Best of Class Anti-spam Solution

VisNetic MailPermit

permission-based email system

Introduction

Challenge-response anti-spam technology is currently at the center of a critical debate over spam, anti-spam tools and the fate of email as a valuable means of business communication.

Detractors fear challenge-response (CR) systems will block legitimate email traffic and may create an "endless loop" between organizations employing CR systems, adding to the overall network traffic burden created by spam. Proponents of challenge-response technology see it as the single, most effective means to combat spam bots.

In this paper, we will present a brief overview of anti-spam technology, including challenge-response. Each potential weakness debated by anti-spam authorities will be addressed. Finally, we will introduce VisNetic MailPermit, a solution that combines the 'best of class' anti-spam technologies—including a proprietary challenge-response system—to overcome these weaknesses.

Rules-Based Anti-Spam Technology

Anti-spam technology to date has been largely rules-based. While effective at identifying and blocking high percentages of spam, rules-based tools fail to block 100% of spam and often block legitimate email. Additionally, as the number of variables considered and rules established increases, so does the ability of spammers to defeat them.

How it works

In short, rules-based anti-spam tools employ header and text analysis to determine whether an email message is spam. Specifically, rules-based technology examines email headers for the circuitous routing spam typically takes and the subject line and body for words or phrases commonly found in spam. It then compares its findings to pre-defined or administrator-defined filters. Messages matching anti-spam filter criteria are deemed spam; all others are believed legitimate.

Early rules-based software utilized algorithms for header and text analysis. Running message content through a series of hard-coded programs and procedures, these solutions resulted in Yes/No decisions—identifying each message as spam or not-spam. Subsequent action on each message—accept or reject—was dependent on whether or not it was categorized as spam.

More recently developed or enhanced rules-based solutions utilize heuristic analysis to determine the likelihood that a message is spam. Unlike algorithmic analysis, heuristic processes are not reliant upon pre-established formulas. Instead, this trial-by-error method applies rules derived from experience to determine whether a message may be spam. Heuristic analysis does not simply produce Yes/No results. Rather, each message is assigned a weight or score indicating the likelihood that the message is spam. Messages are then forwarded to a mail server or email client where administrators or users can specify score thresholds and subsequent required action—accept, reject or quarantine for independent review.

A growing number of rules-based solutions now implement Bayesian filtering, a technique combining heuristics and probability analysis, enabling the software to learn and relearn how to recognize spam by scanning the mail you've read and the mail you've rejected. Probability calculation is based on each message's most unusual characteristics. Over time, the anti-spam engine learns what type of email to deliver - and what to reject or delete.

In addition to a more sophisticated method of analysis, today's rules-based anti-spam solutions include white list capabilities as well as the ability to check messages against various real-time black hole lists, or RBLs—public catalogs of known spammers and open relay servers that spammers use as conduits for their messages. White lists allow mail recipients to identify pre-approved or welcomed senders, exempting them from analysis. RBL checking adds the benefit of collaborative anti-spam efforts to an individual or organization's anti-spam solution.

Disadvantages of Rules-Based Technology

Rules-based anti-spam solutions suffer from two common disadvantages. One, spammers continue to devise methods to evade rules-based anti-spam solutions. Consequently, rules-based solutions fail to trap 100% of inbound spam. Two, they often incorrectly categorize legitimate email as spam.

Simple tactics employed by spammers include disguising words to fool key word and phrase filters defined in rules-based solutions. Spammers also continue to develop techniques to disguise their identities and the origination point of their messages, two attributes that are frequently used in filtering out spam. More sophisticated techniques aspire to overcome existing algorithms and heuristic analysis. These attempts, albeit largely unsuccessful to date, seek to create spam messages that will achieve the lowest possible score.

Critical to enterprise organizations is the false-positive outcome—the incorrect categorization of legitimate email as spam. If a legitimate email contains header, subject or body content that matches the criteria set by an anti-spam rule or variable, the email will be incorrectly handled as spam. For example, a legitimate email from a customer may include a line of yelling—or all capital letters. This email would match a common rule that looks for one or more lines of yelling and would consequently be handled as spam—rejected, deleted or quarantined. In this instance, the organization's opportunity to save a potentially dissatisfied customer is lost.

