



Published on Free Software Magazine (<http://www.freesoftwaremagazine.com>)

Configure Exim with anti-spam

Spam's off! Make it so with Exim and SpamAssassin

By [Ryan Cartwright](#)

A few comments on my article [The perfect network server in issue 17](#) requested some more in depth follow-up pieces. This is what I hope to be the first of those. It focuses on Exim, the mail transfer agent (MTA), specifically setting it up with spam scanning. It is based on setups I currently use, hosted on Debian GNU/Linux.

This is an intermediate article, I'm going to assume you are familiar with mail delivery technologies and terminology. If you don't know what SMTP, MX records and reverse DNS lookups are, you might like to do a little background reading before coming back to this article. I'll try to stick to a descriptive narrative style but some of it will inevitably involve technical language.

What Exim is not

Exim is a mail *transfer* agent. It receives messages, usually by SMTP, figures out what to do with them according to its configuration and transfers them to another location based on that information. The new location may be a locally based mailbox, another server or another daemon running on the same box.

Exim does not do POP3, IMAP, shared calendars or make the tea

It does *not* deliver mail to client machines, and does *not* handle mail user agent message creation functionality. In short, Exim does SMTP and related stuff. It does not do POP3, IMAP[1], LDAP, shared calendars or make the tea. It's important to note that Exim is not capable of pulling mail: it expects all messages to be delivered to it [2].

Debian Exim packages

Debian has a few packages for Exim. At the time of writing the current major version is 4. The key packages to install are `exim4-base`, `exim4-config` and one of the two `exim4-daemon-?` packages. If you are installing a brand new server, then you can install Exim as a task during the Debian installation stage (see [The perfect network server](#) for information on that). If you already have a Debian server then you can install it using `tasksel install mail-server` or with your favourite package manager.

The two Exim4 daemon packages within Debian called `exim4-daemon-light` and `exim4-daemon-heavy`. Installing either will install the `exim4-base` package as a dependency but not `exim4-config`. So, you should install `exim4-config` as well as it comes in handy.

The light daemon package is a perfectly adequate Exim install but leaves out things like SQL data lookups, virus/spam scanning integration and Secure Password Authentication SMTP. All of those are included in the "heavy" daemon and, as you'll want to do spam scanning, you'll need it. Note that `tasksel`, and thus the

Configure Exim with anti-spam


Debian installer, will install `exim4-daemon-light` but you can replace it simply by installing the heavy package afterwards. So, after all that, `apt-get install exim4-daemon-heavy exim4-config` does the trick.

Exim initial configuration

I'm not aware of any GUI tools to configure Exim, beyond address and account management. I started a fairly long debate on the necessity of server GUIs in the last article, but as they say, "write what you know"... so I'll stick to configuring Exim via the shell.

Having installed `exim4-config`, you might as well put it to good use

Having installed `exim4-config`, you might as well put it to good use. To use `exim-config`, run (either as root or sudo) the command `dpkg-reconfigure exim4-config`. I won't detail every step of this because each has decent explanatory text within—and I'd take up most of the article installing Exim. Here are a few pointers on the latter steps.



The debconf way to set-up Exim

Mailbox format

The Exim `local_delivery` transport is used to deliver messages to local mailboxes. By default that is in `mbox` format. There is an option to use the popular `Maildir` format, which you can set here. I prefer Maildir, so I choose to use it here. For a discussion on which is best I suggest you put your flame-proof suit on and hit Google or add a comment. Briefly, `mbox` stores all messages in a single text file. Maildir stores each message as an individual file within sub-directories of your main maildir. An advantage of Maildir is that it doesn't require file locking, so there are fewer delays.

