

**Spam and Virus Scanning with Exim 4
using Exiscan and/or SA-Exim**

Mini-HOWTO

©2003 Tim Jackson (tim@timj.co.uk)
V1.0.10 (25-Feb-2004)

Table of Contents

1. Disclaimer/Copyright/Distribution	4
1.1. No Warranty	4
1.2. Copyright	4
1.3. Licensing/Distribution Permissions	4
2. Introduction - what's all this?	5
3. Mail Scanning Concepts and Options	6
3.1. Introduction	6
3.1.1. Accept all messages, then scan	6
3.1.2. Scan incoming messages as they arrive (while the SMTP session is open).....	6
4. What Software To Use?	8
4.1. Exim patch	8
4.2. Spam scanning software	8
4.3. Virus scanning software	8
5. Installing the software	9
5.1. Installing SpamAssassin	9
5.1.1. Installing using RPMs	9
5.1.2. Installing without RPMs	9
5.1.3. Testing	9
5.1.4. Security reminder	10
5.2. Installing Exiscan	10
5.2.1. Installing as part of an Exim RPM	10
5.2.2. Installing without using RPMs	10
5.3. Installing SA-Exim	10
5.4. Installing Clam Antivirus	11
5.4.1. Compiling as an RPM	11
5.4.2. Security reminder	11
5.4.3. Virus database updates	11
5.5. Installing Sophos Antivirus/sophie	11
5.5.1. Installing as RPM	12
5.5.2. Virus database updates	12
5.6. Compiling Exim as an RPM	12
5.6.1. Setting up the Makefile	12
5.6.2. Setting the spec file build options	12
6. Configuring the software	14
6.1. Configuring SpamAssassin	14
6.1.1. For Exiscan users	14
6.1.2. For SA-Exim users	14
6.1.3. For all users	14
6.2. Configuring Exiscan	15
6.2.1. Overview/setting up DATA ACL	15
6.2.2. MIME checking settings	15
6.2.3. File attachment blocking settings	15
6.2.4. Spam scanning settings	15
6.2.5. Virus scanning settings	16
6.2.6. Ending the ACL	16
6.3. Configuring SA-Exim	17
6.3.1. The SA-Exim Config File	17
6.3.2. Setting Exim ACLs	17
6.4. Configuring Clam Antivirus	17
7. Getting it all running/Testing	19
7.1. Starting the daemons	19
7.1.1. SpamAssassin	19
7.1.2. Clam Antivirus	19
7.1.3. Sophie	19
7.1.4. Exim	19
7.2. Testing	19
7.2.1. Spam scanning	19
7.2.2. Virus scanning	21
7.2.2.1. Manual testing with Clam Antivirus	21
8. Appendix A – Building RPMs from source	22
8.1. Using tarballs (the quick way)	22
8.2. Using a build tree (the conventional way)	22
8.3. More Information	22
9. Appendix B – Removing headers	23
9.1. If you don't already have a system filter	23
9.2. If you already have a system filter	23
10. Appendix C – The GNU Free Documentation License	24

1. Disclaimer/Copyright/Distribution

1.1. No Warranty

Whilst this document is provided in good faith in the hope that it will be useful, the author provides NO WARRANTY WHATSOEVER as to the accuracy or otherwise of any of its contents. Use of any information provided in this document is entirely at your own risk.

1.2. Copyright

This document is copyright ©2003-2004 Tim Jackson (tim@timj.co.uk).

1.3. Licensing/Distribution Permissions

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "Appendix C – The GNU Free Documentation License".

2. Introduction - what's all this?

This document briefly describes how to install and configure the following items of software, so that they work together to provide not only e-mail handling, but realtime spam and virus scanning:

- Exim 4.x (this version refers to 4.24/4.30 but most parts are applicable to earlier versions of Exim 4)
- Exiscan/Exiscan-ACL (this document refers to Exiscan-ACL release 14)
- SA-Exim (this document refers to version 3.1)
- SpamAssassin
- Clam Antivirus/Sophos Antivirus/sophie

In here I hope to summarise much of the information already available from various sources on the Web, therefore helping those new to using these tools to understand the options available, and successfully get a basic configuration up and running. It is not entirely comprehensive, though, and is definitely not a substitute for reading the software documentation.

The instructions should be moderately generic, though they are geared to those using GNU/Linux and an RPM-based packaging system (e.g. Red Hat Linux) and as such provide some detailed information for those users. If you're using something else (e.g. deb packaging), you will need to refer to other sources for detailed installation advice, although the general configuration details given here should remain similar or the same. Details on how to acquire the various items of software are provided inline.

In all cases, if you're having trouble with something, read the software documentation thoroughly, understand how the instructions given there relate to what you've done, and troubleshoot based on that. If that fails, there are mailing lists available for the various pieces of software:

- Exim: <http://www.exim.org/mailman/listinfo/exim-users>
- Exiscan-ACL: <http://duncanthrax.net/exilist/exilist.cgi.pl>
- SA-Exim: <http://lists.merlins.org/lists/listinfo/sa-exim>
- SpamAssassin: <http://spamassassin.org/lists.html>
- Clam Antivirus: <http://www.clamav.net/ml.html>

This document assumes you are using Exim 4. If you are still using Exim 3, you are strongly recommended to upgrade, and I have no idea whether (or how well) any of these techniques will work.

Important Note: Whilst this document does detail (to a certain extent) how to install the software mentioned (since it is necessary to compile and install from scratch), it does assume a basic familiarity with Exim and SMTP. If you are new to Exim, you are strongly recommended to install Exim on its own, and become familiar with how it works, before attempting the projects described in this document. You should also be aware that to achieve the results described you *may* need to be capable of patching and recompiling software – whilst not excessively complicated (especially if you use the RPM packaging discussed), it is not recommended for complete beginners.

