INGENISYS

Ingenious Systems Research

MAILVICE SPAM SINCE 2004 FREE POSTMASTER

FORTIFICATION
INVESTIGATION
VERIFICATION
SIMPLIFICATION

Anti-Spam vs. MailVICE

Author: Steven V. Jackson, Principal
Ingenious Systems Research, LLC

In way of discussing the current climate of anti-spam technology, the following uses the format and excerpts from Wikipedia found at http://en.wikipedia.org/wiki/Anti-spam_techniques_(e-mail) on September 25th, 2007. Although it does not cover all attempted or envisioned methods, it is a good starting point for evaluating MailVICE.

End User Approaches

Address Munging

From Wikipedia

Posting anonymously, or with a fake name and address, is one way to avoid "address harvesting," but users should ensure that the fake address is not valid. Users who want to receive legitimate email regarding their posts or Web sites can alter their addresses so humans can figure out but spammers cannot. For instance, joe@example.net might post as joeNOS@PAM.example.net.invalid, or display his email address as an image instead of text. Address munging, however, can cause legitimate replies to be lost. And if it's not the user's valid address, it has to be truly invalid, otherwise someone or some server will still get the spam for it. See http://www.2kevin.net/munging.html

MailVICE Comments

This is a lot of hassle for a user to deal with just to try to hide out from the bad people on the Internet. With MailVICE, it doesn't really matter whether or not you use your REAL email address on the wide open Internet. Send yourself e-greetings all day if you want. It still won't break your inbox.

Avoid Responding to Spam

From Wikipedia

Spammers often regard responses to their messages—even responses like "Don't spam me"—as confirmation that an email address is valid. Likewise, many spam messages contain Web links or addresses which the user is directed to follow to be removed from the spammer's mailing list. In several cases, spam-fighters have tested these links, confirming they do not lead to the recipient address's removal—if anything, they lead to more spam. (Continues on Wikipedia...)

MailVICE Comments

Since the user never sees the spam in the first place, MailVICE keeps responses to spam from happening. This turns out to be vastly beneficial to the provider because it doesn't run a flag up the pole that reads "waste even MORE of our bandwidth".

Contact Forms

From Wikipedia

Contact forms allow users to send email by filling out forms in a web browser. The web server takes the form data, forwarding it to an email address. The user never sees the email address. Contact forms have the drawback that they require a website that supports server side scripts. They are also inconvenient to the message sender as they are not able to use their preferred e-mail client. Finally if the software used to run the contact forms is badly designed they can become spam tools in their own right. Additionally many spammers have taken to using contact forms to send spam to the intended recipient.

MailVICE Comments

When properly implemented, a good website provides forms as a convenience to the user more than spam protection for the target. The complaints that Wikipedia lists are a bit outdated, but the use of contact forms as a spam control method is not necessary with MailVICE protecting the target address. As with any address, whether 10 or 10,000 people send to a popular address such as "sales" or "info", the inbox is still completely usable.

Disable HTML in e-mail

From Wikipedia

Many modern mail programs incorporate Web browser functionality, such as the display of HTML, URLs, and images. This can easily expose the user to offensive images in spam. In addition, spam written in HTML can contain web bugs which allow spammers to see that the e-mail address is valid and that the message has not been caught in spam filters. JavaScript programs can be used to direct the user's Web browser to an advertised page, or to make the spam message difficult to close or delete. Spam messages have contained attacks upon security vulnerabilities in the HTML renderer, using these holes to install spyware. (Some computer viruses are borne by the same mechanisms.)

Mail clients which do not automatically download and display HTML, images or attachments, have fewer risks, as do clients have been configured to not display these by default.

MailVICE Comments

It's a good idea to protect users from the proliferation of viruses by closing up complex holes in web rendering engines as described. But then again, it's also a good idea to forbid a user from browsing the Internet if you want to be completely safe (and completely disabled). Either way, MailVICE stops things like mail-bugs and other malicious email-born devices from freely flowing into the inbox simply because you can't get in if the user doesn't WANT you in.

