *Reputation of the sending IP address.* Many spam filters compare the sending Internet address against a list of known good and bad sender addresses. For example, say the reputation database indicates that email from the address range 1.2.3.0 to 1.2.3.255 is likely to be spam. If an incoming connection is from IP address 1.2.3.4, it is probably from a spammer. Think of an IP reputation database as a more sophisticated form of blacklists and whitelists.
- *Profiling the sender's operating system.* By sending carefully designed network packets to probe the sender's machine, one can deduce which operating system the sender is running. This can be extremely useful, especially when dealing with spam sent by botnets. For example, it is extremely unlikely for legitimate email to be received *directly* from a PC running Windows ME.

- *Standards compliance.* Legitimate senders of email tend to have their connections configured in a way that is strictly correct and adheres to published standards (for example, RFC 1912). Spammers are less likely to be as careful.
- *Throttle.* Many spam senders simply give up when sending to mail servers that deliberately slow down the connection.

### ***Behavioral Analysis***

These techniques act by understanding the behavior of the connection while receiving the message. Because some of these techniques can delay delivery, they are often employed selectively—for connections that are deemed suspicious.

The techniques generally examine the standards compliance of the sender—mainly, compliance with the SMTP email transfer protocol. They are designed to uncover how the sender behaves when faced with certain responses.

- *Greylisting.* SMTP allows a receiving system to interrupt a transfer while reporting a temporary error—for example, `451 4.7.1 Please try again later`. The correct behavior for a sender would be to disconnect and retry the transfer later. Some spam senders will not bother to retry, thus removing the need to decide whether this connection is sending spam.
- *Nolisting.* The recipient domain publishes fake primary and tertiary MX records, leaving only the secondary record referring to a true mail server. Legitimate senders will attempt to contact the primary and then retry with the secondary. Many spammers will either only try the primary or try the tertiary first. Either way, it is a sign that the sender is a spammer.
- *Greetpause.* A receiving mail server is expected to reply to an incoming connection with a coded informational greeting—for example, `220 mail.example.com ESMTP Service ready`. The sending machine is expected to wait until it has received this greeting and must not start sending any information before then. If the receiving spam filter deliberately delays the greeting, and the sender does not wait correctly, the receiver can conclude that the sender is probably a spammer and reject the connection.

### ***Content Scanning***

These techniques act after having received the message. They scan the content of the message, including its headers. As such, content scanning is the least efficient of the three categories.

- *Format standards compliance.* Sloppily designed email headers or HTML body text are possible signs of spam.

- *Reputation of responsible domains.* Email headers and other metadata contain domain names that belong to the sender, and the spam filter can look these up in the reputation database. The domain information can be forged, but the filter may be able to verify the domain using sender-authentication techniques. For example, if the message purports to come from paypal.com and the spam filter fails to verify the message's DKIM digital signature, it is probably a phishing message. (For more information, see [www.ferris.com?p=320238](http://www.ferris.com?p=320238).)
- *Reputation of call-to-action text.* Spammers usually need to include some way for you to contact them (for example, a Web site where you can buy their fake pharmaceuticals, an email address, or a phone number). Reputation databases can store opinions about the text used in these calls to action.
- *Statistical content analysis.* By performing a statistical analysis of the relationship between words displayed in the message, a spam filter can learn to distinguish between spam and legitimate email (using techniques such as Bayesian classification, Markovian discrimination, or support vector machines).
- *Heuristics.* Spam filters often use rules of thumb about features that distinguish spam from legitimate email. Examples include the proportion of capital letters to lowercase or evidence that a spammer has misconfigured the spam-sending software (for example, mistakenly included text such as %RANDOM\_WORD).
- *Conversation tracking.* If the spam filter can identify the message as a reply to a local user, it is almost certainly legitimate.

## Controlling Spam: How and Where

Spam can be filtered at the desktop, on the mail server, at the network perimeter, or outside it. Filtering at or outside the perimeter is more accurate and reduces the load on an organization's infrastructure.

Spam can be filtered by installed software, prebuilt appliances, or a hosted service. Hosted services also go by other descriptions, such as managed, on-demand, cloud, or software as a service (SaaS).

Because of the inherent economies of scale in hosted services, they may cost less than an on-premise approach (based on an analysis of the *total* cost of ownership).

Some customers resist the hosted option because of privacy or regulatory concerns, for example. Such concerns are usually unfounded or based on misunderstandings.

### *The Key Role of Reputation Services*

A number of anti-spam techniques rely on *reputation*. A reputation service is an Internet-hosted database maintained as a service. In addition to being an important part of almost all hosted services, it is equally useful to an on-premise spam filter.

A good reputation service keeps threats off the organization's network and protects the organization's infrastructure from floods of spam.

The best such services are more than simple blacklists of the network source IP addresses of messages. They also track the reputation of the sending domain and any domains referenced in the message body. In addition, they may track the reputation of any file attachments.

Effective reputation services should be usable by a variety of anti-spam techniques. For example, if an IP address is known to belong to a zombie PC, that information is relevant not only for sender reputation but also for call-to-action reputation—zombies can provide services to hide the Web sites referred to in spam.

Reputation services should employ a feedback loop so that the receipt of spam from a sender automatically lowers the sender's reputation. But reputation databases should be global in their outlook, not just fed by spam received in one region.

## **Conclusion: An Arms Race**

Spammers and anti-spam technologists have for some time been locked into a sort of "arms race." Spammers continue to seek new ways to evade spam filters, but the best spam-control technologies are quick to counter these developments.