3. Mail Scanning Concepts and Options

3.1. Introduction

In these days of ever-increasing junk e-mail, it's increasingly necessary to employ a variety of methods to reduce the amount of unwanted 'spam'. There are many ways of doing this (DNS-based blocking lists, sender verification etc.) but this document looks at one particularly powerful and common requirement: to be able to make decisions on whether or not to accept particular messages based on their content. This does not by any means preclude or supercede the use of other solutions such as DNSBL's; in fact, you'll find the most effective spam-stopping solutions involve a combination of different methods.

There are two primary ways of handling mail scanning:

- a) Accept all messages, then apply some sort of filtering, and bounce any that you don't want
- b) Scan incoming messages as they arrive, and reject at SMTP DATA time

These are explained and discussed below.

3.1.1. Accept all messages, then scan

This is the simplest way, and does not require patching or recompiling Exim. The normal route is to accept the mail and then set up some kind of router which passes the mail to an external scanner (typically SpamAssassin). This has been covered in detail before by many people (including on the exim-users list), and details on how to set it up can be found at the following URL:

http://dman.ddts.net/~dman/config_docs/exim4_spamassassin.html

The main problem with this method is what to do with a mail once it's classified as spam. You could:

- discard it (to /dev/null), but then there's no indication to the sender, which is not good if you get a "false positive" (i.e. legitimate mail is classed as spam).
- bounce it, but the sender is often forged, in which case the bounce will either bounce (and end up frozen in your queue), or go to some innocent bystander (so-called "collateral spam"). **Big problem!**
- move it into a separate mailbox, but then someone has to check it, and the sender has no idea their mail might be delayed (unless you send a notification, but then see the above point about bouncing and it all gets rather clumsy).

None of these are ideal, which is why many people choose the second method, which is to scan at SMTP time.

3.1.2. Scan incoming messages as they arrive (while the SMTP session is open)

This is a more advanced way of rejecting spam and viruses, the main focus of this document. What happens in this case is that after the remote server has connected to your Exim server and sent the message envelope and body, your server scans the complete mail and makes a decision on whether to accept or reject *itbefore* sending the final SMTP confirmation. If it chooses to reject the mail, it can therefore send an SMTP reject code (550) and the problem of what to do with the mail is left entirely to the originating server. Frequently, the "originating server" is not really a mail server at all, but some "spamware" software or perhaps an open proxy, in which case it is unlikely a bounce will be generated. So this method also helps to reduce "collateral spam". This is increasingly important; most common e-mail-borne viruses at the current time fake "From" headers of e-mails, implicating innocent people. Recent virus outbreaks have consequently resulted in blizzards of bounces and "you sent us a virus" warnings to completely innocent parties – it is every server administrator's responsibility to reduce the risk of sending "collateral spam" as much as possible.

This (scan at SMTP time) is the approach that will be documented here. Before you begin, you should be aware that this approach can be costly on resources, and so it isn't recommended in situations where CPU time (in particular) is at a premium. If you're handling high volumes of mail, you're likely to need hefty resources to implement a solution along these lines.

There is, however, one main inherent disadvantage of this method: where spam is concerned you will, largely, lose the ability to provide per-user options on filtering (blacklists/whitelists etc.), because each mail that comes in is scanned only once, and can only be rejected or accepted, but it might have a number of different

recipients all on your server.

There are some 'hybrid' solutions/workarounds you could use to mitigate this limitation, such as:

- a) all mails over a certain (high, e.g. 25) SpamAssassin "spam score" (see later) are discarded at SMTP time, but for any mails that are accepted, you then have customisable per-user filtering after that.
- b) you conditionally scan mails based on one or more recipients having selected scanning (i.e. if one of the recipients wants scanning, everyone gets it whether they like it or not)
- c) using the method outlined by Tom Kistner on <http://www.exim.org/pipermail/exim-users/Week-of-Mon-20030317/051275.html> whereby you allow only one recipient per mail (by temporarily rejecting second and subsequent recipients for a mail at SMTP RCPT time) but this is considered by some to be rather clumsy, and it will delay delivery to the second or subsequent recipients (because the originating server will delay for some period before retrying the mail with the other recipient(s)). A more powerful variant of this, as outlined by Alan Flavell, is to allow multiple recipients per SMTP session, but only if their scanning preferences match. See <http://www.exim.org/pipermail/exim-users/Week-of-Mon-20031006/061151.html>

Either way, you could also use the Exim RCPT ACL to provide a certain degree of per-user-configurable black/whitelisting based on sender address or enable the use of DNSBL's selectively, but that is beyond the scope of this document.

4. What Software To Use?

You can consider that you'll need three types of software (other than Exim itself) to do spam and virus scanning at SMTP time. You'll need:

1. Some kind of Exim patch to handle the mail during transmission and pass it on to the spam and virus scanning software
2. Spam scanning software
3. Virus scanning software

The options are discussed below.

4.1. Exim patch

You have some choice here. There is a 'universal' scanning package for Exim called Exiscan (<http://duncanthrax.net/exiscan-acl/>). This integrates virus, spam and other forms of scanning into Exim. So you can just use this if you want, and this is probably the simplest solution. In any case we'll use this for virus scanning. However, if you want more options and more detailed control of *spam* scanning (only), and/or wish to take advantage of SpamAssassin's `report_safe` option, you might prefer to use SA-Exim (<http://marc.merlins.org/linux/exim/sa.html>). In particular, SA-Exim lets you save rejected messages to a file, which you might find handy if you're jittery about rejecting 'possible spam'. Exiscan and SA-Exim can co-exist peacefully, so you can use Exiscan for virus scanning and SA-Exim for spam scanning if you want.

It's suggested that you read the introductory information given on the websites for the two pieces of software, to decide which you prefer for spam scanning. If in doubt, you can always install both, and then switch between them at leisure by adjusting your configuration.

