Inbox Pros

SPF

The Value of Sender Policy Framework to Security and Deliverability

An Inbox Pros Whitepaper

The value of SPF Authentication to Security

“SPF,” or Sender Policy Framework, is the most basic form of email sender authentication and it has been adopted worldwide as the cornerstone of every sender’s identity and reputation.

When an ISP or secure mail server receives an email, it does not simply deliver it to the inbox, it first authenticates the sender of the email, and then it passes it through content filters, blacklists, and etc. before deciding if the message reaches the inbox, the spam folder, or is rejected based on simple rules, heuristics, Bayesian probabilities, and/or historical reputation data.

Authentication is the first step because it is arguably the most important. The email space has historically experienced considerable abuse at the hands of spammers and phishers, so ISPs take sender identity extremely seriously. Any recipient mail server is legally empowered to reject any email for any reason, and they will if they feel the message is in any way a possible threat to their users.

An email sender has three main ways of identifying themselves to the recipient mail server:

1. The “From:” name. This is the same name that the user sees when they read the message, and it includes the sender’s “domain” or the part after the “@” symbol.
    a. This domain is registered in the DNS, but it is possible for the sending mail server to insert a fake “From:” name, a tactic used by phishers and spammers.

2. The “Return-Path” (aka the “Envelope-From”) domain. This is the domain that the sending mail server uses to identify itself, separate from identifying the sender.
    a. If the sender is using an ESP, the Return-Path often uses the ESP’s name, not the sender’s, to announce its identity to the recipient server. This means that a malicious sender could use their ESP’s “Return-Path” to help anonymize themselves.

3. The originating IP address. This is the IP address of the sending mail server, and an “rDNS” lookup tells the recipient server who the IP address belongs to.
    a. While it is possible for a malicious sender to “spoof” an IP address when sending an email, it is very difficult and easily detected; however, a malicious sender could be renting time on a legitimate server.

SPF resolves these security problems by providing a mechanism for a legitimate sender to publically state to the ISPs and other recipient mail servers which IP addresses are allowed to send their mail.

The very first authentication an ISP uses is an SPF check. This involves the ISP gathering the "From:" domain and looking it up in the DNS to find that domain's public SPF record. The SPF record will list the authorized IP addresses, which the ISP then compares to the originating IP address of the email. If the originating IP address of the email is not in the SPF record, the ISP knows to treat the email with suspicion or even reject it.

If there is no SPF record, that means that the owner of the domain has announced that any IP address may send mail with their name on it. This makes it much easier for malicious actors to pose as that sender. It also prevents the legitimate sender from using other, more specific authentication methods such as DMARC.

The value of SPF Authentication to Deliverability

ISPs consider emails that do not pass SPF to be potential phishing or spamming attempts, which they consider abusive and harmful to their customers. Similarly, allowing in messages from domains without authentication increases the potential risk to their customers by making successful phishing and spamming much easier.

When an ISP checks an SPF Record, it is looking for multiple things.
Firstly, it is looking for approval of the IP address that the email in question came from.
Secondly, it is looking at the closing "all" statement, which informs the ISP of how suspiciously the domain wishes it to treat non-matching emails.

For this reason, it is recommended to identify and authenticate only the IP address(es) allowed to send from the domain and to use a hardened "all" statement that instructs the ISP to treat non-matching emails with extreme suspicion.
Note that ISPs are under no obligation to obey the "all" statement and may treat the email according to whatever they wish, but in actual practice ISPs treat the "all" statement as a minimum level of suspicion.

If an ISP were to check for an SPF record and find nothing, it would consider the sender to be very lax in security and therefore much more likely to inadvertently allow a malicious actor to take advantage of them and abuse the ISP's customers. Additionally, as without an SPF record an ISP has much more difficulty telling the difference between legitimate and fake/abusive mail, if abusive mail is detected the ISP will often simply filter the entire domain, quarantining or even rejecting all of its messages, good and bad.

By implementing a hardened SPF record, a domain is not only identifying its messages as legitimate, it is giving the ISP confidence that the domain does not pose a security risk to its customers. ISPs have many criteria, some public and some secret, for judging if an email is unwanted or should be sent to the spam folder, but the potential for harm trumps them all.

How to deploy an SPF record

An SPF record is a TXT record that is uploaded to a domain's DNS entry and contains the IP address(es) of each server that is allowed to send email on behalf of that domain.

In order to deploy an SPF record, the first step is to identify the IP address(es) used by every server that is allowed to send email from the domain in question. If the server is using a shared pool of IP addresses, an "include:" statement can be used instead to point an ISP at the SPF record that is already created for that IP pool; this is common when using ESPs.

Next, access the DNS portal for the domain in question, usually through the registrar's website. In the domain, create a new "TXT" record. The syntax of the TXT record will always begin with the phrase "v=spf1" in order to differentiate the SPF record from any other TXT records in the domain's DNS entry.

The final result can be as simple as this: "v=spf1 ip4:192.168.0.1 -all"

Or, it can get more complicated, such as this SPF record used by AOL for their domain aol.com:

"v=spf1 ip4:204.29.186.0/23 include:spf.constantcontact.com include:aspmx.sailthru.com include:mail.zendesk.com ~all"

The above record includes CIDR IP ranges ("/24" "/22" etc.) in order to avoid listing out every IP address, because if the SPF record is longer than 255 characters, the ISP will not be able to read it.

Additionally, it has several "include:" statements which reference other SPF records, whose IPs AOL wishes to authorize. It is important to note that these other SPF records may have their own "include:" statements and etc, and so ISPs have announced that they will not check more than 10 total "include:" statements in any single SPF record in order to keep their servers from being bogged down accidentally or maliciously by nesting lookups.

Finally, the SPF record ends with "~all" or "softfail" which tells an ISP to treat an email that claims to be from aol.com but has a sending IP not included in this record with suspicion.

Contact Inbox Pros

Contact Inbox Pros to get an in-depth analysis of your current practices and a custom-tailored complete solution that draws on extensive research, long experience, and unmatched industry contacts to improve your security and deliverability.

Please use our Contact Form or email us at info@inboxpros.com to learn more about how we can help your business.

To learn more, check out our blog, and follow us on social media!

Blog

Facebook

YouTube

Twitter