INGENISYS

Ingenious Systems Research

Sender Network Verification without Global Modification

Breakthrough invention renders traditional anti-spoofing technologies obsolete.

Author: Shiras Michael Walker, Jr., Principal
Ingenious Systems Research, LLC

Can you imagine being able to remove yourself from the spam war while the rest of the world battles it out? How about going years without ever seeing a single piece of spam and only hearing the horror stories about the latest anti-spam countermeasure wreaking havoc across the internet? Can you imagine how you would feel knowing that you couldn't care less about it?

Receiving millions of spam messages but only delivering thousands of real messages *is* like fighting a war, and up until just a few years ago there was really no way to avoid the fight. Let's face it; spammers are not going away when proposed solutions to the problem aren't really "solutions" at all but mere empty promises (the "almost fixes it" 99%'s), utter hyperbole (the "100% effective" but not's), or hopeful thoughts ("by the time we really fix it your grandkids will thank us for it").

Ironically, what we are describing in this document is **not** an "anti-spam" service, but because of what we have invented, you can finally declare victory and withdrawal yourself from the seemingly never ending spam battles. Being "pro-email" rather than "anti-spam", MailVICE stops *your* war against spam, phishing attacks, and other inbox intrusions by adding the missing component to the email equation: **sender network verification *without* global modification**. This technological breakthrough removes the burden on end users and system administrators alike and returns email to the inexpensive and powerful communications medium it was intended to be – without magic tricks, smoke and mirrors, or flawed anti-spoofing protocols.

Want to read about how a couple of guys from East Tennessee beat Microsoft, Google, Yahoo, IBM, Symantec, AOL, Barracuda Networks, Postini, MailFoundry and every other anti-spam initiative at defeating spammers?

Simply turn the page!

The Anti-Spam Research Group of the Internet Research Task Force is attempting to develop a technology that merges the best attributes of several competing email anti-spoofing technologies such as DKIM and SPF. These technologies are designed to combat email spoofing by modifying the global DNS infrastructure. Simple as the changes might look on the surface, modifying millions of email domains and thousands of Internet service provider's networks to support these changes is a daunting task. These technologies require changes to be made to *every* domain and ISP in order to be useful, and seems about as impossible as changing the entire underlying email protocol.

The question these technologies are trying to answer is a simple one: Are the senders who they say they are? The answer, however, is not so simple. If a receiving mail server could somehow prove that an incoming email originated from where it said it did, then this would raise the probability that the message is valid. To provide this Information, the proposed technology would need to know what mail servers are allowed to send email for each domain on the Internet. Past attempts to glean the answer from current DNS data have failed, so this proposed technology plans to attempt a retrofit of new data requiring a massive initial adoption and ongoing maintenance of additional DNS information that has never existed before. The proposal's solution to spoofing requires changing the world.

On the other hand, what if the answer to the question already exists without the need to change anything? If a mail server alone could answer the same question and get the same answer as the proposed technology could only provide in theory, then why try to change the world to get *less*?

The core engine within MailVICE contains anti-spoofing technology that rivals the unattainable final goal of all current and proposed anti-spoofing technologies. Whereas the proposed mechanism requires a global modification to one of the Internet's core subsystems, MailVICE uses readily available data to prove whether or not a sending server has the right to send email on behalf of a domain. MailVICE has the ability, *today*, to properly identify almost any domain's legitimate mail servers without requiring a global modification to any Internet subsystem.

The secret behind the MailVICE core is a series of complex queries that dig up the required information in a very thorough and methodical investigative manner. Even though this process is computationally intensive, the globally distributed Forensic Sender Test (**FST**) has the capability of processing this information for every email that traverses the Internet, *in real-time*.

Proving a server has the right to send email for a domain is just the first step in a comprehensive anti-spam solution. Providing the answer to the ultimate question - *Is this email a spam?* –requires sender/recipient trust to be established. The same proven security used in most Instant Messaging protocols is missing from SMTP, requiring an additional authentication layer be added on top of it. Most "trusted sender"-type mechanisms are prone to spoofing which is where MailVICE's core **FST** technology comes into play. Not only does MailVICE establish a trust relationship between the sender and recipient, but it adds the sender's true sending network to the trust via **FST**, creating a sender/sending network "verified" pair.

Using the core **FST** sender verification technology in multiple ways yields **100%** (yes, one hundred percent) effective spam protection. On one hand, **FST** is verifying the sender's legitimate email servers and using that information to prove email legitimacy, and on the other hand MailVICE is proving whether or not the sender has the permission required to deliver email to the MailVICE user's inbox. So what happens when a long lost buddy from yesteryear tries to send an email to a MailVICE user?

Another interesting exploitation of the **FST** technology is the ability to use a unique challenge/response system that sends unverified senders verification messages only if the incoming message is not spoofed. This prevents MailVICE from sending millions of verification emails to often bogus FROM addresses, wasting global bandwidth and mail server resources. Moreover, MailVICE is able to determine human from robot effectively and only sends verification messages to people who send legitimate email. The long lost buddy sends the MailVICE user an email, verifies himself, instantly becomes a trusted sender, and the email is ultimately delivered - all without MailVICE user intervention.