Disposable e-mail addresses

From Wikipedia

Many email users sometimes need to give an address to a site without complete assurance that the site will not send out spam. One way to mitigate the risk is to provide a *disposable* email address—a temporary address which forwards email to a real account, which the user can disable or abandon. A number of services provide disposable address forwarding. Addresses can be manually disabled, can expire after a given time interval, or can expire after a certain number of messages have been forwarded.

MailVICE Comments

That's a lot of work just to hide from the world. Our take on this is exactly like "Address munging" above - with MailVICE it's pointless.

Reporting Spam

From Wikipedia

Tracking down a spammer's ISP and reporting the offense can lead to the spammer's service being terminated. Unfortunately, it can be difficult to track down the spammer—and while there are some online tools to assist, they are not always accurate. Occasionally, spammers employ their own netblocks. In this case, the abuse contact for the netblock can be the spammer itself and can confirm your address. (Continues on Wikipedia...)

MailVICE Comments

This method is just what employers do NOT want: even MORE time spent on spam. Some home users might have the time to pursue this kind of Don Quixote ventures, but most do not. With MailVICE, the user doesn't even feel inclined to bother.

Responding to Spam

From Wikipedia

Some advocate responding aggressively to spam—in other words, "spamming the spammer". The basic idea is to make spamming less attractive to the spammer, by increasing the spammer's overhead. There are several ways to reach a spammer, but besides the caveats above, it may lead to retaliations by the spammer. (Continues on Wikipedia...)

MailVICE Comments

As with reporting spam, this ultimately proves more bothersome for the recipient than the spammer. Email providers like the same things about MailVICE here as they do with "Avoid responding to spam" above.

Automated techniques for e-mail administrators

Authentication and Reputation

From Wikipedia

A number of systems have been proposed to allow acceptance of email from servers which have authenticated in some fashion as senders of only legitimate email. Many of these systems use the DNS, as do DNSBLs; but rather than being used to list nonconformant sites, the DNS is used to list sites authorized to send email, and (sometimes) to determine the reputation of those sites. Other methods of identifying ham and spam are still used.

Authentication systems cannot detect whether a message is spam. Rather, they allow a site to express trust that an authenticated site will not send spam. Thus, a recipient site may choose to skip expensive spam-filtering methods for messages from authenticated sites.

MailVICE Comments

According to the link associations on Wikipedia, this is where technologies that would like to emulate MailVICE's Forensic Sender Test (FST) such as SPF are classified. We agree that being able to establish a server's right to send mail on behalf of a particular sender is important. In fact, we assert that it is the key to the entire formula. After the spoofing problem has been solved, as Wikipedia suggests, using content filtering becomes irrelevant. The missed point here is that it can't be done on a server or even domain basis to be effective; it must be on a sender to recipient basis. We actually believe that the idea of forcing trust of entire servers is a back door to allowing spam buy-ins.

Challenge/Response Systems

From Wikipedia

Another method which may be used by internet service providers, by specialized services or enterprises to combat spam is to require unknown senders to pass various tests before their messages are delivered. These strategies are termed challenge/response systems or C/R. Some view their use as being as bad as spam since they place the burden of spam fighting on legitimate email senders.

MailVICE Comments

First, the Wikipedia description of the bad side of Challenge/Response is rather anemic. The much bigger problem is that these systems spew bogus challenges all across the Internet. MailVICE doesn't have that problem, however, because of the FST. We never send a challenge to a spoofed sender, which is around 90% of all messages. If we're challenging bob@somewhere.com it's because bob sent an email not because we though a spoof was worth challenging (maybe even confusing Bob into answering the challenge!).

As for the other "problem" mentioned in the Wikipedia article, we don't view it as much of an imposition to issue a sender a one-time request to type six letters compared to the benefit to the recipient. It pretty much qualifies under the "I'd do it for you" clause. But factor in that MailVICE automatically trusts all addresses that have been sent to and the inconvenience factor drops even further.

