Top Ten Worst Ways To Stop Spam

The Best Ways to Ensure the Worst Protection

What is an Anti-Spam Solution?

Implementing an anti-spam solution is not unlike ordering a meal at a restaurant. Few would order a hamburger without the fixings, buns, and fries. Yet, with so many options to choose from in the anti-spam realm, it's hard to distinguish between technologies that will be the whole “meat and potatoes,” and those that will leave you with only a soggy bun.

This white paper addresses technologies that, in and of themselves, are insufficient for comprehensive email security. While many of these technologies are good, if not excellent, when used in conjunction with other technologies, they should never, ever, be used independently.

Number Ten – The Firewall

Numerous firewalls contain some form of anti-spam technology, pre-integrated and ready to run on your network. While firewalls sit at the edge of your network and prevent network-borne attacks from reeking havoc on your intranet, it seems quite intuitive to deploy an anti-spam solution at this level. In fact, it isn't a bad idea at all. Take note that different firewalls offer different qualities of anti-spam technology, ranging from none at all, to pathetic, to actually quite good. One firewall worth mentioning is the SpamBlocker technology, which comes optional on WatchGuard's Firebox X-Series firewalls. The Firebox uses technology from Commtouch, an absolutely brilliant technology that is both extremely fast and highly accurate.

So why did we rank firewalls as number ten on this list? Two reasons:

- First, most firewalls are not very good at catching spam. Those that are can provide great border protection, but they still leave something to be desired.
- Second, filtering spam requires horsepower and firewalls just don't have the guns to fight spam comprehensively. Most organizations desire more from a complete email security solution such as user-level controls over email security preferences, quarantines, and a large slew of other bells-and-whistles.

To put it simply, firewalls are generally too bare-boned and under-powered to provide comprehensive email security.

Number Nine – Desktop Solutions

A non-negotiable rule of thumb: If your organization has ten or more users, a desktop solution is a horrible choice. Like all flavors of anti-spam solutions, some desktop solutions are good, some not. Nonetheless, in organizations of over nine users, the complexities wrought by maintaining the software on individual desktops makes for a real money-sucker. According to Ferris Research, organizations that employ desktop solutions will end up spending about 65% more per user in software and maintenance costs compared to those who do not use a desktop solution.

Number Eight – Challenge-Response

Challenge-response is actually one of the most accurate ways to stop spam. It relies on the fact that machines are stupid and humans are smart. When a message is sent to a user which is protected by a challenge-response system, the system sends a message back, typically asking the sender to perform a task such as type in a few letters that are presented. Spammers would be unable to manually address these millions of challenges, making spamming unprofitable.

However, challenge-response doesn't relieve the burden of dealing with spam. It simply shifts it from recipient to sender. Suddenly the individual must manually make sure each of his messages arrive to the intended recipients. And what about legitimate bulk senders? Imagine the headaches of the legitimate bulk senders having to manually answer to hundreds or thousands of challenges. It's just not feasible.

Number Seven – SpamAssassin

SpamAssassin is a well known and well respected open-source package supported by the Apache group. It's especially popular among ISPs and web hosts (mainly because it's free). The shortcomings of SpamAssassin are a direct result of its own benefits. Because it's free, and widely used, spammers run their messages through it to determine if they can sneak past those millions of SpamAssassin implementations. The end result is unacceptably low capture rates.

While SpamAssassin can be tweaked to be tougher on spam, it also winds up being tougher on good messages too, causing legitimate messages to be quarantined or blocked. For individuals, it's not the end of the world if you lose a message from grandma. But in the business world, when you lose a million-dollar order from a customer, it's cause for some serious frustration.

Number Six – Pattern Matching

Techniques that look for recurring patterns throughout a message face an uphill battle. Pattern matching commonly utilizes expressions designed to look for patterns that are commonly occurring in current spam messages. It is well known that spammers will attempt to obfuscate messages (ex: \/!@gra f0r 1e$$), and most pattern matching engines can compensate quite well for this. However, with billions of diverse spam messages surging throughout the internet, creating and maintaining a database of up-to-date patterns has proven too daunting a task for highly effective spam filtering. Pattern matching also can pose a significant risk of false-positives in the event patterns are not carefully created.

