Monthly report

September 2004

Sender Policy Framework (SPF)

Statistics for September

MessageLabs currently scans over 70 million emails per day on behalf of its clients.

In September, MessageLabs scanned more than 1.45 billion emails worldwide for spam, of which over 1.05 billion or 72.14% (1 in 1.39), were stopped as spam (404.68 per second).

During the same period, we also scanned over 1.78 billion emails for viruses, Trojans and other malicious content, and more than 86 million or 4.83% (or 1 in 20.69) were intercepted (33.27 per second).*

The changing face of spam with SPF

As the spoofing of email addresses remains a widespread and growing problem, efforts to strengthen the existing protocols will continue unabated. Earlier this month, the Internet Engineering Task Force disbanded its MTA Authorization Records in DNS (MARID) working group in the absence of a consensus regarding whether to ratify the proposed Sender-ID specification, which combines Microsoft's Caller-ID and the Sender Policy Framework (SPF).

Instead of adopting Sender-ID as a standard, the proposal remains 'experimental'. MARID was originally formed to develop a standard approach to address the sender authentication problem.

Another blow was dealt when AOL also announced that it had 'serious technical concerns' with Sender-ID, and that it would only use SPF for checking inbound mail for its subscribers.

With a fundamental lack of agreement over the technology to be used, the creation of a unified standard will likely be left to the vendors as they continue to develop their own approaches. Those approaches may attempt to get around proposed Microsoft licensing and patent issues that are incompatible with many existing open source licenses.

Advocates of the original SPF solution believe that further adoption has been hindered over the past six months, as many in the industry have been waiting for the new standard to emerge.

The remainder of this newsletter will now focus on SPF, the predominant sender authentication system in use today.

