Vendors Fight Spam's Sudden Rise

Neal Leavitt

Companies have been complaining about spam for years, and vendors have come up with different ways to fight the deluge of unsolicited e-mail. In fact, Bill Gates—touting Microsoft's plan to attack the issue—told an audience at the 2004 World Economic Forum in Switzerland that technology would reduce spam levels so much that it would no longer be a problem within two years.

Not only was that prediction wrong, but spam levels have actually risen dramatically since October 2006 after generally increasing only moderately during the prior two years, as Figure 1 shows.

Antispam vendor Message Labs reported that spam volumes rose from 82.1 percent of all e-mail in September 2006 to 89.4 percent in November. The annual average was 68.6 percent for 2005 and 86.2 percent for 2006.

The current spam flood has been so great that it has increased overall e-mail volumes by a third during the past 12 months, noted Message Labs senior antispam technologist Matt Sergeant.

Ferris Research estimated that spam is costing organizations $75 billion globally—including $20 billion in the US, $7 billion in Germany, and $3.5 billion in the UK—in antispam product purchases, lost productivity caused by overloaded e-mail systems, and users spending time wading through junk e-mail in their inboxes.

Many companies are spending still more money trying to improve their antispam capabilities by increasing their server capacity and bandwidth, said Sergeant.

There has been talk of changing basic e-mail technology—such as the Simple Mail Transfer Protocol—that has been in place for a long time, but experts say that won't occur for years, if ever.

Spammers have learned to counter vendors' efforts to recognize and filter out, quarantine, or mark as a potential problem unsolicited e-mail, noted Doug Bowers, security vendor Symantec's senior director of antiabuse engineering.

This has yielded an arms race of sorts between spammers and antispam vendors, according to Dmitri Alperovitch, antispam vendor Secure Computing's principal research scientist.

Governments are implementing regulations designed to curb junk e-mail. However, many industry observers doubt these measures will succeed.

DEFEATING ANTISPAM MEASURES

The volume of spam has boomed as it has become a source of revenue for senders, who use it to advertise products and run scams. For example, spam is a critical part of phishing schemes, which direct victims to fake yet seemingly authentic Web sites for banks or credit-card companies and convince them to enter personal information such as Social Security numbers.

Spam originates in countries throughout the world, making it difficult to find and stop. According to Ed Moyle, an analyst with market research firm Security Curve, 27 percent of spam originates in the US and 26 percent in China. Other significant spam sources include Brazil, France, India, Russia, South Korea, and the UK.

Spammers, said Moyle, have been clever in figuring out ways to counter techniques for identifying potential junk mail. For instance, in recent months, spammers have been targeting their e-mail to specific audiences. As an example, they have hit IT and legal firms with spam that contains numerous technology or law-related terms and buzzwords in an effort to dupe traditional antispam filters customized for these types of businesses.

Identifying e-mail senders

One technique used in antispam software is determining whether an e-mail message comes from a source on a blacklist of known spammers.

Blacklist owners look for IP addresses that are the source of large spam volumes or that are owned by known spammers, noted Sergeant. Companies then subscribe to the service, which owners regularly update, and configure their e-mail servers to deny inbound mail from IP addresses on the blacklist.

Some antispam products use e-mail reputation systems. In process-

Published by the IEEE Computer Society

ing large amounts of customers' e-mail, they rate sources on a continuous scale—representing the risk of accepting communications from each—based on their historical message-sending patterns. Antispam systems can apply these scores to reject mail from high-risk sources, explained Alperovitch.

Recently, Barracuda Networks observed a new trend of *pulsing zombies*, in which sophisticated spammers send out a large burst of e-mail through a particular computer, leave it dormant for a while so that vendors will remove its IP address from blacklists or improve its reputation score, and then resume using the machine, noted Stephen Pao, the company's vice president of product management.

A key way that spammers defeat blacklists or reputation systems is to use increasingly sophisticated botnets, which they can buy or rent from underground or organized crime sources.

Botnets are networks of thousands to hundreds of thousands of *zombie* computers hijacked without the owners' knowledge and then used to distribute huge amounts of spam or to launch denial-of-service or other attacks. Secure Computing reports 450,000 new zombies daily.

Spammers update zombies' software and order them to launch attacks via file-sharing channels and technologies such as Internet relay chat, instant messaging, and Internet telephony, said Marcus Sachs, a computer security researcher at SRI International, a nonprofit research corporation, and director of the Internet Storm Center at the SANS Institute, a computer-security training and research organization.

When this happens, the junk e-mail that zombies send comes not from the spammer's IP address, which might be on a blacklist, but from those of hijacked computers, which usually aren't.

Analyzing e-mail content

Many antispam techniques try to recognize unsolicited e-mail by analyzing words within messages. These approaches either look for specific words or statistically analyze sets of words to determine the likelihood that a message is spam, explained Christine Drake, product manager for antispam and antivirus vendor Trend Micro.

