SPAM DNA FILTERING®

The Spam Problem

We estimate that 85% of all email traffic on the Internet is spam. Spam is the most complex problem facing the Internet today. The problem has already led to millions of dollars in lost productivity and additional infrastructure costs for businesses and service providers.

Businesses who maintain in-house email servers are fighting a losing battle to protect their systems from spam because the complexity of the problem is constantly evolving. Spammers are growing wiser on a daily basis, learning new methods to manipulate common spam defenses and obtaining more sophisticated software to trick spam filters and penetrate email inboxes.

Solving the Problem

In order to win the war against spam, businesses must evolve their spam defenses faster than spammers evolve their techniques.

Our Spam DNA Filtering® system accomplishes this by gathering real-time spam intelligence from a number of sources and then actively using this intelligence to block the spam. We track tens of thousands of live spam email characteristics ("DNA"), which alone identify the majority of spam. In addition, a number of third-party spam databases, several DNS checks, and several message-formatting tests are used when analyzing each email.

We aggregate all of this data into a collection of several thousand constantly evolving spam tests that are performed on every email that enters the email hosting system. The results of these tests are combined together to identify more than 98% of spam with virtually zero false-positives.

Weighted Tests

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

- No single test can identify all spam.
- Some tests will falsely identify legitimate email as spam.

Therefore, we allow 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, that point value is added to the overall weight. If the total weight of the email is greater than a certain threshold, the email is flagged as spam. Each customer can set his or her own spam filtering sensitivity level, which is used by this weighting system.

Spam "DNA"

We collect tens of thousands of spam samples through the use of dummy mailboxes ("spam traps") as well as data submitted by customers via the "BlackList Sender" button within webmail. Spam messages are normalized to remove content obfuscation and then broken down into identifiable components, which are used to code the DNA. Spam DNA is similar to anti-virus "fingerprints," and can accurately identify most spam based on specific content found in spam emails. When a match is found, spam points are added to the email.

Zombies, Open Relays, and Known Spam Sources

A "zombie" is a computer that has been taken over by a spammer, and which is used to send out bulk mailings without the computer owner knowing. Normally this occurs because the computer owner opened a virus, which gave a spammer a back door into their system. This is one of the most common sources of spam.

Another prevalent source of spam is "open relays"—insecure mail servers that can be used freely by spammers. Spammers use automated tools to scour the Internet in search of vulnerable mail servers, and then hijack those servers to increase the amount of spam they can send.

To combat this problem, there are several third-party organizations that maintain databases ("blacklists") that list the IP addresses of these compromised machines. There are also databases that list known professional spammers. We have arrangements with approximately 15 of these organizations, so that our system can download full copies of the blacklists hourly and incorporate them into the weighted spam filtering system. When an email arrives from a blacklisted IP address, spam points are added to the email.

DNS and RFC Violations

Spammers tend to be careless in how they send email. Therefore it is important to scrutinize each inbound email to see if it followed the rules defined by current Internet standards. For example, below are just a few of the tests that examine the sending mail server and the message:

- Did the mail server falsely identify itself in the "HELO/EHLO" data?
- Does the mail server have a missing or invalid reverse DNS record?
- Is the domain missing "A" and "MX" DNS records or using illegitimate values?
- Was there an SPF violation? (Was the email sent from a mail server that is not authorized to send mail using the sender's domain name?)
- Are message headers improperly formatted or missing required data?

There are more than 1,000 of these tests. Blocking based on any test alone would block a large amount of legitimate email; however, these tests are extremely effective when used in conjunction with a weighted filtering system.

Elusive Spammers

Spammers are very aware of the filtering techniques used by top-tier email service providers such as us. This has led them to develop creative tactics and advanced software in an attempt to bypass filtering systems. For instance, many spammers use binary encoding to hide their text and HTML email from signature-based filters. It is also very common for spam to include invisible HTML code intermixed with the visible content, and subtle variations in wording and punctuation, as well as purposely misspelled words.

We use pre-processors to "normalize" these emails by decoding binary MIME segments, removing hidden HTML, ignoring punctuation, and recognizing alternate characters used in misspelled words. After the email is normalized, the content is scanned. Additionally, since there is no legitimate reason why an email would contain such formatting tactics other than to bypass spam filters, spam points are added to emails that use this trickery to obscure the email content.

Combining the Tests

After rigorous testing, a final weight is assigned to each email. The weight is compared against a user-defined threshold to determine if the email should be identified as spam. If the email weighting is lower than the threshold, the email is deemed to be spam-free and is delivered normally. If the weight is greater than the threshold, the email is identified as spam.