Example: simplified SMTP conversation

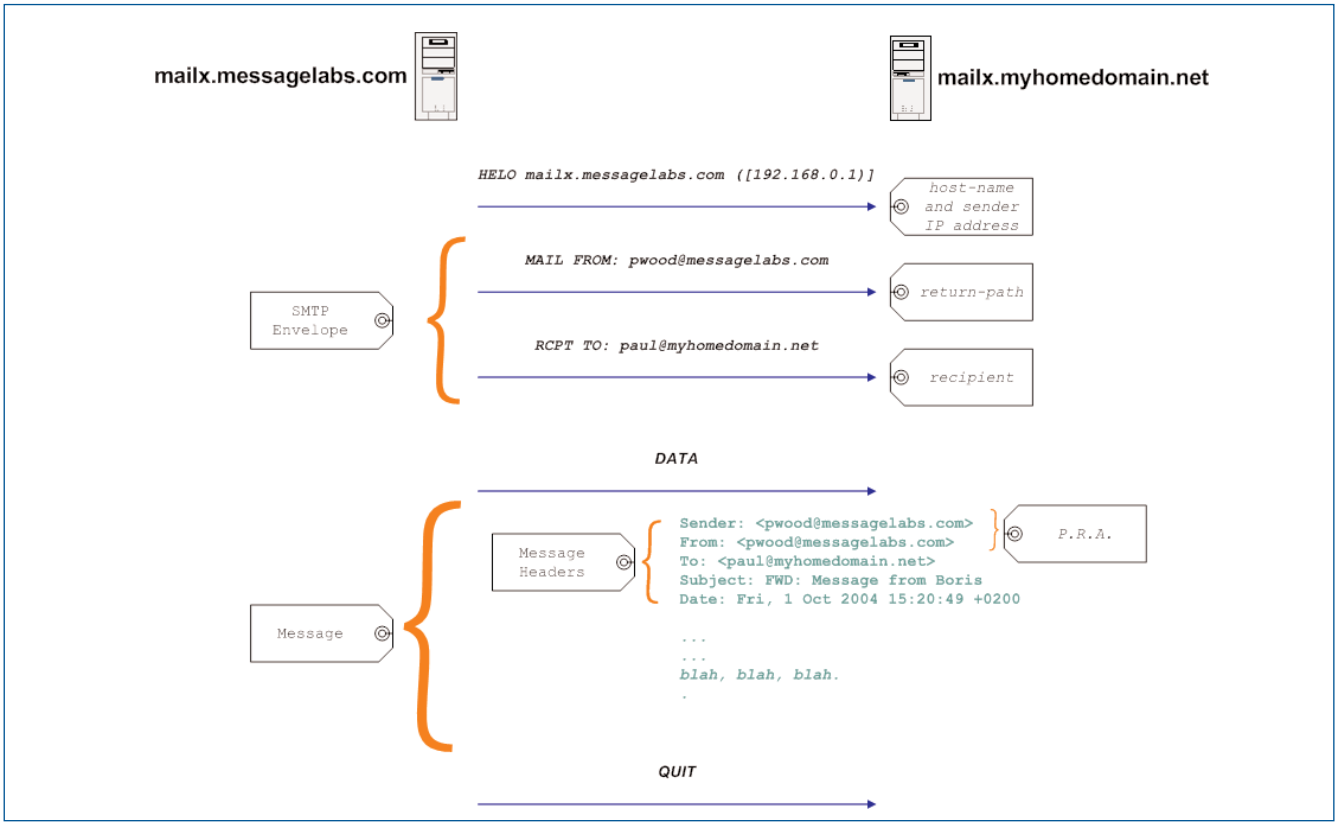

There are already plans to update the SPF system with a 'Unified SPF' proposal, which allows senders and receivers to publish one single record that can be used to check for the main identity types. These identities include the 'host name', 'return-path', and 'Purported Responsible Address' (PRA), the technique that Sender-ID uses to check the record against the most recent sender email address.

This unified approach will be backward compatible with the current 'Classical SPF' system, providing a single SPF record, which addresses the needs of the other competing schemes, further enhancing integration between them.

As these types of systems mature, and as one becomes the dominant player, it is likely that adoption will grow. Organisations such as MessageLabs will be able to absorb many of the implementation issues, leaving businesses only needing to provide the authentication information via their DNS. Everything would then be handled by the control towers within MessageLabs' own infrastructure, for example.

One of the detractions for SPF concerns the forwarding of email (for example, when using ~/.forward files in UNIX), and preserving the integrity of the return-path data contained within the email envelope. There are mechanisms available to support the rewriting of the envelope, but this is considered by many to break the current email protocols.

This really only becomes an issue if an email is passed through a 'forwarder' or a go-between, such as a university alumni address forwarding service, as the email may appear as being forged when processed by the destination server. To overcome this, the email intermediary would be required to rewrite the envelope to match the intermediary sender domain.

The problem arises because SPF validates the sender via the 'MAIL FROM:' domain, also known as the 'return-path'. In the SPF environment, a forwarded mail should be marked with a 'return-path' matching the forwarding agent, rather than via the original sender address. The return-path is commonly used for distributing Non Delivery Report (NDR) bounce messages if the email is undeliverable; in the SPF universe a NDR is delivered back along the same path through which it arrived, rather than by direct delivery.

With SPF, a Sender Rewriting Scheme (SRS) is required to ensure the appropriate changes are made to the envelope so that each go-between can correctly relay bounce messages to the original sender. Additionally, precautions are taken to ensure that spammers are not able to hijack bounce-back messages, guaranteeing that the return-path was added by the agent and not by another email server, and ensuring that they are only valid for a limited duration. Eventually, it is likely that most MTAs will support some form of SRS.

Essentially, SPF will allow an administrator to apply a domain-level policy to dictate how email from this domain is delivered to the rest of the Internet, and if a server in another domain is forwarding an email from their domain, this is likely to break the policy.

However, forwarding email ***will*** work with SPF for the following scenarios:

- You are ***forwarding*** someone else's mail as an attachment or as inline text, ***but*** keeping the 'From:' line as your own. For instance:

```
From: <paul@myworkdomain.com>
Subject: FWD: Message from Boris
```

- You are in one location, say work, and you and want to send an email from your ***home*** email address ***using*** your ***work*** email server. This will work with SPF ***if*** your work email server is configured as an authorised sender in your home email address' DNS SPF record. For instance:

```
From: <paul@myhomedomain.net>
X-Server-From: mailx.myworkdomain.com
Subject: Message from home email address
```

- Under this scenario, the outgoing mail server mailx.myworkdomain.com ***must*** be listed in the DNS SPF record for myhomedomain.net as an authorised mail server.