Other techniques help identify tricks used to obscure words that indicate spam, such as replacing letters by numbers or symbols or inserting spaces between letters, as in v1@gr@ or v i a g r a. Vendors can train these systems to improve their performance.

Spammers have countered with *image spam*, a technique that embeds traditional spam text into one or more image files. Content filters then see one or more images, rather than text it can analyze.

In early 2006, image spam accounted for about 10 percent of the overall spam volume. Since October 2006, Secure Computing said, image spam increased 200 percent and now accounts for 30 percent of all spam and one in every four e-mail messages.

Antispam vendors have tried using optical character recognition (OCR) to scan image spam and recognize letters. However, this approach has been largely ineffective.

To begin with, OCR can be notoriously inaccurate, in part because the software can't always correctly recognize letters and symbols in the many fonts and formats in which they can be shown. Because of this, said Alperovitch, "Spammers have quickly learned to vary fonts and sizes, and employ different foreground and background colors to make it difficult for OCR software to extract the textual message from a spam image."

Adding dots also sometimes keep software from detecting words in images, Bowers said. And, he added, spammers have begun using *captcha* (completely automated public Turing test to tell computers and humans apart) technology, which creates text with distorted characters that can be identified by humans but not by computerized systems.

In response, Trend Micro recently submitted a patent application for Adversarial OCR, a technique that looks for only certain words that indicate spam, rather than trying to recognize all words in an image. This is more efficient and targeted, said Drake.

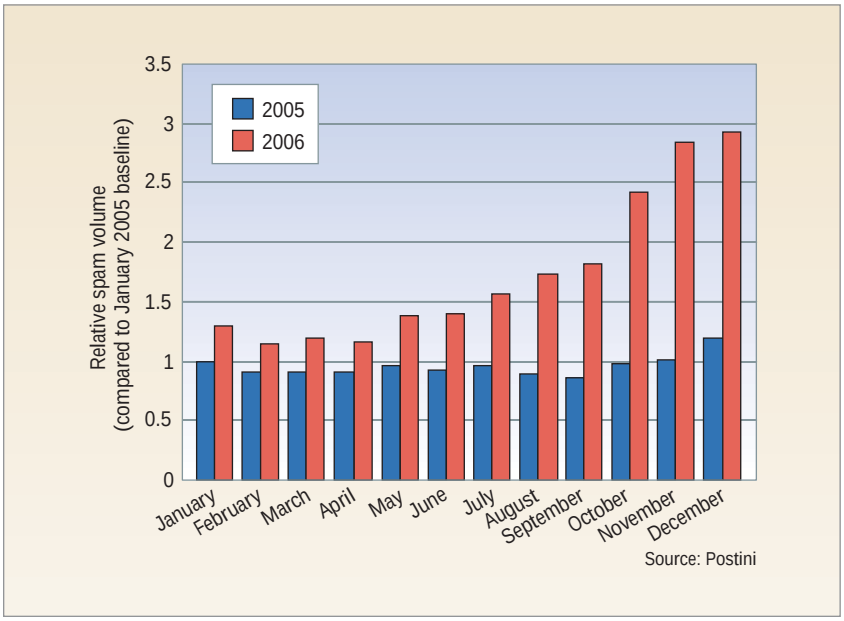

Figure 1. After increasing only moderately during the prior two-year period, spam volumes have increased dramatically since October 2006.

Detecting spam as it hits the Internet

Some antispam systems identify spam by detecting when large numbers of the same message reach the Internet. These systems record the content of the message and stop subsequent identical e-mails. Many of these systems, noted SRI's Sachs, use fake e-mail accounts to identify quantities of spam with the same subject line and content, then direct customers' mail servers to drop these messages when received.

Products such as messaging-management vendor Postini's Connection Manager record and compute hash values for e-mail messages' content and send the values to a central server to record how many times it sees them in other messages. Once the system sees messages with the same hash value a certain number of times, it declares the message to be spam.

However, Alperovitch noted, such systems are prone to false positives when encountering newsletters and other legitimate material sent to large numbers of people.

To fight these systems, spammers use sophisticated software to insert text and modify pixels to randomize the content of e-mail messages so that they don't look like the same message to antispam products.

"It's fairly straightforward to write a simple image-manipulation program that can do this," said Adam Swidler, Postini's solutions marketing senior manager.

Pump and dump

According to Pao, many people have been spammed by stock-touting scams, also known as *pump-and-dump* schemes.

In most of these cases, spammers buy many shares of penny stocks—which don't trade on one of the major stock exchanges and which generally sell for less than $1 per share—and promote them in e-mail messages. If even a few dozen investors pick up on those stocks, the momentary uptick can create a windfall for the spammer.

A joint Purdue University/Oxford University study last year showed that online investors who fall for these schemes can lose up to 8 percent of their investment in just two days and in the long run lose an average of $52.50 for every $1,000 invested.

