

Re: hardware raid solutions?

Source: <http://linux.derkeiler.com/Mailing-Lists/Ubuntu/2006-07/msg01791.html>

- *From:* "Eric S. Johansson" <esj@xxxxxxxxxxx>
 - *Date:* Tue, 11 Jul 2006 21:19:04 -0400
-

David Abrahams wrote:

"Eric S. Johansson" <esj@xxxxxxxxxxx> writes:

sorry, I said that wrong. this is for an open-source project.

[1] python knowledgeable help wanted. Job descriptions available on request.

Will this do a magically-better job than my current combo of SpamBayes and the SpamCop blacklists applied by my sysadmin?

yes. better quality filtering, more opportunities to eliminate false positives, fewer special case hacks, and most importantly, less work for you.

note: this filter system runs inside of postfix using the pre and post queuing filters stages

geek view starting from the very front:

----- front-end (can be distributed among multiple relay servers) -----

1) Blacklist test?

If blacklisted, does it contain a very large proof of work stamp (i.e. 10 minute) for emergency bypass a.k.a. Brown listing? yes, pass. no, 5xx return code the message

2) does the to: address exist on the mail server. Enter in local Brown list if not

3) does e-mail arrived too fast for a given address? if so, 4xx return code the message the message

----- backend

4) proof of work stamp test, make a stamp, passed directly to inbox

5) friends "automatic white" list. if I know you, you pass

6) slow white list (match a pattern, don't go to jail)

6a) sumo filter. If the message is bigger than 50k, pass)

Re: hardware raid solutions?

7) content filter (CRM 114 with three band interpretation of pR)
messages are directed either to inbox, dumpster, or spam trap
where the human can interpret whether the messages are spam or not.

the spam trap user interface is a simple mechanism for recategorizing messages they could not be analyzed by any the other stages. the recategorization trains a content filter for at first higher, then different accuracy.

I am rather proud of the fact that myself and a friend made a spam trap user interface that is simple enough that an ordinary administrative assistant could manage the spam trap handling for a company of roughly hundred people in only 10 minutes per day at most. That is, as long as CRM 114 behaves itself. :-)

an interesting side effect of the friends list and the content filter is that if you look at the score of any message passing the friends list, if it scores as spam, that's an extremely high probability indication that the content filter is having problems. Typically I have found that around 5%–15% of the messages coming through the friends filter are considered spam.

on the outbound side, all messages are given stamps if they are going to people you don't know. Proof of work stamps make a great introducer and this level of effort is virtually invisible in most organizations but the benefit is high especially when you consider that spam assassin and a few other tools have the stamp recognition code in place.

There are other features I haven't yet exploited. For example, the output of the dumpster (clearly spam) can be used to build a brownlist database and if one is feeling the really clever, sharing the brownlist with others will make for a more inclusive and accurate brownlist.

Interpreting the brownlist entries according to what CIDR block they share with other brownlist entries could be used as a trigger to force spam trap interpretation until one has a sufficient number of "spots in the cidr block" to declare the block bad.

but the one place where I absolutely need the most help is closing the feedback between the e-mail client and my system. When spam leaks through, I have no way to easily communicate from the client to the antispam gateway. I just don't have the time to acquire the knowledge to do something that will work well for the ignorant user.

so that's it in a nutshell. Feel free to contact me off list. Any further discussion and I will take it to new thread.

—

ubuntu-users mailing list
ubuntu-users@xxxxxxxxxxxxxxxxxxxxx
<https://lists.ubuntu.com/mailman/listinfo/ubuntu-users>