Checksum-based filtering

From Wikipedia

Checksum-based filter exploits the fact that the messages are sent in bulk, that is that they will be identical with small variations. Checksum-based filters strip out everything that might vary between messages, reduce what remains to a checksum, and look that checksum up in a database which collects the checksums of messages that email recipients consider to be spam (some people have a button on their email client which they can click to nominate a message as being spam); if the checksum is in the database, the message is likely to be spam.

The advantage of this type of filtering is that it lets ordinary users help identify spam, and not just administrators, thus vastly increasing the pool of spam fighters. The disadvantage is that spammers can insert unique invisible gibberish—known as hashbusters—into the middle of each of their messages, thus making each message unique and having a different checksum. This leads to an arms race between the developers of the checksum software and the developers of the spam-generating software.

MailVICE Comments

This type of filtering is simply a more processor intensive version of pattern matching. The "arms race" mentioned actually applies to ANY content filtering method. Here is as good a place as any to mention the ultimate end of such an arms race. The spammers eventually make their messages look more and more legitimate and the filters get more and more aggressive. False-positives become so frequent that email loses all value as a communication medium.

As if to prove this point, one of our own test accounts with Yahoo turned up a message just days ago that was a perfect fake invoice for an online purchase enclosed in a PDF attachment. We have no doubt that the message and the attached PDF were generated for this one transmission from a pallet of pieces giving the message a high degree of statistical uniqueness. If a spam filter were to learn to stop that message, they would've also learned to stop all REAL sales receipts!

Country-based filtering

From Wikipedia

Some e-mail servers expect to never communicate with particular countries from which they receive a great deal of spam. Therefore, they use country-based filtering - a technique that blocks e-mail from certain countries. This technique is based on country of origin determined by the sender's IP address rather than any trait of the sender.

MailVICE Comments

Although this might work well for some corporate mail servers, it's not generally acceptable for an ISP to attempt this method. MailVICE does not directly employ a method like this because one of our design directives is that we will take any email at least as far as the user's "vice". (We do, by the way, realize that the word "vice" should logically be spelled with an "s" in this whimsical context. That, however, would not satisfy our motif. ;-)

DNSBLs

From Wikipedia

DNS-based Blackhole Lists, or DNSBLs, are used for heuristic filtering and blocking. A site publishes lists (typically of IP addresses) via the DNS, in such a way that mail servers can easily be set to reject mail from those sources. There are literally scores of DNSBLs, each of which reflects different policies: some list sites known to emit spam; others list open mail relays or proxies; others list ISPs known to support spam. Other DNS-based anti-spam systems list known good ("white") or bad ("black") IPs domains or URLs, including RHSBLs and URIBLs. For history, details, and examples of DNSBLs, see DNSBL.

MailVICE Comments

MailVICE does employ various DNSBLs, particularly the real-time lists that rely on the use of seed boxes. Unlike most products, if an IP is listed, it doesn't mean that we'll discard the mail. As a matter of fact, if mail from a listed IP is from a trusted sender, the mail is delivered as long as the network source from the FST is valid. The RBL results are basically used to help sort the held mail and further limit the proliferation of bogus challenges.

Enforcing RFC standards

From Wikipedia

Enforcing technical requirements of the Simple Mail Transfer Protocol (SMTP) can be used to block mail coming from systems that do not comply with the RFC standards. A lot of spammers use poorly written software or are unable to comply with the standards because they do not have legitimate control of the computer sending spam (zombie computer). By setting restrictions on the MTA a mail administrator can reduce spam significantly.

MailVICE Comments

We would argue that the claim that it reduces spam significantly may have been true at some point, but as with all other aspects of filtering, spammers have "caught up". MailVICE DOES require reasonably tight compliance with the standards just because it's the right thing to do. We do not, however, enforce obscure deviations from normal practices in order to complete mail transactions. Our objective is, after all, for our users to receive the mail they want just as much as it is to stop the mail they don't want.

HELO/EHLO checking

From Wikipedia

(The Wikipedia description is merely a series of examples by which the HELO/EHLO [opening command of an SMTP transaction] is checked to resolve in a simple DNS A lookup to the IP of the connected machine. Wikipedia claims that this eliminates 25% of all spam.)

