

Spam Filtering Technology

A Growing Problem

As of November 2002, we estimate that 65% of all email traffic on the Internet is unsolicited bulk email ("Spam"). Spam is one of the fastest growing, most complex problems facing the Internet today. The problem has already led to millions of dollars in lost productivity and added systems costs for corporations and small businesses.

Businesses who maintain in-house email servers are fighting a losing battle to protect their systems from this epidemic because complexity of the problem is constantly evolving. As each week goes by, Spammers are growing wiser, learning new tricks to fool common Spam defenses and obtaining more sophisticated software to penetrate more mailboxes.

Is There A Solution?

Yes, luckily there is: Evolve your Spam defenses as fast as Spammers evolve their techniques.

Easy Email Spam Filtering accomplishes this by gathering real-time Spam intelligence from a number of sources, and then actively using this intelligence to block unwanted email. Easy Email tracks tens of thousands of live Spam email characteristics ("DNA"), which by itself identify 85% of all Spam. In addition, approximately 25 third-party Spam databases, several DNS checks, and several message-formatting tests are used when analyzing each email.

In all, approximately 45 constantly evolving tests are performed on every email that enters the system. The results of these tests are combined using a methodical weighting system that successfully identifies over 95% of Spam.

Weighted Tests

There are two important factors to consider when dealing with Spam:

- No single test can identify all Spam.
- Tests will sometimes falsely identify legitimate email as Spam.

Therefore, Easy Email allows no single test to cause an email to be flagged as Spam. Instead, multiple tests are used in conjunction via a weighting system, where each test is assigned a point value. When an individual test fails, its point value gets added to the email's overall weight. If the total weight of the email is greater than a certain threshold, the email gets flagged as Spam. Normally it takes a minimum of two to four tests to fail, depending on the severity of the tests, before the total weight reaches the necessary threshold.

Spam "DNA"

Easy Email collects thousands of Spam samples through the use of dummy mailboxes ("Spam traps") as well as data submitted by clients. Spam messages are then broken down into identifiable components, which are used to develop Spam DNA. Spam DNA is similar to anti-virus "fingerprints", and can accurately identify most Spam based on specific content that would only be found in certain emails.

Advanced pattern recognition technology allows Easy Email to simultaneously apply thousands of heuristic algorithms to each email in search of parts of the email that are identifiable by the Spam DNA. When a match is found, the email fails this heavily weighted Spam test.

Open Relays and Known Spam Sources

Open relays are the most common source of Spam. These are improperly configured mail servers that allow outbound mail to be sent by just about anybody. Spammers use automated tools to

Spam Filtering Technology

search the Internet for vulnerable servers, and then hijack these servers to increase the amount of Spam they can send.

To combat this problem, there are several third-party organizations that maintain active databases that "blacklist" open relays. Databases also exist that blacklist other known Spam sources such as proxies, insecure web forms, and dial-up IP addresses.

Easy Email tests each email against approximately 25 of these well-known blacklist databases, including those from Osirusoft, SpamCop, SpamHaus and NJABL. These databases are updated multiple times each day. Each blacklist test that fails adds an additional weighting to the email.

DNS and RFC Violations

Spammers tend to be sloppy in how they send email. Therefore it is important to scrutinize each inbound email to see if it followed all of the rules defined by current RFC's. The following tests examine the mail server that delivered the email as well as the domain name used in the sender's return address:

- Did the mail server falsely identify its self in the "HELO/EHLO" data?
- Is the mail server missing its reverse DNS record?
- Is the domain missing "A" and "MX" DNS records?
- Is the domain missing "postmaster" and/or "abuse" addresses?
- Does the domain not accept delivery status notifications?

The email's message headers are also examined:

- Are the message headers improperly formatted or missing required data?
- Are the message headers in a format consistent with Spam?
- Is the return address in a format consistent with automated mailers?

Blocking based on any of these tests alone would block a large amount of legitimate email. However, when used in conjunction with a weighted filtering system, these tests are extremely effective.

Geographical Routing

This test analyzes the Internet route that an email travels, and looks for highly inefficient routing that is very common in Spam. For example, an email might fail this test if it is sent from an account in the U.S. to another account in the U.S., but is routed through a server in Korea.