*Author: Richi Jennings*

*Editor: Mona Cohen*



# Trend Micro Interview

*We spoke to John Maddison, vice president of Core Technology Solutions at Trend Micro, and asked for his perspectives on spam and spam control. Here is what he had to say:*

## ***What combination of anti-spam techniques does Trend Micro use?***

For connection analysis, we use the industry's first and largest email reputation service, which is correlated with Web and file reputation in our Smart Protection Network. This, plus our connection management and behavior analysis, blocks email at the source. For content scanning we provide another layer of reputation services for calls to action, as well as advanced heuristics and multilingual spam detection. This is all supported by statistical analysis.

## ***One of your more unusual features is support vector machine learning. In what ways do you find this better than the more common statistical methods?***

Spammers are directly attacking traditional anti-spam learning systems, such as naïve Bayesian. The learning infrastructure that we built into our support vector machine is combating these attacks very effectively. We selected this approach over Bayesian because it has better accuracy, predictability, performance, and maintainability.

## ***As a hosted anti-spam provider, what are your main challenges?***

Often, customers have misconceptions about hosted offerings with concerns about availability, control, privacy, performance, and accuracy. However, our hosted email security solution is more effective than many popular on-site anti-spam solutions, and we have a very aggressive service-level agreement, which provides money-back commitments to availability, latency, spam blocking, false positive rates, zero virus infections, and support. These guarantees are generally much better than what can be achieved by on-site solutions and for a lower cost.

## ***How does Trend Micro improve the experience of email end users and messaging managers today? What's next?***

As threat volumes continue to increase exponentially, a signature-based approach isn't sustainable. With the cloud-client architecture of our Smart Protection Network, lightweight clients query in-the-cloud technologies. This reduces the burden on endpoints while providing immediate access to threat intelligence. And the accuracy of our in-the-cloud, correlated databases constantly increases as we receive ongoing feedback from our automated, global feedback loops. Our customers protect their investment by deploying a solution that will also protect them against tomorrow's threats.

# About This Report

## *Trend Micro's Sponsorship of This White Paper*

[Trend Micro](#) commissioned this white paper with full distribution rights. You may copy or freely reproduce this document, provided you disclose authorship and sponsorship and include this notice. Ferris Research independently conducted all research for this document and retained full editorial control.

## *Ferris Research*

Ferris Research is a market research firm specializing in messaging, electronic content control, compliance, e-discovery, and data leak prevention. To help clients track the technology and spot important developments, Ferris publishes reports, white papers, bulletins, and a news wire; organizes conferences and surveys; and provides customized consulting.

In business since 1991, we enjoy an international reputation as the leading firm in our field and have by far the largest and most experienced research team in our core competencies. Our clients include 300 of the world's 1,000 largest organizations as well as computer vendors from the largest corporations to small startups.

While other analysts have come and gone, we have published more than 200 formal reports and 1,100 short bulletins since 1991. Our news service covers more than 2,000 highly specialized announcements annually. In short, our technology and industry depth helps you understand today's products, where they have come from, and where they are going.

Ferris Research is located at One San Antonio Place, San Francisco, Calif. 94133, USA. For more information, visit [www.ferris.com](http://www.ferris.com) or call +1 (650) 452-6215.

## *Free News Service*

Ferris Research publishes a free daily news service to help you keep current on messaging, electronic content control, compliance, e-discovery, data leak prevention, and related topics. To register, go to [www.ferris.com/forms/newsletter\\_signup.php](http://www.ferris.com/forms/newsletter_signup.php). In addition to our daily electronic newsletter, you will receive periodic emails announcing new Ferris reports or webcasts. To opt out and suppress further email from Ferris Research, click on the opt-out button at the end of each email.

# Recent Ferris Research Reports

Email Archiving: On-Premise vs. SaaS  
Total Cost of Ownership of Microsoft-Based Unified Communications  
Symantec Buys MessageLabs: Transaction Assessment  
Proofpoint Buys Fortiva: Transaction Assessment  
Return Path Buys Habeas: Transaction Assessment  
Synchronica Gets Funding: Transaction Assessment  
Synchronica Buys AxisMobile: Transaction Assessment  
McAfee Buys Reconnex: Transaction Assessment  
Corporate Penetration of Instant Messaging: A Survey of 136 Organizations  
Ensim Unify: Automating the Health of a Microsoft Infrastructure  
Email Support Staff Requirements and Costs: A Survey of 136 Organizations  
Kazeon Information Access Platform  
E-Discovery Best Practices: The IT Perspective  
Email Products: Market Shares, Versions Deployed, Migrations, and Software Cost  
Archiving of Electronic Information: Key Laws and Regulations  
The Central Role of Classification for Compliance and Content Control  
Email Disaster Recovery: The Current Landscape  
Exchange 2007 Implementation Issues  
Archiving of Electronic Information: Key Laws and Regulations  
Email Archiving: Purchase Motivations and Product Selection Criteria  
Email Archiving: Best Practices  
Email Archiving: Recent Product Announcements  
Sherpa Software's Archive Attender for Exchange  
Email-Borne Virus Trends  
Continuous Availability With Neverfail  
SocialText's Corporate Wiki  
Enhancing Outlook Web Access  
Spam and Other Email Threats: Market and Technology Update  
Email Sender Authentication  
Four New Messaging Products and Services  
Achieving Regulatory Compliance With Email and Internet Content Security Policy Enforcement  
Key Messaging Issues: 2007 and Beyond  
Meeting the Challenge of Email Discovery  
Spam Control: The Current Landscape  
Planning and Implementing an Email Archiving Solution