4.2. Spam scanning software

If you're going to be doing spam scanning, the most important part of the jigsaw is a piece of software which can scan message contents and make some kind of recommendation as to the 'spamminess' of the message. For the purposes of this document, we will use the most common software, SpamAssassin (<http://spamassassin.org/>). You should bear in mind, however, that there are competing solutions (bogofilter, spamprobe etc.) which might be better suited to your needs. In that case, however, these instructions won't help you.

4.3. Virus scanning software

You've got a number of choices here. Exiscan supports a number of virus scanners directly, and supports virtually any scanner (or multiple scanners) indirectly via it's generic command-line scanner configuration.

Some common choices, all of which provide a memory-resident daemon (this is good to ease resource load and generally speed things up) are:

- Clam Antivirus (free software, widely used)
- Sophos Antivirus (commercial; either using the SAVI library in conjunction with a daemon such as 'sophie', or with the commandline-based 'sweep' scanner)
- Kaspersky Antivirus (commercial; has it's own daemon)
- ScannerDaemon (free software; from the OpenAntivirus project – <http://www.openantivirus.org/>)

ClamAV and ScannerDaemon are both free (in both senses of the word) software and can provide a totally free virus scanning solution. Alternatively, you can download a fully-working version of Sophos (with limited trial licence) from the Sophos website (<http://www.sophos.com/>), and I believe some kind of trial is also available for Kaspersky (<http://www.kaspersky.com/>).

Please see the Exiscan documentation for further information about your options when selecting a virus scanner.

This document details how to install either Clam Antivirus or Sophos Antivirus (with the 'sophie' daemon).

5. Installing the software

5.1. Installing SpamAssassin

You can download the SpamAssassin source from <http://spamassassin.org/>. The latest stable version at the time of writing is 2.63, and you're advised to use that (and keep up to date with new releases). If you like living on the edge (and beating the latest spam techniques), you might like to try the latest CVS version.

5.1.1. Installing using RPMs

You can either download a pre-built binary package, if there is a suitable one available for your system, or use the spec file provided as part of SpamAssassin. If you're unsure about how to build an RPM from source, see Appendix A.

NOTE: The spec file supplied with SpamAssassin 2.43 and 2.50 doesn't work properly when building on default installs of Red Hat 7.2. See http://www.hughes-family.org/bugzilla/show_bug.cgi?id=1115 for details and workarounds.

You'll also need some Perl modules that may not be installed by default. Please see the SpamAssassin INSTALL file for full details of the required modules, but for Red Hat 7.3, these may (depending on your installation) include:

- perl-Time-HiRes (available on Red Hat 7.3 CD 3)
- perl-Digest-MD5 (available on Red Hat 7.3 CD 1)
- perl-MIME-Base64 (available on Red Hat 7.3 CD 1)

You may also need some others, such as:

- Digest::SHA1
- Digest::HMAC_MD5
- Test::Simple (ignore the Test::Harness warnings)
- Net::DNS

You can either install these additional modules from CPAN using the 'normal' Perl method (`perl -MCPAN...`), or if (like me), you prefer to install them as RPMs, you can use the 'cpanflute2' utility available from <http://www.freshrpms.net/> (it's called `perl-RPM-Specfile`). Other similar utilities are also available and will probably do the job equally well.

5.1.2. Installing without RPMs

Follow the instructions in the SpamAssassin distribution.

5.1.3. Testing

In either case, assuming you've managed to get SpamAssassin installed, you should now follow the 'testing' instructions in the INSTALL file provided with SpamAssassin to test it, namely:

```
spamassassin -t < sample-nonspam.txt > nonspam.out
spamassassin -t < sample-spam.txt > spam.out
```

(The `sample-spam` and `sample-nonspam` files are provided as part of the SpamAssassin distribution, and can be found in the documentation directory of SA.)

With a bit of luck, all should be working. If not, check you've got all the necessary Perl modules installed (see the INSTALL file) as that's the main cause of problems.

Once it's working, you can move on to installing Exiscan and/or SA-Exim.

5.1.4. Security reminder

Before you start the daemon (we'll come to that later, after configuration), whichever way you installed SpamAssassin, ensure that your firewall is blocking port 783/tcp from the world at large.

5.2. Installing Exiscan

Skip this section if you're only doing spam scanning (not virus scanning), and you've chosen to use SA-Exim to handle things.

Exiscan is a patch to Exim, and integrates heavily with it (for the technically-minded, it keeps clear of `local_scan()`, instead patching Exim to create similar, custom, interface for itself). You'll need a version of Exiscan designed specifically for the version of Exim that you're running, and you can download Exiscan from <http://duncanthrax.net/exiscan-acl/>.

You may be confused by references on the website and mailing lists to Exiscan, because there are two different Exiscan types:

- 'Original' Exiscan – this was the only option until Exim 4.14. It uses custom Exim configuration options starting with `exiscan_*`. This is now deprecated, is not maintained for current Exim releases and unless you have a good reason, you should use Exiscan-ACL for new installations.
- Exiscan-ACL – introduced publically with the release of Exim 4.20, this is an overhauled version of Exiscan which generalises much of the configuration so that it integrates neatly into the Exim ACL system. This is a very slick idea which provides maximum configuration flexibility and familiar rules (if you're used to Exim ACL's).

Just be aware that discussions and other documentation you may find on the Internet might be referring to either version! This document now deals solely with Exiscan-ACL; the last version which discusses 'original' Exiscan is v1.0.8.

Also in beta is a new release of Exiscan (4.30-15 at the time of writing) which has more advanced support for dealing with individual MIME parts. It is not currently covered in this document, but will be in a future release.

5.2.1. Installing as part of an Exim RPM

If you want to compile Exim (with Exiscan included) as an RPM, you can do this in two ways:

- Semi-official RPM packages (both binary and source) are maintained by Nigel Metheringham and are available on the Exim FTP site (<ftp.exim.org>). At the time of writing, the binary packages now include Exiscan by default
- You can use my unified Exim spec file which achieves a similar goal to Nigel's, but is slightly different and also includes SA-Exim support; see the chapter entitled "Compiling Exim as an RPM".