Single or multiple config files

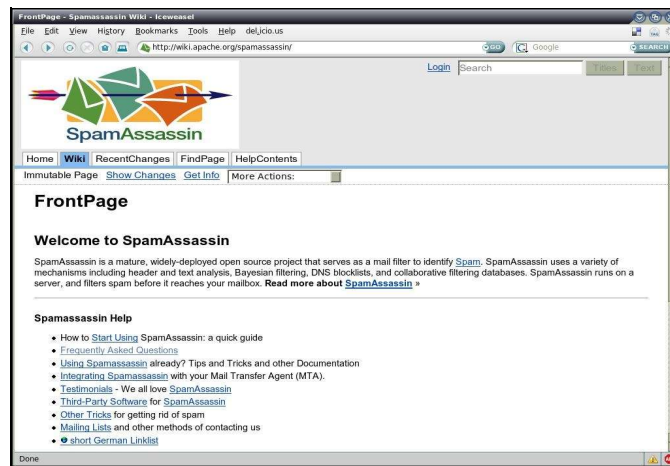
The next step is, as far as I know, unique to Debian (and its derivatives). You can choose to have the various routers and transports within separate config files or all lumped in one file. The former is the more usual Debian way of doing things, the latter is more common outside of the Debian world view. If you have experience of Exim on other distributions then go for the latter, otherwise the former may help you find things better. Both types will contain the same configuration options—just in different files. I have used multiple files and refer to them here. If you choose a single file then all your configuration options will be found in `/etc/exim4/exim4.conf`.

SpamAssassin

The automatically generated configuration will create a perfectly usable Exim server: any local user will be able to have mail delivered to their mailbox by Exim ready for to be picked up. What it won't do as yet is any spam scanning. I'll look at spam scanning now. The traditional anti-spam approach for GNU/Linux is SpamAssassin. What's good is that v4.50 Exim has a built-in interface for SpamAssassin. This was formerly the `exiscan_acl` plugin.

SpamAssassin requires you to enable it before it will run

Once installed (e.g. `apt-get install spamassassin`), SpamAssassin requires you to enable it before it will run. Edit `/etc/default/spamassassin` and set `ENABLE = 1`. You can also have SpamAssassin's rules updated on a nightly basis by setting `CRON = 1` in the same file. Once you've done that, you can start the `spamd` daemon with `/etc/init.d/spamassassin start`.



There are plenty of good SpamAssassin resources on the web. Including this wiki. By default SpamAssassin listens on port 783. If you install `exim4-daemon-heavy`, Exim will be already set to look for it on that port. If your SpamAssassin listens on a different port you should change the setting in `/etc/exim4/conf.d/main/02_exim4-config_options` accordingly.

Access Control Lists (ACLs)

You need to tell Exim how and when to use SpamAssassin. This can be done either in a router (which is run on messages after they've been received), or within an Access Control List (ACL). I'll do it within an ACL. ACLs are run against messages during SMTP sessions. Exim usually runs an ACL during the SMTP `RCPT command` and one after `DATA` command. Spam checking is generally made on the body, so you need to edit the latter ACL. This assumes SpamAssassin is installed and running. The settings we need to change are within the `/etc/exim4/conf.d/acl/40_exim4-config_check_data`. Uncomment and amend the default setting as follows:

```
warn
spam = Debian-exim
condition = ${if <{${message_size}{100k}{1}{0}}
message = X-Spam_score: $spam_score\n\
          X-Spam_score_int: $spam_score_int\n\
          X-Spam_bar: $spam_bar\n\
          X-Spam_report: $spam_report
```

Configure Exim with anti-spam

This ACL runs every message below 100k in size through SpamAssassin as it is received. The 100k limit eases the server load by not checking larger messages with attachments that are more likely to contain malware than spam. The ACL adds some additional headers to the message defining the spam “score”. I will use these in filters later.

Rejecting suspect spam

This configuration will warn you about problems, but will still relay suspicious messages. Why not deny relaying to them? It’s certainly possible to have the ACL reject messages above a certain threshold. I’m not sure that this is necessarily a good idea though.

SpamAssassin is good at its job but it will produce false positives

SpamAssassin is good at what it does but it will produce false positives. In particular I’ve found badly written and configured HTML e-mail newsletters to be an issue. Whilst it’s tempting to shut out such messages, it may not be practical for your users. It is perhaps better to accept the message with an additional header and act upon that later. To do that in Exim you use filters.