- You are in one location, say work, and you want to send an email from your ***home*** email address ***using*** your ***home*** ISP's email server. This will work with SPF provided your work network ***allows*** you to connect to your home ISP's server ***and*** you can authenticate yourself to that server (for instance, by logging-in and checking your email first, or by using the SMTP AUTH protocol). For example:

```
From: <paul@myhomedomain.net>
X-Server-From: mailx.myworkdomain.com
Subject: Message from home email address
Subject: Message from home email account
```

- Again this would work fine. I just have to login to the

myhomedomain.net POP3 server and check my mail before I send something, or use SMTP AUTH to login to the SMTP server.

SPF does ***not*** work for the following scenarios:

■ You are in one location, say work, and you want to send an email from your ***home*** email address, but you ***cannot*** connect to your ***home*** ISP's server ***and*** your work server is ***not*** listed as a valid outgoing mail server in your home domain's DNS SPF record.

This scenario is more likely than the counter-example above, unless you can control your home domain's DNS records.

■ You are sending an email to an address that forwards your mail onwards ***without*** using SRS. If you send an email to alias@alumni.alma-mater.edu and that address ***forwards*** automatically to paul@myworkdomain.com, then SPF checks at myworkdomain.com may ***fail*** and cause your mail to bounce if alumni.alma-mater.edu do not rewrite the envelope 'MAIL FROM:' so that bounces go back through them:

```
From: <paul@myhomedomain.net>
Return-Path: <paul@myhomedomain.net>
Subject: Message to alumni alias account
```

Using SRS, this may be rewritten as something like:

```
From: <paul@myhomedomain.net>
X-Actually-From: <paul@myhomedomain.net>
Return-Path: <SRS0+BG34=T9=myhomedomain.net=paul@
alumni.alma-mater.edu>
Subject: Message to alumni alias account
```

• Where SRS0+BG34=T9= ... is the way SRS may rewrite the return-path, creating a 'hash' and an 'expiry' time, which only has meaning on the server that created the hash, such that it cannot easily be spoofed

Subsequent forwarding may result in the return-path being further rewritten, using a special SRS1 marker, which includes the original forwarder domain, and changing the domain at the end to refer to itself, for example:

```
SRS1+alumni.alma-mater.edu=BG34=T9=myhomedomain.net
=paul@myworkdomain.com…
SRS1+alumni.alma-mater.edu=BG34=T9=myhomedomain.net
=paul@forward1.com…
SRS1+alumni.alma-mater.edu=BG34=T9=myhomedomain.net
=paul@forward2.com
```

• Where SRS1+alumni.alma-mater.edu= ... is the way subsequent SRS changes to return-path may be handled when additional forwarding is required

For the most part, this is acceptable. Only forwarders need to be concerned by SRS, and most servers and email programs support authentication (there are a number of methods that can be used to identify the sender, including 'login-before-send', and SMTP AUTH), so this may be implemented with very little pain.

Although not designed as an anti-spam measure in itself, sender authentication schemes make it possible for email systems to check which computers are authorised to send email on behalf of a particular domain, allowing the server to confirm whether the email did originate from the domain it claims to be from. By checking not only the return-path but also the originating server's address, if the originating server isn't authorised to send email on behalf of the domain in the return address, then it's probably spoofed and should be treated with suspicion.

This should make SPF an excellent proposition to prevent domain spoofing from virus outbreaks such as MyDoom.A or Sobig.F; however, it should perhaps be noted that email domain authentication will not in itself eliminate phishing type fraud. While currently these types of emails may try to pretend to be from legitimate domains, once SPF becomes commonplace, they do not have to keep doing that. In the very near future, we should expect to see emails from bogus domains such as c **1**tibank.com, paypa **1**.com and e **e**bay.com , as SPF allows both legitimate and illegitimate senders to create valid DNS records.

Once a sender authentication standard is in operation, there would be no requirement for business end-users to do anything more. Businesses failing to adopt one of these solutions may be concerned that one day all of their emails

could be rejected by servers that only accept email from systems that support sender authentication. It is unlikely that this scenario would ever come to pass, but their email will almost certainly become more aggressively filtered.

This prospect has already resulted in many spammers publishing SPF records for their domains.

The temptation by many advocates of SPF as an anti-spam measure has been to lower the threshold of any email that passes a sender authentication test, raising it for any that fails. All well and good in theory, but in practice spammers are already beginning to take advantage of this tactic.