5.2.2. Installing without using RPMs

To install Exiscan without using RPM, untar the source file, and use the "patch" command to patch the Exim source tree. You'll then need to recompile Exim.

5.3. Installing SA-Exim

Skip this section if you're not doing spam scanning, or you have chosen to use Exiscan to handle your spam scanning.

SA-Exim uses Exim's `local_scan` interface to scan messages as they are received, which is what we want to do. You can download it at <http://marc.merlins.org/linux/exim/sa.html>.

SA-Exim can be used in two ways:

- As a simple `local_scan.c` replacement

- As a loadable module, using Derek 'dman' Hudson's `dl_open` patch for Exim

We'll discuss the simple case, which is the former (simple `local_scan` replacement). If you'd like to have SA-Exim as a loadable module using the `dl_open` patch, please see the SA-Exim documentation.

Please note: to avoid confusion, any references to a 'local_scan' config option you may see in mailing list postings imply that the 'dlopen' patch is installed. So, to reiterate, there are two distinct options:

Choice 1 (with dlopen patch)

- Compile Exim with dlopen patch
- Compile SA-Exim as `.so`
- Use "local_scan" option in Exim config file to tell the dlopen patch which `local_scan` modules to load

Choice 2 (without dlopen patch - this is what my RPM packaging does)

- Compile Exim without dlopen patch, but with `sa-exim.c` in the Exim source tree
- That's it. No config options required.

Unless you have plans to plug in other `local_scan` modules, I would personally opt for choice 2 at this stage - it's one less thing to worry about. Again, these instructions assume choice 2.

SA-Exim provides a file to replace the (dummy) `src/local_scan.c` in the default Exim source tree, so you'll need to replace this and recompile Exim. You'll also need `sa-exim.h` (do "make `sa-exim.h`" in the SA-Exim source directory). If you're installing Exim as an RPM, you can do all this automatically with my unified Exim spec file; see the chapter entitled "Compiling Exim as an RPM".

5.4. Installing Clam Antivirus

If you want to install Clam Antivirus as your scanner, you'll need to firstly download the source from www.clamav.net and then compile and install. If you're doing this without using RPMs, follow the instructions given in the Clam AV documentation. If you want to compile it as an RPM, read on.

These instructions apply to ClamAV 0.67, and have changed slightly since 0.5x.

5.4.1. Compiling as an RPM

Some distributions, including PLD, include Clam Antivirus packages, so if you're using those you can probably just go ahead and install the RPM or other package as supplied. If you're using another distribution, you may need to compile it, one way or another. I have an RPM spec file available on my website, to work under Red Hat/Fedora, and you can download it at <http://www.timj.co.uk/linux/rpms.php>.

When you install the RPM, an unprivileged user 'clamav' will be created, which is the user that we will set ClamAV up to run as. You will need to add 'clamav' as an additional user in the 'mailnull' group or whichever group Exim runs as, in `/etc/group`.

5.4.2. Installing without an RPM

Check you have created an unprivileged user 'clamav' for the daemon to run as, and add it as an additional user in the 'mailnull' group (or whatever Exim runs as) in `/etc/group`.

5.4.3. Security reminder

Before you start the daemon (we'll come to that later, after configuration), whichever way you installed ClamAV, ensure that your firewall is blocking port 3310/tcp from the world at large.

5.4.4. Virus database updates

ClamAV includes 'freshclam', a utility to automatically download virus database updates. If you installed as an RPM, this should set it up to run once a day as a cron job (by putting a script in `/etc/cron.daily`). If you

didn't use RPMs, check your installation.

5.5. Installing Sophos Antivirus/sophie

Sophos Antivirus itself consists primarily of three parts:

- Virus database
- Command-line scanner
- libsavi shared library

The virus database itself is of course the most important part and what you are (primarily) paying for. Beyond that, however, you have a choice: you can either use the command line scanner directly, or you can use some third party software (e.g. a daemon), interfaced via the SAVI interface using the libsavi shared library.

In this documentation, I'll focus on using a third party daemon (sophie) with the libsavi library, since this is the most efficient way of scanning, though it's perfectly possible to use the command line scanner with Exiscan – see other documents for details.

The most common daemon used with Sophos is the 'sophie' daemon, which is free software and available from <http://www.vanja.com/tools/sophie/>.

To install sophie without using RPMs, follow the instructions with the package. To install it as an RPM, read on.

5.5.1. Installing as RPM

To install sophie and Sophos Antivirus as RPMs, you can use my spec files found at <http://www.timj.co.uk/linux/>.

5.5.2. Virus database updates

Either way, you may find the sophos-update script from <http://englanders.cc/~jason/patches.php> useful to auto-download updates to the antivirus database.

5.6. Compiling Exim as an RPM

If you're a fan of RPMs, you may wish to compile Exim, together with any associated patches (Exiscan and/or SA-Exim), as an RPM. To do so, there are a variety of spec files available. One obvious choice is the spec file/SRPM provided by Nigel Metheringham on ftp.exim.org. However, I'll focus here on the spec file I have created, which includes (optional) support for Exiscan-ACL and SA-Exim.

Before you start, you'll need to have decided whether you're using SA-Exim, Exiscan or both, since both of them are effectively compiled into Exim and therefore form part of the Exim RPM. (N.B. this needn't necessarily be strictly true if you want to install SA-Exim as a shared object using the dl_open patch (which allows 'pluggable' local_scan objects), but I'm not covering that here).

You can download my 'universal' spec file from <http://www.timj.co.uk/linux/exim.php>.