The study reported that about 15 percent of all e-mail, or about 100 million messages per week, is pump-and-dump related.

Despite vendor efforts, spam levels have jumped dramatically since October 2006.

The US Securities and Exchange Commission began prosecuting these illegal schemes in earnest in 2005. Thus, most pump-and-dump activity has moved overseas, largely to Russia, Eastern Europe, and Asia.

FIGHTING BACK

Antispam vendors, businesses, governments, and antispam organizations are taking steps to try to stem the rising flood of unsolicited e-mail.

Upgrading antispam efforts

Antispam companies are beefing up their servers and research in an effort to better and more quickly recognize spam.

Businesses are also increasingly using managed antispam services, in which a company's e-mail goes through a filtering server that is often in an antispam vendor's facility.

"Managed services require no hardware or software on client premises, thereby removing the need for onsite installation, maintenance, and additional infrastructure complexity," said Message Labs' Sergeant. "It's a predictable cost structure with a limited need for internal resources for ongoing management and support."

Barracuda installs its spam firewall in front of a company's e-mail server at the customer's site and processes mail before it reaches users.

Some companies are using dedicated antispam appliances—by vendors such as Barracuda, Cisco Systems, Secure Computing, and Symantec—rather than antispam applications running on a central server, explained Natalie Lambert, a senior analyst for Forrester Research.

Governmental efforts

Worldwide governmental efforts to combat spam have had limited success.

The Finnish government says it has helped reduce the volume of spam so much that it represented only 30 percent of all e-mail received in the country last year, down from 80 percent in 2003.

The government has set up several enforcement agencies that work together to fight spam. For example, the Finnish Communications Regulatory Authority supervises message filtering and the Consumer Ombudsman and Agency supervises illegal marketing. The government also passed a spam law in May 2004 establishing guidelines for online commercial practices and penalties for violations.

"We still need to beef up our enforcement activities in the European Union (EU)," said European Commission spokesperson Martin Selmayr, "and that will require well-equipped national regulators." The EC is planning to review the matter during the next few months to determine whether more laws are needed.

As part of 2002's ePrivacy Directive, the EU mandated that businesses can't send e-mail marketing messages without prior consent, except to people with whom they have had a business or commercial relationship. The directive also said that when messages are sent, recipients must be given a way to opt out of future communications.

In 2003, the US Congress passed the Can-Spam (controlling the assault of nonsolicited pornography and marketing) Act, which requires e-mail marketers to meet a series of requirements, including giving

potential recipients a way to opt out of getting unsolicited commercial messages.

Nonetheless, spam has grown in volume.

Some vendors are skeptical about whether legislative efforts are making inroads. "It's doubtful whether government regulations can have any real effect on spam, partly because of jurisdictional problems, partly because many such efforts to date have been heavily watered down by direct-marketing lobbyists," said Era Eriksson, senior content-filter researcher for security vendor F-Secure.

International antispam efforts

Six of the world's largest antispam organizations—the Organization for Economic Cooperation and Development, Asia-Pacific Economic Cooperation, EU Contact Network for Spam Enforcement Authorities, International Telecommunication Union, London Action Plan for Spam Enforcement, and the Seoul-Melbourne Memorandum of Understanding—formed StopSpamAlliance.org last year to share information on matters such as best spam-reduction practices and antispam laws.

This marks the first coordinated international effort to deal with spam. "A lack of international cooperation and governmental efforts has helped spammers," said Sergeant.

Despite the best efforts of antispam vendors, companies, and others, spam will continue to grow, said Forrester's Lambert. "There is so much money in it, that it is simply a fact of life," she explained.

In fact, Secure Computing predicts that spam will represent 95 percent of all e-mail by the end of this year.

"At the end of the day, spammers are driven by profit," explained Bowers. "They'll continue to evolve their techniques to evade the latest and greatest antispam technology. Spam will be a threat until this equation changes."

Thus, like antivirus companies, antispam vendors will find themselves constantly reacting to new spam techniques.

"The cycle will continue until it becomes unprofitable for spammers to send out unsolicited e-mail," said Security Curve's Moyle. "If the return was so small that it did not outweigh the risks of violating the law, the effort in sending the mail, and so on, I think they'd stop doing it. In the meantime, no amount of technology or legislation will prevent spam, given the strong economic incentive to send it out."

"Current techniques for combating spam are like trying to fight pollution by issuing gas masks," said SRI's Sachs "We are putting the onus of protection on the victims rather than taking steps to stop the polluters. Until we get serious about catching and prosecuting the spammers, this problem will not go away." ■

Neal Leavitt is president of Leavitt Communications (www.leavcom.com), a Fallbrook, California-based international marketing communications company with affiliate offices in Brazil, France, Germany, Hong Kong, India, and the UK. He writes frequently on technology topics and can be reached at neal@leavcom.com.

Editor: Lee Garber, *Computer*, l.garber@computer.org