Filtering messages

Many clients allow you to setup a filter to detect and examine a header. You could, for example, set a client-filter to check the X-Spam-Score header for a value above 9.9 and dump any messages that match directly in the client’s trash or a separate spam folder. Your users may feel slightly better for not seeing higher scoring messages though. Even when client-side filters work, many have found that reducing the amount of delivered spam goes down well with users.

Exim has an extensive filtering system which can be used by individual users on messages delivered to them or globally on all messages relayed by the server. The latter is a system filter and you can have one per Exim installation. Specify the location of your system filter within any of the main config files. I put mine at the top of the `/etc/exim4/conf.d/main/02_exim4-config_options`. Add something similar to this:

```
system_filter = "/etc/exim4/system.filter"
system_filter_user = Debian-exim
system_filter_group = Debian-exim
system_filter_pipe_transport = address_pipe
system_filter_file_transport = address_file
system_filter_reply_transport = address_reply
```

The filter itself is a series of conditions and resulting actions. Here’s a sample condition for a system filter:

```
if $h_X-Spam_score_int is above 99
  and foranyaddress $recipients ($thisaddress contains "@mydomain.com")
  save /var/mail/suspect_spam
  mail to $thisaddress
  subject "[ SPAM Witheld ] $h_subject:"
  from "Company Mail Server <no-reply@mydomain.com>"
  text "This Message is sent automatically by e-mail software, please do not
  reply to it\n\nThe server suspects that a message sent to you by $h_From:
  is spam. It scored : $h_X-Spam_score: .\n The spam report is:\n
  $h_X-Spam-report: \n.\n\nThe original message has not been sent to you, but
  stored on the server.\n If you were expecting it, please contact your system
  administrator with the reference\n $message_id
  finish endif
```

Configure Exim with anti-spam

The second line will loop through all the recipient (To, CC, BCC) addresses for one that contains mydomain.com. This is done to ensure we only run this on inbound messages. Once matched, it appends the original message into a local file and sends a notification to the intended recipient, rewriting the subject. Note the use of the other additional headers in the notification text.

Now you need to handle messages scoring between 4.9 and 9.9. You can do this with another condition which rewrites the Subject header.

```
if $h_X-Spam_score_int is below 99
    and $h_X-Spam_score_int is above 49
    and foranyaddress $recipients ($thisaddress contains "@mydomain.com")
    then
        headers add OldSubject "$h_Subject"
        headers remove "Subject"
        headers add Subject "[ SPAM ] $h_OldSubject"
        headers remove "OldSubject"
finish endif
```

Header manipulation happens on-the-fly in Exim filters; so you needed to save the existing subject before using it in the new one.

Final words

Exim is a little like chess: straightforward to learn, a lifetime to master. SpamAssassin is really the same, so coming up with a comprehensive guide article that covers all questions is almost impossible. What I have tried to do is give some introductory guidance on configuring both on Debian GNU/Linux. Think of it as whetting your appetite. If you are looking to deploy an Exim server with SpamAssassin and other features I've not had room to mention, such as virtual users, virus and malware scanning etc. I would suggest you do some more reading before you start. The Exim website has comprehensive documentation but is perhaps not best suited to the beginner. A better place to start will be <http://www.exim-new-users.co.uk>.

Bibliography

[1] Have a look at something like [Courier IMAP,Cyrus](#), [Qpopper](#) et al for POP3 or IMAP servers.

[2] Try [fetchmail](#) for something that pulls mail from another server and feeds it into Exim for local delivery. The two work very well together.

Biography

[Ryan Cartwright](#) (/user/8833" title="View user profile.): Ryan Cartwright is IT Manager for [Contact a Family](#) (<http://www.cafamily.org.uk>), a UK National charity for families with disabled children where they make significant [use of FLOSS](#) (<http://www.cafamily.org.uk/oss>). He is also a FLOSS advocate and you might find him on the [GLLUG](#) (<http://gllug.org.uk>) mailing list.

Copyright information

This article is made available under the "Attribution-NonCommercial-Sharealike" Creative Commons License 2.5 available from <http://creativecommons.org/licenses/by-nc-sa/2.5/>.

Configure Exim with anti-spam

Source URL:

http://www.freesoftwaremagazine.com/articles/exim_and_anti_spam_spamassassin