If you're going to use SA-Exim, depending on what link you followed to download it, you may (or may not) need to rename your SA-Exim downloaded file to "sa-exim-x.y.tar.gz" (where x.y is the version number). This isn't normal practice when building RPMs from spec files, but in this case SA-Exim downloads from Marc Merlin's front page don't always have a version number in the filename which can be confusing when you've got multiple versions hanging around, so my spec file specifies the use of a version number in the filename.

If you have no idea what to do with a spec file, first read Appendix A to find out how to build an RPM from spec and sources.

5.6.1. Setting up the Makefile

The first thing you need to do is to open for editing the downloaded file "exim-4.xx.Makefile", (where xx is the Exim minor revision you're using), or just take the default one from the Exim source distribution and rename it to exim-4.xx.Makefile. (The version supplied on my site is a fairly simple and 'clean' base configuration, with just a few simple changes from the default to make it match standard settings on Red Hat and include a couple of commonly used options).

Now edit the Makefile to suit your needs, referring to the comments in it and the Exim documentation where necessary

5.6.2. Setting the spec file build options

You now need to set some options in the spec file based on whether you are going to build SA-Exim and/or Exiscan, and whether you opted to include Eximon in the Makefile. Follow the instructions in the spec file, namely:

```
# Set this to 1 if you have enabled eximon in your Makefile
%define eximon 0

# Set this to 1 if you want to compile in the SA-Exim patch
%define build_saexim 1

# Set this to 1 if you want to compile in Exiscan-ACL (http://duncanthrax.net/exiscan-acl/)
%define build_exiscan_acl 1
```

You should ensure that the downloaded source files for Exiscan and/or SA-Exim are in place in your SOURCES directory, then ensure the versions in the spec file correspond:

```
# Versions of SA-Exim and Exiscan to use
%define sa_exim_ver 3.1
%define exiscan_rel 14
```

Now issue "rpmbuild -ba exim4xx.spec" and, if all goes well, you should have an RPM of Exim with your selected options compiled in, in the 'RPMS' directory of your RPM build tree. (A source RPM, containing all source files plus the spec will be in the 'SRPMS' directory).

6. Configuring the software

Before you start, bear in mind that what I give here is very much a 'quick start' guide to important/common options, and is **not a substitute for reading the documentation** for the various pieces of software. The options provided are examples and the appropriate settings may vary significantly depending on your particular environment.

6.1. Configuring SpamAssassin

SpamAssassin itself works pretty well 'out of the box', although you'll need to configure Exiscan/SA-Exim properly.

6.1.1. For Exiscan users

There is one option – `report_safe` - you'll probably need/want to change if you're using Exiscan-ACL. You do this in the local config file, (typically `/etc/mail/spamassassin/local.cf`):

```
report_safe 0
```

(This option changes the message body of suspected spam by writing the SpamAssassin report in the body and attaching the original message as a MIME part. Some people like this, some people don't. It is supported by SA-Exim v3.1+ but not currently by Exiscan-ACL)

6.1.2. For SA-Exim users

You may want to change the `rewrite_subject` option – it defaults to 0, but if set to 1 adds “**** SPAM...” at the start of the subject line for messages found to be spam. Otherwise, the spam reports will only be found in special message headers (e.g. X-Spam-Flag, X-Spam-Status, X-Spam-Level, X-Spam-Report). In an ideal world this would be unnecessary, but unfortunately many common but useless clients such as Micros*** Lookout and Lookout Express can't filter on non-standard headers, so you may need to set this option to 1 if those users want to be able to filter 'possible spam' into a separate folder.

NOTE: this setting is ignored and you need to configure this separately if you're using Exiscan-ACL!

6.1.3. For all users

SpamAssassin in general is very configurable, so it's worth reading the manual to find out settings which you might want to customise for your site. However, one key setting to bear in mind is **required_hits** – default 5.0, this sets the number of points which a mail must score to be *marked* as spam. Note that this *isn't* the score at which mail will be rejected outright; that is set separately in SA-Exim or Exiscan, depending on which you are using.

You should also note that there are many more useful things you can put in the global configuration file, including black/white lists and your own rules.

You should ensure that the SpamAssassin daemon is started using the options `'-d -a'`, *not* the `'-c'` option. If you installed from RPM, the easiest way to do this and make sure it happens each time you start SpamAssassin is by creating the file `/etc/sysconfig/spamd` and putting this line in it:

```
OPTIONS="-d -a"
```

Furthermore, if you want to take advantage of SpamAssassin's auto-whitelist and Bayesian auto-learning features (please read the SpamAssassin documentation for more information about these), you'll probably need to set up the directory where this is stored. Find out the UNIX user that Exim runs as (possibly 'mail', uid 12) and locate the home directory for that user (from `/etc/passwd`). This is typically `/var/spool/mail`. Then, in the home directory, do:

```
mkdir .spamassassin
chown mail.mail .spamassassin
chmod 700 .spamassassin
```

Substitute the Exim user/group for “mail.mail” if it's not mail/mail. SpamAssassin's auto-whitelist database file

will be stored in this directory once you start using it.

6.2. Configuring Exiscan

You need to enter some Exiscan options into your Exim config file (normally `/etc/exim/exim.conf` if you installed using my RPM specs). This guide now only refers to Exiscan-ACL; if you're using the 'original' Exiscan, either upgrade or read v1.0.8 of this guide.

6.2.1. Overview/setting up DATA ACL

Exiscan-ACL requires you to set up a DATA ACL and insert rules in there. If you haven't already, define an ACL to be used (I'll call it 'acl_check_data' here) in the main part of your Exim config:

```
acl_smtp_data = acl_check_data
```

and, down in the ACL section of your config file (somewhere after the "begin acl" line), start that ACL:

```
acl_check_data:
```

Now, the beauty of Exiscan-ACL is that you can customise the exact rules to suit your own circumstances. However, that means that there are an infinite number of possibilities for configuring it. I'll outline here some common rules to get you started, but you're going to need to make sure you understand how ACLs work if you want to go beyond these (and it's worth doing, because you'll appreciate the amazing flexibility of Exim in general!)