A second geographical test uses recent Spam statistics to identify countries that have a high probability of sending Spam. If an email travels through one of these countries, a variable weight is added to the email based on the country's current Spam rate.

Elusive Spammers

Spammers are very aware of the filtering techniques used by top-tier email providers. This has led Spammers to develop creative tactics and advanced software in an attempt to bypass filtering systems.

For instance, many Spammers now use binary encoding to hide their text and html email from signature-based filters. Easy Email's pre-processors decode binary mime segments, rendering this trick useless against the Spam DNA filters. To take things a step further, because there is no

Spam Filtering Technology

legitimate reason why a text or html email should be binary encoded (other than to bypass Spam filters), an additional weighting is added when this type of email is discovered.

More aggressive Spammers are also engaged in "Polymorphic Spam Attacks". These are attacks where many copies of the same email is sent, but with each email containing subtle differences in punctuation or spacing designed to circumvent content filters. Easy Email uses its Spam DNA to recognize similar pieces of these different email versions and trap them in Spam filters.

Spam is often times sent through several different mail servers before arriving at its final destination. This is done to disguise the original source of the email, in an attempt to avoid Spam blacklists. However, unlike most filtering systems, Easy Email traces the email back to its origin and scans each hop along the way. If any server that the email was routed through fails a blacklist or DNS test, additional weighting is added to the email.

Habeas SENDER WARRANTED EMAILSM (SWESM)

Easy Email has integrated technology from Habeas (www.habeas.com) that helps to ensure that legitimate email does not get blocked as spam. Habeas has created an enforcement system called SENDER WARRANTED EMAILSM (SWESM) that enables individuals and companies to warrant that the email they are sending is not Spam. The SWESM technology is based on existing trademark and copyright laws, which allows Habeas to obtain enforceable injunctions and judgments against Spammers through the Courts.

What this means is that any email sent under the Habeas system is warranted (guaranteed) not to be Spam, and meets the exacting standards for a HABEAS COMPLIANT MESSAGESM. Habeas has some of the strictest standards in the industry, including only permitting mailing lists that are verified opt-in. When a Habeas SENDER WARRANTED EMAILSM arrives to the Easy Email system, it is delivered to the recipient with confidence that it is Spam free.

Combining the Tests

After rigorous testing, a final weighting is assigned to each email. A threshold is then used to determine if the email should be identified as Spam. If the email's weighting is lower-than the threshold, the email is deemed to be Spam-free and is delivered normally. However, if the email's weighting is equal-to or greater-than the threshold, the email is identified as Spam.

Each domain administrator using the Easy Email email platform can control the sensitivity used to determine if an email is Spam, and can override that sensitivity for individual users. A high sensitivity causes a low threshold to be used and catches almost all Spam. This is the optimal setting for most companies. A low sensitivity causes a higher threshold to be used, and catches only the most obvious Spam. This can be used as a safeguard if there are concerns that legitimate email is getting blocked. In most cases this is not necessary.

What To Do With Spam

Once a Spam email has been identified, there are then several actions that can occur. The following actions are available to each domain administrator. These can also be overridden for specific users:

- Delete the email immediately. This is usually not advised, just incase a legitimate email gets falsely identified as Spam.
- Deliver to a Junk Mail folder. This will allow each user to review their Spam in their own WebMail folder.

Spam Filtering Technology

- Deliver as an attachment. This hides Spam from the user, and delivers them an alert message notifying them that they received the attached Spam email.
- Deliver to an alternate email address. This is useful if a company wants to have one administrator review all of the Spam that is received.
- Add text to the beginning of the subject. This will allow each user to set up custom filtering rules inside of MS Outlook or other email programs.
- Add text to the beginning of the body. This will discreetly tell the user that the mail is Spam, and may be beneficial in cases where a company initially is paranoid about using the more aggressive actions.
- Add text to the end of the body. There is not much use for this action, unless you want to tell the user "You have just read Spam!"

When domain administrators use the "Deliver to a Junk Mail folder" action, an additional feature becomes available. An automated cleanup script can be run each day to delete old Spam from users' Junk Mail folders. The domain administrator can specify how many days to keep Spam in folders, or how many total Spam emails to keep before old ones are deleted.