

Re: SpamAssassin

Source: <http://unix.derkeiler.com/Newsgroups/comp.os.vms/2007-02/msg01358.html>

- *From:* helbig@xxxxxxxxxxxxxxxxxxxxxxxxxxxx (Phillip Helbig---remove CLOTHES to reply)
 - *Date:* Mon, 26 Feb 2007 23:02:01 +0000 (UTC)
-

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

In article <erif9h\$fqv\$1@xxxxxxxxxx>, 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.

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

Re: SpamAssassin

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).

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.