What to Do With Spam

Once a spam email has been identified, there are several actions that occur based on user preferences:

- **Delete the email**. In this case, users will never see the spam.
- **Deliver to Spam folder**. This will allow each user to review the emails that have been tagged as spam. This folder can be viewed from webmail or IMAP, and settings are available to automatically purge old spam from this folder after a certain number of days or number of emails.
- **Tag the subject**. The word “[SPAM]” will be added to the beginning of the subject line, and then delivered normally. This allows each user to set up custom filtering rules inside of desktop mail programs, such as Microsoft Outlook.
- **Deliver to an alternate email address**. This is useful if a company wants to have a single administrator review all of the spam that their users receive.

False-Positive Prevention

Every effort is made to ensure that legitimate emails are never falsely identified as spam (“falsepositives”). To prevent this, several filtering tests have been incorporated which are designed with the reverse approach of identifying characteristics found in legitimate email instead of identifying spam. These tests work in conjunction with the weighting system to help eliminate false-positives. Specific details of these tests cannot be provided, in order to keep this information out of the hands of spammers.

Additional false-positive protection is also acquired from ReturnPath (www.returnpath.com) and has been bundled into the Spam DNA Filtering® system. ReturnPath’s Bonded Sender and Sender Score systems are described below.

Bonded Sender

The Bonded Sender Program is aimed at eliminating the possibility of legitimate bulk email being flagged as spam by anti-spam filtering systems. Bonded Sender works with bulk mailers such as eBay, AT&T, and other reputable companies to ensure the integrity of their email campaigns. Senders must adhere to strict email standards, undergo a qualification process, and post a financial bond that enforces financial penalties for non-compliance. We allow bulk mail from these legitimate organizations to bypass the spam filters.

Sender Score

ReturnPath’s Sender Score system provides us with access to a comprehensive email reputation database. The data is compiled from more than 40 million mailboxes across a number of email providers, including Webmail.us. Sender Score looks at things such as complaint data, mail volume, bounced messages, unsubscribe functionality, security practices, identity stability, and more. Think of it as a credit score for email senders. When an email arrives, we are able to tell if the sender is a “good” sender or a “bad” sender and add or subtract spam points accordingly.

Safe Lists

While every effort is made to ensure that legitimate emails are not identified as spam, a small number of false-positives are unavoidable. To solve this problem, domain administrators and email users can specify trusted email addresses that should always bypass the filtering system. This feature should be used when specific emails sometimes gets identified as spam, such as opt-in newsletters or emails from colleagues whose mail servers are blacklisted or configured improperly.

Exclusive Filtering

For customers desiring maximum spam protection, the Exclusive filtering level can be used to block email from all email addresses not appearing on the Safe List (described above). This will cause email from all unknown senders to be flagged as spam and the user-defined spam action will be taken.

Individual User Flexibility

Domain administrators can define the default spam filtering settings for all of their users. Users then have the ability to adjust their spam filtering to their liking, or leave it at the defaults. In addition, there is both a domain-level and a user-level Safe List. Domain administrators can add email addresses and domains to the Safe List, and it takes affect for all users in the domain. Users can add create their own personal Safe List, and it takes effect just for their account. Both Safe Lists work in conjunction with each other to ensure that mail from trusted senders always gets delivered.

Network-Level Spam Security

Our email hosting system is guarded against intrusion at its border firewalls and is monitored 24x7 by Tier 1 Engineers. Global filtering rules block traffic from the most abusive spam sources on the Internet and throttles traffic from emerging spam sources. This removes the potential for spammers to send a denial of service ("DoS") attack to the Webmail.us system.

Directory Harvest Attacks

A directory harvest attack is an attempt by a malicious person to find out the email addresses that exist within a domain. Spammers do this by sending a series of connections to a SMTP server pretending to deliver mail to a large quantity of randomly selected name combinations, and collecting the responses from the server. The SMTP responses normally indicate whether or not each email address exists, thus allowing a spammer to compile a list of valid email addresses.

We protect customers from Directory Harvest Attacks by automatically disconnecting spammers who send mail to too many invalid recipients. Subsequent connections are then throttled so that the spammer cannot establish new connections at a rapid rate. This renders Directory Harvest attempts useless and greatly reduces the chances of customer email addresses ending up on bulk mailing lists.

Abuse and Blacklist Prevention

Serious measures are taken to keep abusive users off of our email hosting system and to keep our system out of the anti-spam blacklists. Acceptable Use Policy compliance is strongly enforced in order to maintain the integrity of the email service for the benefit of all customers. Monitoring systems detect when a customer tries to send out a bulk mailing, and cuts it off before it is too late, alerting engineers who then contact the customer.

Our mail server IP addresses are checked against all known anti-spam blacklists hourly. In the rare case that one becomes blacklisted, a couple of things happen. First, engineers receive an alert from the monitoring system. Immediately, engineers take action to diagnose and stop whatever caused the blacklisting to occur. During this time, the IP address that was blacklisted is removed from use so that no outgoing mail is sent from it. If an entire range of IP addresses become blacklisted, fallback routing is used to route mail through an alternate SMTP server cluster with clean IPs. Once the issue has been resolved, we contact the blacklisting organization in order to get the IP address removed from the blacklist.