6.2.2. MIME checking settings

You can configure Exiscan to check whether a message contains invalid MIME. Exiscan can give you a number indicating the severity of MIME errors in the message and you can use this to reject badly malformed messages (which usually indicates malicious content). For example:

```
deny    message      = This message contains malformed MIME ($demime_reason).
        demime       = *
        condition    = ${if >{$demime_errorlevel}{2}{1}{0}}
```

This will reject messages with severe (\geq level 2) MIME errors.

6.2.3. File attachment blocking settings

It's generally considered prudent to block certain types of file attachment which are rarely used 'for real' but are often used as virus carriers, as a first line of defence. This is also an extremely effective way of stopping new viruses or variants which haven't yet made it into your virus scanner's database. Such attachments are identified by their file extension and a simple subset to block might be `.bat/.com/.exe/.pif/.scr/.vbs`.

To set up Exiscan to block the above extensions, use something like the following:

```
deny    message      = This message contains an attachment of a type which we do not
accept  (.$found_extension)
        demime       = bat:com:exe:pif:prf:scr:vbs
```

6.2.4. Spam scanning settings

Skip this if you're not doing spam scanning, or are doing it with SA-Exim

First, you need to define your spam scanner's TCP socket location. Assuming it's on the local machine and listening on port 783, add the following option in the *main part of the Exim config (not the DATA ACL!)*:

```
spamd_address = 127.0.0.1 783
```

Back in the DATA ACL, there are a wide variety of options available for spam scanning. Let's assume you want the following:

- All mails (whether or not they are spam) should have an X-Spam-Score header , containing the SpamAssassin score for that mail
- All mails (whether or not they are spam) should have an X-Spam-Report header, containing SpamAssassin's report on the mail

Assuming that's the case, add your first spam ACL rule in the DATA ACL:

```
warn message = X-Spam-Score: $spam_score
      spam = nobody:true
warn message = X-Spam-Report: $spam_report
      spam = nobody:true
```

If you want to add a "**** SPAM ****" marker to the start of subject lines if a mail contains spam (see the discussion above about `rewrite_subject` in the SpamAssassin configuration for more info) then add this ACL rule:

```
warn message = Subject: **** SPAM **** $h_Subject
      spam = nobody
```

Now, to do our SMTP-time rejection, we need to add a 'deny' ACL rule. Decide the score at which you wish to reject and multiply it by 10 to find the rejection threshold. Assuming you choose 12 (a reasonable starting value), you should add a rule similar to the one below: (the '120' is the rejection threshold as calculated above)

```
deny message = This message scored $spam_score spam points.
      spam = nobody:true
      condition = ${if >{$spam_score_int}{120}{1}{0}}
```

6.2.5. Virus scanning settings

Skip this if you're not doing virus scanning

You need to configure Exiscan to pass messages to your virus scanner. First, in the *main part of the Exim config (not the DATA ACL!)* define your virus scanner using the `av_scanner` option. The Exiscan manual details how to use this, but here are common settings:

For Clam Antivirus:

```
av_scanner = clamd:127.0.0.1 3310
(You could also use clamd's UNIX socket – use: av_scanner = clamd:/path/to/socket).
```

For Sophos with sophie:

```
av_scanner = sophie:/tmp/sophie
```

Now you need to set up your DATA ACL rule to do the virus scanning. Something such as the following rule should work well:

```
deny message = This message contains a virus or other harmful content
($malware_name)
      demime = *
      malware = *
```

The 'demime' line here is only necessary if you want Exiscan to unpack (separate out) the different parts of MIME messages (that is, messages with attachments) before passing to your virus scanner. If your virus scanner is MIME-aware (Sophos is, Clam AV isn't), you may not need to do this but it's probably a good idea to give it a go anyway - "belt and braces".

6.2.6. Ending the ACL

Finally, close the `acl_check_data` ACL with a final 'accept' line:

```
accept
```

so that if mails don't contain viruses/spam, they will be accepted. Don't forget this, or you will end up rejecting everything!

6.3. Configuring SA-Exim

SA-Exim is a highly configurable piece of software, and has an exceptionally well-commented config file. The exact options that you want will very much depend on your personal preferences and environment. However, I'll give you some pointers here towards options you may typically want to look at.

6.3.1. The SA-Exim Config File

The SA-Exim config file can typically be found residing at `/etc/exim/sa-exim.conf`

The most important 'basic' options are:

- `SAEximRunCond` – equivalent to Exiscan's ACL 'condition', this determines when and if SA-Exim will scan a message. You could just set it to '1' to scan all messages, but there are some circumstances when you might not want to – see below.
- `SAEximRejCond` – this determines whether mails are actually rejected if they are found to be spam. You don't really want to reject messages addressed to `postmaster@yourdomain/abuse@yourdomain`, so the default configuration for this option checks some flags which can be set in an Exim ACL to avoid this – see below.
- `SAPermreject` – this is the 'threshold' spam score. If a message scores higher than this, it will be rejected. You might be a bit nervous about this setting, so if you like you might set it higher and slowly reduce it. 12.0 (the default) is unlikely to catch any real mail, though, and in fact you can probably safely pull it down to 11.0 or even less without rejecting real mail. Experiment!
- `SAPermrejectSavCond` – this determines whether messages that have been rejected as spam are saved anywhere. If you want to save them, set this to 1 (or use some kind of condition) and then set `SAPermrejectsave` to the location where you want to save messages.
- `SAMsgpermrej` – this is the message used to reject spam with if it exceeds the threshold. If a real mail is inadvertently bounced as spam, the sender should see this message in the bounce which they get.