As always in the anti-spam arena, any new technology approach is closely followed by advancements in spammers' strategies. Although it may seem that the adoption of SPF by spammers is a bad thing, it could be considered as somewhat reassuring; in order for SPF to become a useful tool in combating spam, effective domain reputation and accreditation schemes are also required to ensure that SPF records are valid and to track which domains are responsible for generating spam. Domain reputations can be profiled over time, white-listing and accreditation makes it possible to discriminate between the good and bad senders, and using techniques for grey-listing to track the senders whose reputations are unknown.

In one sense, it's a bit like confusing identity cards with burglar alarms. Just because a thief carries an identity card, that doesn't mean you can switch off your alarm. However, if you had a burglar alarm that checked identity cards of anyone in the building and then compared the identities to a list of known burglars, it could be used to correlate spammer identities.

As more spammers adopt SPF as a means to validate emails which may otherwise be blocked, the process of identifying spammers may become easier, as databases of known spammer domains become more trustworthy.

Spammers are already registering their own domains with completely correct SPF information. They often buy these domains in bulk and will change them frequently. Eventually, the spammers may be able to generate new identities faster than they can be tracked!

One threat to these schemes arises from the potential to create policies that are inadequate. For example, a spammer may provide an SPF record for its domains that include the dial-up and broadband addresses of computers running as an open-proxy, such as those typically infected by a backdoor Trojan horse program.

Criminals are already using computer viruses to create armies of 'zombie' computers that are hired out for sending spam. These zombies are typically desktop machines in homes and offices throughout the world, totally unrelated to the spammer, many with typically broadband 'always-on' connections. When they become infected, these zombies can then be used to send spam which will appear to come from a legitimate SPF compliant domain and may not be rejected. Services that check SPF will have to create heuristics to check the SPF record itself is strong enough, and is not covering large tracts of insecure address space.

Many businesses are now being encouraged to publish SPF records for their domains. Even if your mail server doesn't currently use SPF, there are many others who are, and this number is set to increase.

This means that by creating an SPF record for your domains, others can validate the true origin of your emails. In turn, this will offer greater protection against domain spoofing, particularly during a virus outbreak or a 'Joe-job' spam attack. (A 'Joe-job' is where a malicious spammer sends large volumes of spam forged to appear as though it were sent by an innocent user or domain, which may become flooded by the NDR bounces).

You can easily add an SPF policy to your domain by adding a TXT (text) record to your DNS domain zone, for example if using the BIND DNS server, MessageLabs clients may choose to add a record like this:

- myworkdomain.net IN TXT "v=spf1 include:spf.messagelabs.com ~all"

- Where v=spf1 indicates this is an SPF record, and that the SPF check will be made against the SPF record of the spf.messagelabs.com zone

- The ~all indicates that a SPF check should result in more aggressive filtering if the lookup ***doesn't*** match the given criteria

- Should customers wish to be ***even more*** aggressive, perhaps after rigorous testing, they can use –all , which indicates that all their mail comes direct from MessageLabs' control towers, without exception

You can get more information and assistance on setting-up SPF records by visiting the SPF Wizard at http://spf.pobox.com/wizard.html.

Aside from the technological considerations, MessageLabs' approach has not just been to wait and see what new standard emerges; fundamental to the core of any solution will be the development of an effective domain accreditation scheme to support any implementation. MessageLabs has also been working closely with Meng Wong (the designer of SPF), Microsoft and the IETF to try and ensure the implementation of sender authentication is successful.

It has already been seen that one of the major benefits of implementing a managed service approach to email security can be that the managed email security provider will take responsibility for managing the implementation of the sender authentication scheme. Moreover, a managed service approach is more likely to benefit from the economy of scale afforded by such a system when considering the implementation of a reputable domain accreditation scheme that is able to identify the true spammer domains from the core business traffic.

In the coming months, we may begin to see a greater take-up of domain policy systems such as SPF, Caller-ID or DomainKeys; this might mean that, as they compete for market domination and closer integration, other approaches such as 'challenge-response' systems and 'electronic payment' and 'cryptographic puzzles', don't generate the broader adoption that their advocates would have hoped for.

These mechanisms are equally vulnerable to spoofing, and scalability becomes more of an issue as global adoption seems less likely. You may already find yourself responding to 'challenges' for emails that you didn't even send, just to make sure that any important mails are being delivered.

Therefore it is important to maintain a balanced approach when considering the benefits and shortcomings of having such an open communications medium.

*** The information relating to MessageLabs' services contained in this newsletter is based on data generated internally by MessageLabs and has not been subject to an independent review by a third party.**