Adding SPF
To Your Anti-Spam Solution

Introduction

The Sender Policy Framework (SPF) is an open standard that introduces a grammar for domains to describe the email they send. These descriptions are published in the Domain Name System (DNS), and Simple Mail Transfer Protocol (SMTP) receivers may use these descriptions to authenticate email and to better apply local policy. SPF also defines a designated sender scheme that domains can use to describe the hosts they send email through and, because SPF is extensible, other schemes can be added to the language.

What SPF Does

The primary aim of SPF is to reduce send-address forgery. SPF scales well, distributing the burden of whitelisting mail servers across the space of DNS domains. All domains that originate email should consider SPF. Although the first draft for the proposal of SPF standards was published in December 2003 (titled "Sender Permitted From – a Convention to Describe Hosts Authorized to Send SMTP Traffic"), not many domains have yet implemented SPF.

About 30% of spam today uses spoofed email addresses. This means that someone at this moment could be receiving an email appearing to be coming from my email address, even though I never sent that email. The most common evidence of this is when we receive an email from ourselves, and we wonder, of course, how that could be possible.

The fact is that the SMTP protocol was not written with spammers in mind, but rather it tried mimicking the behavior we adopt when we send paper mail. We can very easily write a sender address on the envelope but sign ourselves in a different way on the letter; neither of those has to really be us. About the only way a recipient can know that it might not be us sending that letter is to look at the stamp on the envelope and see that the letter was posted halfway across the world. Email works the same way. The envelope sender is not a certainty – the apparent sender even less a certainty. Spammers exploit and abuse this weakness of the SMTP protocol because many people who are unaware of this fact create a whitelist of their own email address, or even of their entire email domain.

The SPF record was created to help guard against this situation. In the SPF record, you can declare that your emails can come only from a specific set of IP addresses. Most commonly, you will declare that your emails are coming from your mail exchanger (MX) records, which specify how email is to be routed with the SMTP. But that is not the only option, simply because that is not always the case. Nevertheless, you have the ability to tell the world that your emails will come only from the set of IP addresses or subnetworks that you specify in your SPF record.

If the receiving anti-spam system is configured to check for SPF records, it will know immediately if an email is not legitimate because it is not coming from any of the authorized IP addresses. Your own anti-spam system will also benefit because now you can scan incoming emails that are using your domain as a sender and immediately reject them when the IP address is not one of those expected, without fear of rejecting a legitimate email from one of your colleagues.

The VPN Factor

Assuring that this SPF procedure works effectively means that you do not allow your users to go home and send emails with your company domain as a sender, using their Internet service provider's mail relay. Those users should access your company's virtual private network (VPN) and relay their company emails from the company's mail server. This is best practices as well and, combined with the SPF record, will help recipients distinguish your legitimate emails from spam using your spoofed domain.

Other Protective Measures

There are more and newer ways to protect your domain and ensure that recipients know your emails are legitimate. Some advanced anti-spam systems have a broad range of functions using many different techniques that – taken together – provide in-depth defense in a single gateway appliance.

One such function is email envelope pre-scanning, a relatively new technology that makes a sound judgment on whether an email is from a spammer or not – without the need to actually download and scan the email itself. Pre-scanning saves significant amounts of Internet bandwidth because messages can be rejected *before* they're even downloaded. It also saves both CPU and disk utilization because emails known to be spam are rejected, without having to pass it through a very resource-intensive scan process. Moreover, it enables anti-spam systems that have it to refuse acceptance, so the sender is then responsible for non-delivery notifications, thus massively reducing, or even eliminating, the company's non-delivery notification queues.

Another available tool is a mail portal system, which allows end users in organizations with SMTP email servers to have direct control of their quarantined emails. If a user sees an email in the mail portal report that has been incorrectly blocked as spam, it's a simple matter to have that email released and the sender become whitelisted.

Conclusion

There is a significant upward trend in the number of spam emails being received by companies and organizations around the world; in some extreme cases, 98% of their emails are either spam or contain viruses. For a large organization, this can represent several hundred thousand unwanted emails arriving each day.

Having a comprehensive, multi-function anti-spam system is essential these days, and implementing the forgery-fighting properties of SPF adds to an already powerful synergy – providing another valuable layer to the overall solution.