6.3.2. Setting Exim ACLs

To avoid rejecting messages to `postmaster@/abuse@`, you should add the following ACL part to Exim's RCPT ACL (in the Exim config file), somewhere near the top:

```
warn      message      = X-SA-Do-Not-Rej: Yes
          local_parts  = postmaster:abuse
```

This will add a header "X-SA-Do-Not-Rej: Yes" to any messages addressed to `postmaster@/abuse@`, and then (according to the default SA-Exim config), suppress rejection of spam to these addresses. Messages that are classed as spam will still receive the spam headers, however, which is a good compromise – the mail can still be filtered into a separate folder in your (or your users') mail client, but in the event that it's actually a genuine mail (typically an abuse complaint), you will have the opportunity to review this folder and read it.

6.4. Configuring Clam Antivirus

You'll probably need/want to change quite a few of the default settings for Clam Antivirus. I've set up a config file which works well with Exiscan. If you use my RPM packaging as supplied, this will be used by default and you should find ClamAV is ready to go. If not, you can download my ClamAV config file at <http://www.timj.co.uk/linux/sources/clamav.conf>.

First, just to check, make sure you followed the instructions in the earlier section about installing Clam Antivirus, particularly with respect to making sure you created a 'clamav' user, and this is an additional user in the 'mailnull' group, or whichever group Exim runs as.

For reference, the key changes are outlined below:

1. Commented out the line which has "Example" on it.
2. Set the LogFile to `/var/log/clamav/clamd.log`
3. LogVerbose and LogTime enabled. You'll find these useful if trying to track down problems.
4. Set PidFile `/var/run/clamav/clamd.pid`
5. Set "DataDirectory `/usr/share/clamav`" as this is where the RPM normally installs the virus database files.

Check, if you're not installing from RPM.

6. Comment out the "LocalSocket" option. These instructions assume interfacing to ClamAV over a TCP socket, though Exiscan now supports interfacing over a UNIX socket too.
7. Enable the "TCPsocket 3310" option (see 4 above)
8. Enabled the TCPAddr 127.0.0.1 option to only allow connections from the local host
9. Enabled the StreamSaveToDisk and "StreamMaxLength 10M" options
10. Set the "MaxThreads" option to 10
11. Enable the "User clamav" option – better not to run Clam Antivirus as root.
12. Enable the "AllowSupplementaryGroups" option, to give ClamAV the required access to the temporary scan files created by Exiscan.
13. Enable the 'ScanMail' option. (NOTE: this was omitted in a previous version of this document but is important to ensure detection of some recent viruses)

7. Getting it all running/Testing

Assuming you've followed this document through, you should now be ready to start your chosen combination of software and get things working.

7.1. Starting the daemons

7.1.1. SpamAssassin

On many systems, including Red Hat-alike ones where you install SA as an RPM:

- Execute "service spamassassin start". You should then have a process called 'spamd' running, hopefully with the '-d' and '-a' options enabled.
- Execute "chkconfig spamassassin on" just to make sure it will start on reboot.

otherwise:

- Execute "spamd -d -a" and arrange for this to happen on bootup.

7.1.2. Clam Antivirus

If you installed as an RPM:

- Execute "service clamd start". You should then have a process called "clamd" running.
- Execute "chkconfig clamd on" just to make sure it will start on reboot.

otherwise:

- Just run 'clamd'

7.1.3. Sophie

If you installed as an RPM:

- Execute "service sophie start". You should then have a process called "sophie" running.
- Execute "chkconfig sophie on" just to make sure it will start on reboot.

otherwise:

- read the Sophie documentation

7.1.4. Exim

Start Exim in your usual way (typically "service exim start" if you installed using an RPM)

7.2. Testing

7.2.1. Spam scanning

The first test is to check that your spam scanning is working. Assuming that you've followed the instructions here, there are three possibilities for what happens to a mail when it comes in:

- It scores less than SpamAssassin's *required_hits* value, and is accepted as normal
- It scores more than SpamAssassin's *required_hits* value, but less than the rejection threshold set in SA-Exim/Exiscan, in which case it is marked as spam (either in the headers, or in the subject, depending on how you configured SpamAssassin) but accepted
- It scores more than the rejection threshold set in SA-Exim/Exiscan, and is rejected at DATA time.

You therefore need to test these three circumstances. Before you do this, if you have virus scanning enabled via Exiscan, temporarily disable it by commenting out the relevant ACL lines, just so we can eliminate one source of possible problems.

First, telnet to your mailserver on port 25 and send yourself a plain 'test' e-mail:

```
telnet your.server 25
[Wait for connection and the server banner, e.g. "220 your.server.name ESMTP Exim 4.30 ..."]
HELO test
[Server should respond: 250...]
MAIL FROM: <you@your.address>
[Server should respond: 250...]
RCPT TO: <you@your.address>
[Server should respond: 250...]
DATA
[Server should respond: 354...]
From: you@your.address
To: you@your.address
Subject: a test mail

test
.
```

[Server should respond: 250, possibly after a short pause whilst scanning takes place]

With a bit of luck, you should now have a message in your mailbox. Look at the source. It should have one or two additional headers (varying according to your software, options and where you send the mail from), possibly something like this:

```
X-Spam-Status: No, hits=1.1 required=5.0
                tests=NO_REAL_NAME, MSGID_FROM_MTA_SHORT
                autolearn=no version=2.60
X-Spam-Level: *
```

If so, excellent! If not, or you got a temporary error (4xx) when you sent your mail, check your Exim main and reject logs to find out what happened.

