

Spotting Spoofed Spam

Like viruses, spam is constantly changing, and this requires spam filtering technology to be continuously updated to combat these new threats. Depending on the environment, spam filters can have difficulty with spam claiming to be from the recipient's own domain. This paper describes best practices for email system configuration that, in combination with Netbox Blue's advanced spam filtering capability, can eliminate this type of spam.

Background

Due to the long and complicated history of email, mailservers typically do not check if the mailserver that delivered the mail corresponds to the domain in the sender's email address. It has become fairly standard practice to send email using the mailservers of your ISP, even if your email address is provided by a third party. For example, Alice works for Bobcorp. The email client on her iPhone is configured to receive emails via IMAP from Bobcorp's IMAP server, but it sends outgoing messages via the mobile phone carrier's SMTP server.

Spammers routinely take advantage of this situation by "spoofing" the sending email address – traditionally they have used invalid email addresses, but Sender Address Verification (SAV) makes light work of this. Spammers are now evolving their technique by using a variety of real sender addresses. In this document we consider the kind of spam which claims to come from the recipient's own email address, or from another email address at their domain.

The Netbox can eliminate this type of spam altogether, but it does require some configuration. A far cleaner and easier to administer solution is to send all email from your domain via the Netbox SMTP server. This is probably already happening for users on your managed network, but mobile devices and other email clients outside the managed network must be configured to use the Netbox SMTP server for outgoing email, with SSL (or TLS) encryption and SMTP authentication. All modern email clients support these features and are straightforward to configure.

This means you can use SPF records or content filtering rules on the Netbox to block email claiming to be from your domain unless it comes from a trusted mail client – this could include the Netbox SMTP server, your company Exchange server and maybe a limited number of other servers such as company web servers. It also includes users on the road with a laptop, tablet or smartphone if they are not already using the company's internal servers.

Sender Policy Framework

Sender Policy Framework (SPF) is a system specified in the IETF document RFC4408 that is designed to address the problem of spoofing. It is currently an experimental specification, but the proposal has received widespread support and has been implemented by a large number of software and hardware vendors. It works by setting policies for which mailservers may deliver email from a particular domain. Early implementations of SPF stored these policies in TXT records in the domain name system (DNS), but SPF has now been allocated its own DNS record type – providing both types of records is recommended for backwards compatibility.

A simple SPF record for example.com might look like this:

```
v=spf1 mx a:example.safenetbox.biz ip4:123.45.67.89 -all
```

v=spf1 just says that this is an SPF version one record; mx means allow email from all of the domain's mailservers (as specified in their MX records); a:example.safenetbox.biz means allow email from any IP address that the domain name example.safenetbox.biz resolves to (as specified in the DNS A records); ip4:123.45.67.89 means allow messages to be sent from the mailserver with IP address 123.45.67.89. and all means that all other messages should be denied.

SPF allows much more flexible configurations, including policies based on IP addresses, DNS A records (hostnames), MX records and more. For more information on how to configure SPF see <http://www.openspf.net/>.

SPF and the Netbox

How does SPF help with the problem of spoofed spam? By default the Netbox actually ignores the SPF policy. This is because SPF is still experimental, and a number of sites have misconfigured SPF records – which could lead to messages falsely being blocked as spam. Since the guiding principle behind the Netbox’s spam filter is to avoid false positives, Netbox Blue advises caution when SPF filtering for all domains. As SPF deployment increases we will monitor the situation, and we hope to be able to change our default filters to use SPF for all domains at some point in the future. For now, we recommend you at least set up an SPF record for each of your domains, and set the Netbox to use SPF for local domains only. Once you have configured your SPF records properly, this will not cause any email to be

blocked incorrectly, but will instantly stop any spoofed spam messages arriving at the Netbox claiming to come from your own domain.


SPF has additional benefits – if someone spoofs your domain and sends out spam to third parties using your email addresses, blacklists are less likely to include you if they can check the SPF record and see that the messages are spoofed. It will also prevent such forged messages from being delivered to early adopters of SPF, and in time, regular users.

Configuring SPF records is beyond the scope of standard Netbox Blue support packages, but we would be happy to consult on this subject outside our standard Netbox support services.

Alternative approaches

SPF is an important part of the future of spam control, and the sooner organisations implement SPF records the better. If SPF is really not an option, and you are being plagued by spam spoofed from your own domain, you can use the Netbox’s content filtering module to block spoofed spam. In order to do this without blocking legitimate email from mobile and remote users, you must still configure mobile and remote devices to use the Netbox SMTP server for outgoing email, with SSL (or TLS) encryption and SMTP authentication. Once this is done, you can create a content filtering rule to block emails from your domain unless they come from “trusted senders” – i.e. users on your managed network, or users who have authenticated with the SMTP server.

To do this log into the Netbox, go to *Content Scanning* → *Rules* and click *Add New Rule*. Give the rule a name such as “Spoofed Spam”, and choose an action (such as “unwanted content”) to take when the rule triggers. Press update, then click “Add new criteria”. Choose “Sender details match”, and for each email domain configured on the Netbox enter `*@domain.name` – one domain per line. (where domain.name is replaced with the real domain name such as `example.com` or `example.com.au`). Press update, then click *Back to Rule*. Press “add criteria” and this time choose “Client IP”. Tick the “NOT” box, and choose “Any trusted client” from the drop-down list box. When you are done the rule should look like this:



Copyright © 2000-2011 Netbox Blue Pty Ltd®

A final warning

Both of these approaches rely on all outgoing email from your domain going through the Netbox, or other trusted mailservers. Before implementing these techniques please ensure that all users, including remote users and mobile devices are configured to use the Netbox as their SMTP server, otherwise legitimate email may be lost.

Netbox Blue is a leading provider of internet and email security, filtering and management solutions. Netbox Blue provides schools and organisations with the tools to protect their network from internal and external threats, control data leakage and ensure staff and students use the internet productively. The company offers a broad portfolio of products and services including Unified Threat Management appliances, email filtering appliances, soft appliances (for virtual environments) and OEM-ready solutions. The company was established in 1999, is privately held and is based in Australia. Netbox Blue has a presence in more than twenty countries and has partnerships and distribution agreements with some of the world’s largest IT providers.