Rules-based anti-spam solutions that employ heuristic analysis and Bayesian filtering provide the best protection against spammers' attempts to "spoof" legitimate email and against false-positives. Instead of simply accepting or rejecting a message meeting one or more criteria (e.g. a valid point of origin or conversely, one line of yelling), these solutions assign a weight to each condition met, resulting in an overall score that more accurately reflects the likelihood that a message is spam. Users set the score threshold or tolerance and retain the ability to review email not conclusively categorized as spam. A byproduct of that review is a better "trained" anti-spam engine.

Challenge-Response Systems

Challenge-response, or Permission-based Email systems, are relatively new and to date, have been available exclusively through ASPs. Gaining in popularity, these systems provide the most effective anti-spam measure against computer-generated spam. Employing a verification or authorization step that can only be completed by a human being, this technology promises to eliminate nearly 100% of spam from your Inbox, without blocking email from legitimate sources.

How it works

Challenge-response (CR) anti-spam technology shifts the burden of effort from email recipients to email senders, by requiring them to obtain authorization to deliver their email message. If they do not complete the authorization or verification requirement, their email messages are blocked and never appear in your mailbox.

In short, CR systems compare every inbound email message first to a white list—immediately passing email from whitelisted, or approved senders on to the recipient—and then to a black list—automatically blocking and deleting mail from blacklisted, or unapproved senders.

If the sender is on neither list, the CR system automatically sends an email, or challenge, explaining that you are using a "permission-based" system. The email asks the sender to go to a Web page and complete a verification task. Once completed, this sender will not be challenged again. In other words, subsequent email messages from this sender will not be subjected to your anti-spam mechanism.

Answering the critics of challenge-response

Loss of legitimate email correspondence

Critics of challenge-response systems believe that challenges issued to legitimate senders, deemed by opponents as unfriendly and time-consuming, will consequently be ignored, breaking an open—and fast, cost-effective—line of communication between customers, vendors and their partners. Proper deployment of a well-designed CR solution that supports whitelists, authorizes senders via a fast and simple verification task and allows for the inclusion or exclusion of specified recipients from the CR process answers this concern.

First, upon deployment, companies must populate white lists with all known legitimate senders. For a typical small to medium business, white list entries will include active vendors, business partners, resellers and distributors, and existing customers. Enterprise organization white lists may include investors, subsidiaries, etc. Comprehensive white lists created prior to issuing the first challenge will ensure that email from known, legitimate senders will be delivered without interruption.

The more granular control an organization has over their CR white lists, the more effective their use. For example, white lists defined at the system or organization level will result in the automatic receipt of all email from white listed addresses to any individual in the organization. This means that your Sales, Accounting and IT staff

may receive emails touting the latest special promotion from one of your approved vendors—emails that are destined and more suited to your Purchasing staff. While your vendor is spared participation in your CR system and your risk of blocking legitimate email—albeit legitimate only to a single department in the company—is eliminated, one, company-wide white list compromises the initial goal—combating spam. To your Sales, Accounting and IT staff, this vendor's emails are spam and as such, continue to cost the organization real dollars in the wasted resources required to deal with it.

CR systems that support multiple, layered white lists can work to effectively remove legitimate senders from the challenge-response obstacle and still protect the organization from spam. White list capability should be present at the system, domain and user level, with user level taking precedence. In the example presented above, user level white list capability would safeguard Sales, Accounting and IT staff from unwanted vendor email without trapping legitimate business correspondence to Purchasing staff.

What if your organization solicits email correspondence from legitimate sources not yet known and consequently not white listed? Then the second item, a fast and easy verification step is essential. Challenges that require senders to complete a multi-step process can be too much work and too costly for individuals and organizations seeking to correspond with your organization via email. Additionally, a challenge that requires senders to identify an image in one simple step may also work to dissuade compliance as the time required for the image to load may exceed the sender's time-tolerance threshold.

