

Re: SpamAssassin

Source: <http://unix.derkeiler.com/Newsgroups/comp.os.vms/2007-02/msg01390.html>

- *From:* david20@xxxxxxxxxxxxxxxxxxxx
 - *Date:* Tue, 27 Feb 2007 14:09:36 +0000 (UTC)
-

In article <ervot9\$s9n\$1@xxxxxxxx>, helbig@xxxxxxxxxxxxxxxxxxxxxxxxxxxx (Phillip Helbig---remove CLOTHES to reply) writes:

In article <erunca\$88v\$1@xxxxxxxxxxxxxxxxxxxx>, david20@xxxxxxxxxxxxxxxxxxxx writes:

In article <erif9h\$fqv\$1@xxxxxxxx>, helbig@xxxxxxxxxxxxxxxxxxxxxxxxxxxx (Phillip Helbig---remove CLOTHES to reply) writes:

Does anyone here have experience with SpamAssassin? Do you recommend it?

Since noone else has responded I'll have a go. Basically the answer is there is no such threshold. No content scanning anti-spam product is 100% accurate. The best will only claim 98% accuracy. Which at first sight sounds a lot but really means it gets it wrong for 2 out of every 100 mail messages. Those mistakes will either be false positives (mail which is mistakenly considered to be spam but isn't) or false negatives (mail which is spam but is missed). The threshold just changes the ratio of false positives to false negatives. The only way you can guarantee that all legitimate mail gets through is to set the threshold to a ridiculously high level in which case all mail (including spam) will get through.

Yes, in theory. In my case, in practice, the provider flags it as spam for a score of 5 or higher. Looking at well over 1000 messages, NO legitimate email was flagged as spam. Of the stuff not flagged as spam, about 2/3 is. About 60--70% of the total gets flagged as spam. So, for 100 messages a day, I could get rid of 65 or so quickly. Of the remaining 35, I have to fish out a dozen legitimate ones. So it does save time with no appreciable risk.

Re: SpamAssassin

Not quite sure I understand the above
is that

65 spam messages marked as spam
35 messages not marked as spam
of those
12 messages not marked as spam are legitimate mail
23 messages not marked as spam are really spam

ie you have a 23% false negative rate using spamassassin with a threshold of 5
but a zero false positive rate.

That sounds like a pretty appalling rate of mistakes – either 5 is a particularly high score to set as a threshold for spamassassin so that most mail will get through or I'd suspect that the provider isn't keeping the spamassassin rules uptodate so that the spammers are getting their mail through. If it is the latter then things will become interesting when the provider next updates the rules since for a brief period thereafter the situation may be totally reversed with few false negatives and a number of false positives.

My provider would send a message to the (alleged) sender if a message is not relayed to me because it is suspected of being spam. Thus, a legitimate sender could find out that this was a problem. (I'm not sure this is a good idea. What I observe myself is that a spammer sends email to a non-existent address somewhere else with my address (or a non-existent address in my domain) as the forged sender.

The spam scanner at the other end sends me (or, if the user doesn't exist, my Postmaster, who is me wearing another hat) a message that there was a problem. Maybe, however, the spammer really wanted to spam my addresses and is using the spam filter at the other site as a spam relay. (Fine point: for non-existent users, the Postmaster only gets the mail if they are also syntactically invalid VMS usernames, since I drop all the rest during the SMTP dialog, since this is by far the most spam I get (dictionary attack).

Yes this is backscatter and if your provider is accepting mail for your systems and then running spamassassin on it before sending a bounce then they are part of the spam problem not the solution. It is not acceptable nowadays to send a bounce to the supposed sending address after you have discovered a mail message is spam or contains a virus. Some systems can run an anti-spam product during the SMTP dialogue and can reject the mail at the end of the DATA part which is acceptable. There are now some DNSBLs which list sites which produce backscatter eg Spamcop. I personally think that is going too far since there are still some circumstances (account going overquota where the account is not on the receiving mailserver itself etc) which most mail systems cannot cope with other than by producing a bounce.

Re: SpamAssassin

The main advantage for me is: if I choose to drop the spam, then I don't have to have an ALPHA always have the cluster alias, but a VAX (with TCPIP 5.3) would be OK. (A lot of spam is email to non-existent users. These generate bounces which, because the sender is often faked, bounce back. With 5.4, I can reject email to non-existent usernames (at least if they are valid VMS usernames, which most of them are), but that runs only on ALPHA.)

Unless your Dynamic-DNS provider has a list of all your valid email addresses then no anti-spam product it runs can determine that a message is for a non-existent account on your systems.

True, but since no legitimate mail is sent to non-existent users, it could be flagged as spam based on other grounds.

True but then you have potentially changed an SMTP dialogue rejection of an invalid address (assuming the connection is to your ALPHA) into a backscatter bounce message from your provider complaining about spam.

If these are hobbyist VMS systems I'd get PMAS and PMDF. I was going to say that if they were commercial systems then I'd pay for and install MX which runs on Alphas and VAX and would reject mail for non-existent addresses. The price for which was not too excessive. Unfortunately I just had a look at <http://www.madgoat.com/mx.html> and it appears that no new orders are being taken (though existing customers will continue to be supported). The freeware version MX4.2 is still available and I would hope that would support rejecting mail for non-existent users but can't be sure.

If this change in the status of MX is permanent then it is bad news since it leaves nothing filling the small business gap for a real mailserver. Larger enterprises can afford the price of PMDF, hobbyists can get PMDF free. But smaller commercial businesses are left just with the mailservers which come with the IP stacks.

David Webb

Re: SpamAssassin

Re: SpamAssassin

Security team leader
CCSS
Middlesex University
.