MailVICE Comments

Again, Wikipedia's estimates of the effectiveness of this method do not match our own statistics. It is likely as with other methods that spammers have "caught on". Further, Wikipedia doesn't mention that legitimate email is commonly lost using this method due to multi-ring configurations found in many ISP's and secure corporate environments. A server's REAL name could commonly be "mail.companydomain.local" but it's external interface could be delivering mail to the world with no external DNS record of "companydomain.local". MailVICE does not employ or recommend this method.

Greylisting

From Wikipedia

The SMTP allows for temporary rejection of incoming messages. Greylisting is the technique to temporarily reject messages from unknown sender mail servers. A temporary rejection is designated with a 4xx error code that is recognized by all normal MTAs, which then proceed to retry delivery later.

Greylisting is based on the premise that spammers and spambots will not re-try their messages. Instead, they will move on to the next message and next address. This is effective since a re-try attempt means the message and state of the process must be stored inherently increasing the cost incurred by the spammer, but a standard component of any legitimate sender's server.

MailVICE Comments

At least one of the designers of MailVICE finds this method to be the most offensive perversion of email service outside of spam itself. First, the method destroys the timeliness of email creating support issues for the sender's provider and latency for their own recipients. Second, as with all other methods that assume spammers to be sloppy, the effectiveness disappeared very quickly after deployment. The spammers stepped up and now only the legitimate users are penalized. Not only does MailVICE refuse to employ this method, but we encourage anyone who uses the method to stop immediately.

Fake MX Records

From Wikipedia

Virus infected spam bots ignore requirements that the email start at the lowest numbered MX record and move up the list if a failure occurs. They try the highest numbered MX records first thinking that the backup servers have less spam filtering than the low numbered MX servers. Spam bots usually do not retry on failure but move on to the next email address in their list. Thus adding fake high numbered MX records is an effective way to reduce incoming spam. See FakeMX.

One can also reduce spam by having a fake lowest MX record as well. This causes real email to have to retry as well but it only adds a second to the delivery time. Some people report as much as 90% reduction in spam bot spam using this method.

MailVICE Comments

Once again, Wikipedia presents an incredible claim of 90% reduction in spam that doesn't meet the statistics on our servers. This method actually suffers from many of the shortcomings of grey listing. There is even a flip side of this equation. MailVICE relies on the practices that are "standard" in that they are actually employed all over the Internet, not just in an RFC. One of these "standards" seems to be that telco's who are bandwidth providers commonly inject a lowest priority exchanger for email servers on their bandwidth. Should the bandwidth drop, the telco's mail server takes mail on behalf of the bandwidth customer but in many cases, never gives it back. Considering this situation, MailVICE itself doesn't like to roll through dead MX records because it hopes the REAL server will come back online. In short, this method is sloppy and doesn't do much to stop spam anymore.

Greeting Delay

From Wikipedia

A greeting delay is a deliberate pause introduced by an SMTP server before it sends the SMTP greeting banner to the client. The client is supposed to wait until it has received this banner before it sends any data to the server. (per RFC2821 3.1). Many spam-sending applications do not wait to receive this banner, and instead start sending data once the TCP connection is complete. The server can detect this, and drop the connection.

There are some legitimate sites that play "fast and loose" with the SMTP specifications, and may be caught by this mechanism. It also has a tendency to interact badly with sites that perform Callback Verification, as common callback verification systems have timeouts that are much shorter than those mandated by RFC2821 4.5.3.2.

MailVICE Comments

MailVICE doesn't specifically delay the greeting because doing so would simply provide an easier path to denial of service attacks. This is just another method that invites denial of service attacks. We do, however, require that the SMTP transaction be completed in the correct order.