Unlike most other anti-spam services, the MailVICE user experience is designed to wean the user from having to use it. Part of the experience includes the automatic addition of verified senders when a MailVICE user sends email. MailVICE quickly learns who a user wishes to receive email from, and within a few days most users rarely interact with MailVICE at all. This unique feature is possible because MailVICE is logically located in between the user's real mail server and the internet. As outbound email passes from their mail server, through the MailVICE network, and out to its final destination, the TO address in the message is added as a verified sender. This significantly cuts down on user administration since MailVICE safely assumes the trust. Moreover, MailVICE will not only trust the recipient's email address, but will only accept email from that recipient if they send email from a valid network (must pass the **FST** check). To pass the **FST** check, the sender must be sending his email from a network that "has permission" to send email on behalf of his own domain. A MailVICE user may also trust email addresses sending from networks that do not have inherent permission to send on their domain's behalf, however, MailVICE will only deliver email from that person if they send from *that* invalid network, verified via **FST** check.

Ever receive an email from a friend only to find it to be spam *not* sent by that friend? When using MailVICE, people you "trust" won't be appearing to spam you. This is a perfect example of **FST** at work. MailVICE not only knows your friend's email address and knows you want to receive email from him, but also knows the networks your friend is allowed to send you email from. This capability (promised, but yet to be delivered by current anti-spoofing protocols) is what sets MailVICE apart from every other anti-spam technology on the planet.

Because of the logical location of MailVICE in the email circuit, it becomes very easy to capture diagnostic transport layer information that is often times very difficult to manage. As email traverses MailVICE, all transport information is stored with each transaction and is available to both administrator and user. Since MailVICE is the MTA it's in the perfect position for the recording of SMTP transaction data for both inbound and outbound email. Within a couple of mouse clicks the SMTP transaction log can be located for a specific inbound or outbound email; granular logs in seconds.

Being in the "line of fire" and publically exposed to the world, the MailVICE MTA has been designed to take a beating. The MailVICE MTA is a hardened and scalable network application that is easily distributed across thousands of servers utilizing a multi-threaded process that takes full advantage of SMP and threaded I/O. Part of the MTA's job is protecting the real mail servers behind it by acting as an application layer firewall. When combined with layer 3 and 4 firewalling, the MailVICE MTA is designed to withstand multi-gigabit brute force DoS attacks without losing any email. MailVICE was not designed to stop the intake of any email message addressed to a valid user. All inbound email is delivered to either the user's inbox or their "vice", a temporary holding facility where all unverified email is stored. Either way, MailVICE never impedes the intake of any email, and users always have access to their past undelivered email.

MailVICE queues every message sent to any MailVICE user and either delivers the message or stores the message for future retrieval until it expires. Only messages that match the verified pair entries within the user's Verified Senders Directory are automatically delivered. This" closed by default" methodology, coupled with rock solid sender network

verification, allows only legitimate trusted senders to deliver straight into MailVICE users' inboxes. Everything else is stored for the user for a customized length of time, without the need for an inbox-like anti-spam quarantine.

Administration of MailVICE couldn't be any easier. Mail server administrators simply add user accounts to their mail servers. Upon receipt of email addressed to that user or that user logging into the MailVICE interface, the MailVICE account is automatically created. If a user does not log into the interface, advanced MailVICE protection is disabled by default. However, once logged in, full MailVICE protection is enabled after verifying the account using SMTP Auth against the back-end mail server. Periodically, MailVICE will verify the email account still exists by performing additional SMTP Auth queries against the back-end mail server. After three unsuccessful SMTP Auth queries, the MailVICE account is automatically deleted. Administrators also have the ability to control MailVICE behavior and view performance and utilization statistics through a customized interface.

Possibly the strongest MailVICE features are the technologies it doesn't use. Notice that MailVICE does not rely on things such as Bayesian filtering, for instance. In fact, MailVICE does not use any technology that passes judgment on the contents of an email message, or the reputation of the server that sent it. Moreover, MailVICE does not look inside the email envelope whatsoever. Whether the email message contains the word "pharmacy", or the message contains improperly formatted HTML means nothing to MailVICE. The user has the power to decide who they want to accept email from and it is none of MailVICE's business to care what is actually being received. This eliminates the huge risk of losing legitimate inbound email from false positive filtering matches, not to mention the **privacy** aspects of not needing to parse through personal communications. Even if a MailVICE user's verified sender was sending email from a network on 20 spam blacklists, from a server with no reverse DNS, and contained a list of every "spammy" word and phrase imaginable, the message would still be delivered *every time* as long as the sender passed an **FST** check.

The MailVICE system has prevented millions of spam messages from reaching thousands of its user's inboxes. To date, not a single MailVICE user has reported an actual spam delivery or lost a single legitimate email since MailVICE v1.0 went live in August of 2004. Interestingly, MailVICE itself has not been significantly modified since then either, and has withstood *every* anti-spam countermeasure without incident and remains impervious to those countermeasures without reconfiguration or modification in any way. Spammers just can't change how the Internet works which is required to defeat MailVICE. As long as email is delivered using SMTP, there is nothing a spammer can do to get past it.

With MailVICE, there is no reason to change how the Internet works. The spammers can't do it, and the anti-spammers shouldn't have to do it. Until SMTP is replaced with something that has sender network verification built into the protocol, MailVICE is available to provide that functionality today without the need for global change.

Not a single MailVICE user received a spam today, but we know you did.