21
Jan
2010

Some thoughts on best practices for SMTP blocking of e-mail spam

Blocking e-mail spam at the time of SMTP (Simple Mail Transfer Protocol) transfer has become a best practice. There is no point wasting precious bandwidth & disk space and spending time browsing a huge spambox when most of the incoming flow is clearly spam. At LinuxForce our e-mail hygiene service, LinuxForceMail℠, makes extensive use of SMTP blocking techniques (using free and open source software such as Exim, Clam AV, SpamAssassin and Policyd-weight). But we are extremely careful to only block sites and e-mails that are so "spammy" that we are justified in blocking it. That doesn't prevent false positives, but it keeps them to a minimum.

Recently we investigated an incident where one of our users had their e-mail blocked by another company's anti-spam system. In investigating the problem, we learned that some vendors support an option to block e-mail whose *Received* header is on a blacklist (in our case it was Barracuda, but other vendors are also guilty). Let me be blunt: this is *boneheaded*, but the reason is subtle so I can understand how the mistake might be made.

First, blocking senders appearing on a blacklist at SMTP time is good practice. But to understand why blocking *Received* headers at SMTP time is bad, it is important to understand how e-mail transport works. The sending system opens a TCP/IP connection from a particular IP address. That IP address should be checked against blacklists. And other tests on the envelope can help identify spam. But the message headers including the *Received* header are not so definite. We shall see that even a blacklisted IP in these headers may be legitimate. So blocking such e-mail incurs unnecessary risks.

The problem occurs when a user of an ISP (Internet Service Provider) sends an e-mail from home, they are typically using a transient, "dynamic" IP address. Indeed it is possible that their IP address has just changed. Since the new address may have been previously used by someone infected with a virus sending out spam, this "new" IP address may be on the blacklists. So, due to no fault of your own, you have a blacklisted IP address (I will suppress my urge to rant for IPv6 when everyone can finally have their own IP address and be responsible for its security).

Now, when you send an e-mail through your ISP's mail server, it records your (blacklisted) IP as the first *Received* header. So your (presumably secure) system sending a legitimate message through your ISP's legitimate, authenticating mail server is blacklisted by your recipients' overambitious anti-spam system. Ouch. That is why blocking such an e-mail is just wrong. This kind of blocking creates annoying unnecessary complications for the users and admins at both sides. Using e-mail filtering to put such e-mails into a spam folder would be a reasonable way to handle the situation. Filtering is able to handle false positives whereas blocking generates unrecoverable errors.

Do not block e-mail based on the *Received* header!

Posted by CJ Fearnley *in* Security, Systems Management, Tech Notes, Ubuntu, *0 comments*