Hybrid Filtering

From Wikipedia

Hybrid filtering, such as is implemented in the open source programs SpamAssassin and Policyd-weight, uses some or all of the various tests for spam, and assigns a numerical score to each test. Each message is scanned for these patterns, and the applicable scores tallied up. If the total is above a fixed value, the message is rejected or flagged as spam. By ensuring that no single spam test by itself can flag a message as spam, the false positive rate can be greatly reduced. [4]

MailVICE Comments

This is a clear route to angry users. Make that threshold just a little too tight and Joe User can't talk to his dad about his health issues making the support desk's phone ring. Make it loose enough for Joe and you'll prompt a dozen more phone calls about too much spam. This method is actually one of the things that caused MailVICE to be created in the first place.

PTR/Reverse DNS checkings

From Wikipedia

(Wikipedia actually fails to describe this method with any clarity.)

MailVICE Comments

The forward/reverse cross check (FRCC) is another example of an Internet standard that doesn't appear in an RFC. AOL made this check a defacto standard by requiring deliveries into their system to pass this test. Their reasoning was sound in that legitimate servers shouldn't mind having their true identities discovered. MailVICE's FST uses this test as a prerequisite because without passing the test, the results could be forged.

Rule-based filtering

From Wikipedia

Content filtering techniques relied on the specification of lists of words or regular expressions disallowed in mail messages. Thus, if a site receives spam advertising "herbal Viagra", the administrator might place these words in the filter configuration. The mail server would thence reject any message containing the phrase. (Continues on Wikipedia...)

MailVICE Comments

Again, this method is just asking for support problems on a small scale and the annihilation of email as a communication tool on a global scale. The ultimate logical end to better filters is spam that looks more like real email. The ultimate end to spam that closely resembles real email is unreliable email.

Sender-supported whitelists and tags

From Wikipedia

There are a small number of organizations which offer IP whitelisting and/or licensed tags that can be placed in email (for a fee) to assure recipients' systems that the messages thus tagged are not spam. This system relies on legal enforcement of the tag. The intent is for email administrators to whitelist messages bearing the licensed tag.

MailVICE Comments

Anyone but us smell a scam? This is the ultimate in shameless profiteering not only selling snake oil, but peddling a system with holes that can be sold to the highest bidder.

SMTP callback verification

From Wikipedia

Since a large percentage of spam has forged and invalid sender ("from") addresses, some spam can be detected by checking that this "from" address is valid. A mail server can try to verify the sender address by making an SMTP connection back to the mail exchanger for the address, as if it was creating a bounce, but stopping just before any e-mail is sent.

Callback verification can be compliant with SMTP RFCs, but it has various drawbacks. Since nearly all spam has forged return addresses, nearly all callbacks are to innocent third party mail servers that are unrelated to the spam. At the same time, there will be numerous false negatives due to spammers abusing real addresses and some false positives.

MailVICE Comments

This method is absolutely useless without solid spoof detection. All a spammer has to do is use the last successful delivery address as the sending address of its very next spam! Even though MailVICE has air-tight spoof detection with our FST, we didn't see that burdening the sending server of a legitimate sender as "appropriate" since it's almost like penalizing the good guy with wasted bandwidth.

Statistical content filtering

From Wikipedia

Statistical filtering was first proposed in 1998 by Mehran Sahami et al., at the AAAI-98 Workshop on Learning for Text Categorization. A statistical filter is a kind of document classification system, and a number of machine learning researchers have turned their attention to the problem. Statistical filtering was popularized by Paul Graham's influential

2002 article A Plan for Spam, which proposed the use of naive Bayes classifiers to predict whether messages are spam or not – based on collections of spam and nonspam ("ham") email submitted by users. [7] (Continues on Wikipedia...)

MailVICE Comments

This is just a super-charged take on plain old pattern matching with all the same uselessness. No matter how "bleeding edge" you make reading an email, it can not possibly beat the best computer on Earth: the human mind. Otherwise stated, let's say that you hired a real human being to read every single email you received and filter out the junk mail. They couldn't get it right unless they knew every last nook and cranny of your life. How does even a real live person know whether or not you've applied for a mortgage or ordered some online pharmaceuticals? The inbox is a much quieter place once any such invasion of privacy is abandoned in favor of simply concentrating on trust and identity.