CR solutions that require entry of one, short code, quickly loaded on the page, present a less time-consuming challenge to new senders. Moreover, a vanilla page presence devoid of marketing hype is conducive to sender compliance. A concise, easy to understand presentation of the step required increases acquiescence.

Finally, companies seeking to effectively combat spam but still reluctant to challenge potential customers should employ a CR system that allows for the exclusion of specified email boxes from the challenge-response process. This is especially important to companies wherein a single sale will generate significant revenue, increasing the potential loss if the sender fails to complete a CR system challenge. In such organizations, where any customer correspondence—sales leads, orders, service requests—is conducted via email, exclusion of Sales, Service, Support and Info boxes from the challenge requirement should be an option.

Endless Loop between CR Systems

As with the early implementations of email auto-responders, critics fear a break down if two users with a challenge-response systems attempt to communicate with each other. This potential 'endless loop' between CR systems is easily avoided if developers of CR systems implement the ability to automatically add the recipient of outbound messages to agents' white lists. This feature ensures that responses to email sent from an organization utilizing CR technology will automatically be delivered, without challenge.

Additionally, many believe that CR traffic will simply add to the overall network traffic burden created by spam. Admittedly, challenge-response systems that challenge every inbound email will unnecessarily increase email traffic over the Internet. However, anti-spam solutions that utilize challenge-response technology as a second line of defense, behind rules-based technologies for example, will send fewer challenges and result in a lower impact on overall email traffic volumes.

Combining What Works Into An Enterprise Anti-Spam Solution

Deerfield.com introduces VisNetic MailPermit, a server-based, enterprise anti-spam solution designed to overcome demonstrated weaknesses of rules-based anti-spam tools and potential weaknesses identified in early challenge-response systems. VisNetic MailPermit combines the following "best of class" anti-spam technologies:

- **Whitelists and Blacklists**

  VisNetic MailPermit includes whitelists and blacklists at the system, domain and user level, providing granular control and optimum anti-spam protection to all users in an organization. Additionally, users may opt to dynamically populate whitelists with the To address on all outbound messages—reducing administrative burden and lowering the likelihood that an endless loop will be created between MailPermit and another CR system.

- **Heuristic Analysis**

  VisNetic MailPermit includes a tightly integrated implementation of VisNetic AntiSpam, powered by SpamAssassin—the most advanced, rules-based anti-spam technology available today. Known for its comprehensive rule set and highly developed scoring method, this technology also allows for user-defined score thresholds at the system, domain and user levels—providing flexible, granular control of anti-spam measures in your organization.

- **Bayesian Filtering**

  MailPermit's implementation of VisNetic AntiSpam also includes Bayesian filtering, a mechanism by which the software learns what is or is not spam by referencing the highest scored non-spam emails. This learning tool is enabled by default and as the number of email messages filtered increases, so does the accuracy of VisNetic AntiSpam. Also included is a utility that allows administrators the ability to designate good and bad (spam) messages and then instruct the learning mechanism to utilize them.

- **Challenge-Response System**

  VisNetic MailPermit's CR system features:

  - A fast, simple verification step. To complete your challenge, senders simply go to a Web page (linked from within the challenge email) and enter a short, quickly loaded alphanumeric code into a single field.
  - Email confirmation to senders. Once the verification step is successfully completed, senders receive email notification that their original message has been delivered to its intended recipient—eliminating any question about their ability to correspond with the MailPermit user via email.
  - Reminder notice to senders. If a sender fails to complete the verification step, a reminder email is issued. The notice references the sender's original email correspondence, includes instructions for completing the task and warns of the impending deletion of the original email. Scheduled intervals for reminder notices and automatic message deletion are user-definable.
  - Flexible licensing and use of CR technology. MailPermit administrators may opt to only issue challenges for specific email accounts. This flexibility allows for the exclusion of more sensitive mailboxes from your permission-based email system.
  - Economical number of challenges. VisNetic MailPermit differs from other CR systems because it factors in the results of message categorization or scoring by VisNetic AntiSpam. Only when a message reaches a certain probability that it is SPAM will a challenge be issued.
  - User-based verification of suspect email. Via a simple web interface, users may access email stored in their pending folder. This granular control provides added protection against false-positives—allowing users to see pending email—and shields IT administrators from a stream of employee inquiries about missing email.