Number Five – Bayesian Filters

In late 2002, Paul Graham wrote “A Plan For Spam”. This essay illustrated how he had successfully implemented a form of machine learning based upon the mathematical concepts of Thomas Bayes. Using Bayesian analysis, Paul was able to eliminate most spam from his inbox. After publishing his essay, many individuals and software engineers used Paul's tips to develop their own successful filters. Many, if not most, commercial and free anti-spam packages still include some flavor of Bayesian analysis.

While Bayesian filtering is still an important part of today's anti-spam technology, left alone it's no longer effective. Since it relies 100% on the words contained in a message, Bayesian filters can be foiled by injecting random words and phrases, or by using no words at all, and letting an image do the talking.

Number Four – Checksum Signatures

Solutions such as Vipul's Razor take a snap-shot signature of an email message and compares that signature against a clearing house of known spam messages. These signatures are engineered to compensate for the inevitable randomized text inserted by many professional spammers. The end result is a technology which is highly effective at recognizing known spam messages. The deficiency, as you may have already guessed, is that new spam messages can slip right past checksum-based filters, until the message has

been recognized as spam and a signature created. This leaves hours of time, which translates into millions of messages, before the checksum approach is effective.

Number Three - Greylisting

This technique, affectionately known as greylisting, challenges the RFC compliancy of a sender's server. In most cases, when a message is sent from an unknown server, the greylisting technology will reject the message, stating “try again later” to the sender. RFC compliant servers can recognize this command, and will typically retry in about 15 minutes, at which time their message will be accepted.

For efficiency, many spammers use the “fire and forget” methodology. The logic to handle a “try again later” request either doesn't exist or is disabled on many of the scripts and programs which spammers use. Thus, when a spamming computer is told to try again, it simply ignores the request, doesn't retry, and its message is not received.

Greylisting can quickly stop around 80% of spam messages with little effort. Moreover, it works almost flawlessly with all genuine senders. Nonetheless, an 80% capture rate is not adequate. Additionally, as more and more servers implement greylisting, more and more spammers will deploy logic to handle the retry request. It should, therefore, never be considered a “solution” for spam.

Number Two – Real-Time Blackhole Lists / URI Blackhole Lists

RBLs and URIBLs, simply put, are a failed technology. RBLs check the IP address of the sender against a database of known spamming machines. There are dozens of these lists, each offering their own system for identifying good and bad senders. Some meticulously maintain their blackhole lists, while others simply block huge chunks of IP addresses based on the originating web host, ISP, or even country. URIBLs check the URIs (ex: www.domain.com) against a database of known websites which have been included in spam messages.

To make a long story short, many RBLs and URIBLs operators are notoriously slow to update their records, resulting in very low capture rates and dangerously high false-positives. RBLs and URIBLs can, and are, commonly used to help score a message, but they should never be used as a lone decision maker.

Number One – White / Black Lists

Enter the most ancient form of anti-spam. When the Internet and spam was in its infancy, the handful of email users would manage the handful of spammers, by keeping a handful of addresses in their blacklist. Nowadays, spammers almost always spoof email addresses, making a black list nearly 100% ineffective. Moreover, whitelists can create a hole in which spammers can exploit by spoofing their mail as a whitelisted address, allowing it to flow freely into your inbox.

Conclusion

There you have it. You’ve been equipped with the know-how to slash through the weeds of the anti-spam jungle. Now that you know what to avoid, how do you know what to embrace? Unfortunately, even in the corporate anti-spam world, there are many more talkers than doers. Therefore, make a list of what you expect, find vendors that “promise” to meet your expectations (there will be many), and evaluate, evaluate, evaluate. Don’t settle for an OK solution – tomorrow it will be unacceptable… and may even earn its place on our *Top Ten Worst* list.