Tarpits

From Wikipedia

A tarpit is any server software which intentionally responds pathologically slowly to client commands. This is not to be confused with slow speeds sometimes experienced by Telecoms clients. By running a tarpit which treats acceptable mail normally and known spam slowly or which appears to be an open mail relay, a site can slow down the rate at which spammers can inject messages into the mail facility. Many systems will simply disconnect if the server doesn't respond quickly, which will eliminate the spam. However, a few legitimate e-mail systems will also not deal correctly with these delays.

MailVICE Comments

This method is another gimmick that just doesn't do much in the real world. It seems to assume that a spammer would give up just because things are going slowly which is absolutely ridiculous. A throttled connection is wide open for the spammer to make other connections. Practices like this also make the receiving server prone to denial of service attacks through connection flooding as well. Needless to say, this is NOT a MailVICE practice. MailVICE takes mail as fast as it possibly can and it doesn't matter to the end user: If they don't want to talk to the sender, it won't happen.

Transparent SMTP proxy

From Wikipedia

Transparent SMTP proxies allow combating spam in real time, combining sender's behavior controls, providing legitimate users immediate feedback, eliminating a need for quarantine.

MailVICE Comments

Wikipedia seems to be pretty confused about how a "transparent SMTP proxy" might be useful in spam fighting because it mentions things like lack of a quarantine which really has nothing to do with the use of a proxy. Let's say for a minute that they're talking more about using an application layer firewall to protect the real mail server. If that IS what they're talking about, then MailVICE does fit the description. As a matter of fact, one of the reasons MailVICE was created in the beginning was more for the purpose of mail server protection than spam fighting. About 97% of the usual burden on a mail server disappears with MailVICE as a firewall.

X-ASVP eXtensible Anti-spam Verification Protocol

From Wikipedia

X-ASVP is a paper design of a method to enable peer to peer scalable authentication between holders of an Internet e-mail address, with redundancy and reliability provided by secondary (tld) and tertiary (global) providers. Its proponents claim that a prototype implementation exists.

MailVICE Comments

This method is just proof that you can post just about anything on Wikipedia. A quick glance at the base article reveals yet another scheme to redesign the entire Internet. MailVICE is 100% effective without changing anything about the way email works.

Many of the methods listed by Wikipedia act as temporary defenses by exploiting the spammer's use of "whatever works". Although we don't really agree with any of these methods, they were marginally effective as a solution-of-the-day. Administrators that continue to use these methods as if the spammer wouldn't adapt are losing sight of the bigger picture of spam. Spam is an industry that makes incredible amounts of money, so much money that it's even difficult to estimate. With all that money there for the taking, why do administrators not expect the spammers to improve and overcome any shortcomings in their methods quickly? Spam as an industry has a bigger research and development budget than most of these administrators' companies.

Although mentioned in bits and pieces above, an interesting aspect of the MailVICE system is that there is no way for a spammer to defeat the system. No matter how well they conform to standards, no matter how creative they get with their content, and ultimately no matter how "legitimate" they try to become, all they really accomplish is giving MailVICE more information to pass on to the user. There is no point at which MailVICE has to change its methods. The exact same methods have been in use since August of 2004 and there is no conceptual point at which these methods will need to change.

It is generally agreed in the industry that validation of sender identity is the king pin of spam control. The reason for this is not that it creates a way to transparently filter based on legitimacy. Simply throwing away spoofs would do nothing more than turn the huge research and development budget of the spammers toward finding ways to send without spoofing that leaves law enforcement knocking on the wrong door. And they would find it. But add just one piece to this equation and it becomes an impenetrable defense: a registry of acceptable senders. The Forensic Sender Test provides the validation and the rest of MailVICE provides the enforcement of the registry.

If SenderID or Domain Keys were ever globally implemented, solid systems like MailVICE would be everywhere and spam would cease to exist. And as soon as every mail administrator on Earth jumps in and devotes his or her part of the thousands upon thousands of labor hours required to participate in these new systems, everything will be fine. Or, more realistically, they could just get MailVICE and be done with it.