How do these technologies work together?

Upon receipt of each email, VisNetic MailPermit verifies that the intended recipient is a valid user on the mail server. If the recipient is not a valid user, the email is rejected. The From address on email sent to valid users is subsequently verified against whitelists and blacklists established in VisNetic MailPermit, starting with lists defined at the agent level, then domain and system level. Mail from whitelisted addresses is forwarded directly to the mail server. Mail from blacklisted addresses is deleted.

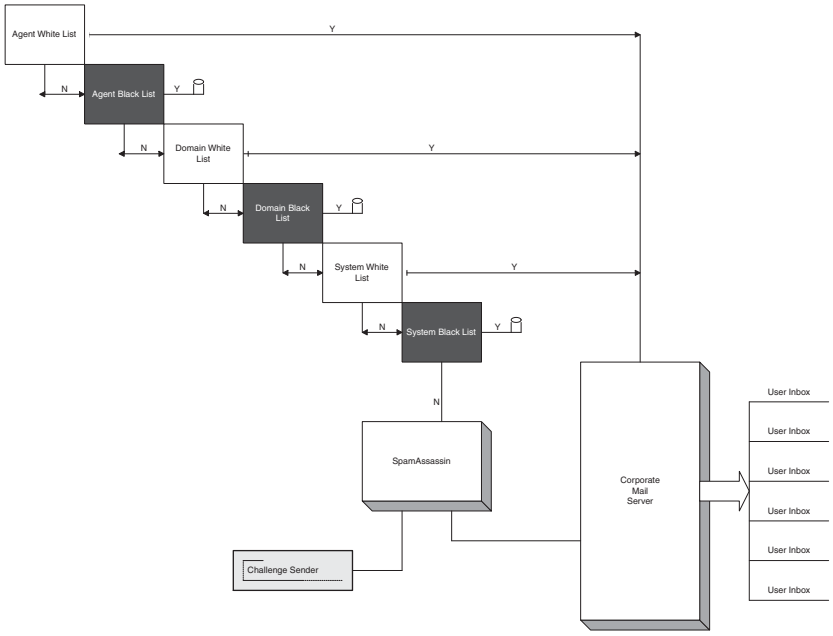

If the From address does not exist on a whitelist or blacklist in VisNetic MailPermit, the email is passed to VisNetic AntiSpam, a rules-based filtering tool powered by SpamAssassin technology. VisNetic AntiSpam assigns a numerical value or score to each message based on the likelihood that the message is spam. Subsequent action is determined by user-definable score thresholds (default thresholds are listed below).

- Email receiving a score of 2 or less has a low likelihood of being spam and is forwarded to the mail server for processing.
- Email receiving a score of 10 or greater has a high likelihood of being spam and is rejected.
- Email scoring between 2 and 10 is passed to VisNetic MailPermit's challenge-response system.

Conclusion

Integration of these methodologies—in the right order—has resulted in an anti-spam solution that is more effective at eliminating spam and reducing false-positives than any other anti-spam solution available today. Furthermore, VisNetic MailPermit is the first server-based anti-spam solution featuring "best of class" rules-based and CR technologies available for internal deployment in SMBs and enterprise organizations.

Licensing and Availability

Initially, VisNetic MailPermit will be available as a plug-in to VisNetic MailServer. A Gateway for SMTP version for use with any SMTP server will be released next. Anticipated public availability for both is August, 2003.

License sizes for the plug-in include a 6-49 user license and a 50-unlimited user license. Gateway licensing will simply include a single domain license and a multiple domain license. Prices start at $99.95 US for the plug-in and $499.95 US for the Gateway.

Contact Information

Deerfield Communications, Inc.
P.O. Box 851. Gaylord, MI 49735

Telephone 989.732.8856
Fax 989.731.9299

www.deerfield.com

sales@deerfield.com
feedback@deerfield.com