Assuming that worked, now try sending a mail that's a bit spammy. Telnet to your server again, as before, but use something like this as the mail:

```
From: test@test
Subject: MAKE MONEY FAST!!! $$$          34fs4

viagra
```

The mail should be accepted, and if you look at the headers when you receive it, you should find something like this (again, will vary depending on whether you are using SA-Exim or Exiscan):

```
X-Spam-Status: Yes, hits=6.0 required=5.0
                tests=CASHCASHCASH,MISSING_HEADERS,NO_REAL_NAME,PLING_PLING,
                SUBJ_HAS_SPACES,SUBJ_HAS_UNIQ_ID
                autolearn=no version=2.60
X-Spam-Level: *****
X-Spam-Checker-Version: SpamAssassin 2.60 (1.212-2003-09-23-exp) on your.server.example
X-Spam-Report: ---- Start SpamAssassin results
6.00 points, 5 required;
* 0.8 -- From: does not include a real name
* 1.7 -- Subject contains lots of white space
* 1.3 -- Subject contains a unique ID
* 0.5 -- Missing To: header
* 0.0 -- Contains at least 3 dollar signs in a row
* 1.7 -- Subject has lots of exclamation marks
---- End of SpamAssassin results
X-Spam-Flag: YES
```

You'll notice the X-Spam-Flag header, which can be useful to filter "possible spam" into a separate folder in your mail client. (The X-Spam-Level header can also be useful for this; it contains a number of asterisks corresponding to the spam rating of the mail).

If you are using SA-Exim, you may also get additional headers useful for debugging such as:

```
X-SA-Exim-Mail-From: you@your.address
X-SA-Exim-Version: 3.1 (built Sat Oct 4 09:02:58 BST 2003)
X-SA-Exim-Scanned: Yes
```

If you don't like these headers, you can get rid of them – see Appendix B.

Finally, let's try to get a message rejected. Telnet in again, but use this as your message:

```
From: test@test
Subject: MAKE MONEY FAST!!! $$$           34fs4
MIME-Version: 1.0
```

```
VIAGRA!!!!
AS SEEN ON NATIONAL TV, WORK FROM HOME!!!
GUARANTEED 100% THIS IS NOT SPAM!!!
```

This should get rejected, so instead of a final 250 code from the server, you should receive a 550. The message is rejected, and won't end up in your mailbox. If you're using SA-Exim and you've enabled saving of rejected messages, you should be able to see it in the save directory.

Note that there is also a 'magic' test phrase (called 'GTUBE') which, if used in a mail, will trigger SpamAssassin to reject it (rather like the widely-used 'EICAR' test string for virus engines:

```
/XJS\*C4JDBQADN1\ .NSBN3\*2IDNEN\*GTUBE-STANDARD-ANTI-UBE-TEST-EMAIL\*C\ .34X/
```

7.2.2. Virus scanning

Assuming your spam scanning is now working (if you're doing it), it's time to test virus scanning. Re-enable it if you temporarily disabled it earlier.

Before you start, you'll need something to test for. A good test is the Eicar test signature, which is a 'pretend' virus recognised by most anti-virus software. You can download it from http://www.eicar.org/anti_virus_test_file.htm, or just paste the next line into a file (this is the 'virus'):

```
X50!P%@AP[4\PZX54(P^)7CC)7}$EICAR-STANDARD-ANTIVIRUS-TEST-FILE!$H+H*
```

First, try sending yourself the file by telnet. Telnet in as before, and send the above line after DATA. Hopefully, you'll get a response something like this:

```
550 This message contains a virus or other harmful content (Eicar-Test-Signature).
```

If so, great! Looks good. Now use an e-mail client to send yourself viruses (through your server) using different forms such as enclosing in a ZIP file etc. It should all be detected and rejected.

7.2.2.1. Manual testing with Clam Antivirus

If you're having problems, and are using Clam Antivirus, you can test if it's working generally by placing a file somewhere and telnetting to port 3310 on the machine running Clam AV. Type "SCAN /path/to/eicar.com" and hopefully you'll get a response something like "/path/to/eicar.com: Eicar-Test-Signature FOUND". If not, you'll need to dig deeper. If you do, then try putting the Eicar file in /var/spool/exim/scan and scanning it then. If that fails, you've got a permissions problem.

8. Appendix A – Building RPMs from source

If you want to build an RPM from source, you need three things: the source to the program (typically a .tar.gz file), an RPM 'specification' (spec) file, and (possibly) any ancillary external scripts/patches which are needed to build the software.

The command to build an RPM is “rpmbuild” on recent versions of RPM, or “rpm” on older versions. You may need to install the package “rpm-build” from your distribution, if it isn't already installed.

There are two primary ways to build an RPM from source:

8.1. Using tarballs (the quick way)

If the piece of software you're using provides a spec file within its tarball, in the standard way, and this is all OK and no external files are needed, you may be able to build an RPM directly from the tarball, like so:

```
rpmbuild -tb <file>.tar.gz
```

Often, this isn't the case, and you'll need to use a build tree.

8.2. Using a build tree (the conventional way)

To build things the 'proper' (conventional) way, you'll first need an RPM build tree. Typically, you'll have an example one installed which is writeable by the root user, in /usr/src/redhat. It's not recommended to build things as root, however, so it's best to reproduce this tree somewhere else and make it writeable by an unprivileged user. Then, in the home directory of the unprivileged user you're going to use to compile things, create an “.rpmmacros” file, containing a single line with this in: “%_topdir /path/to/your/rpm/build/tree”.

Now, to build an RPM, you need to first place the program sources and any ancillary files in the SOURCES directory of your RPM build tree. Next, put the spec file in the SPECS directory of your RPM build tree and change into that directory. Then type:

```
rpmbuild -ba <specfile>
```

where <specfile> is the name of the spec file for the program that you want to build. Remember this may be “rpm -ba <specfile>” on older systems.

If all goes well and there are no errors, the finished RPM should be in the RPMS directory of your build tree, in the i386 directory if you're building on an i386 machine (or possibly the 'noarch' directory if it's non-native; not the case for Exim). You'll also have a source RPM, which includes all the source files and the spec file itself, in the SRPMS directory of your build tree.

8.3. More Information

For more information about RPMs/spec files etc., please visit <http://www.rpm.org/>.

9. Appendix B – Removing headers

This advice applies generally if you wish to remove specific headers from an e-mail, but is mainly applicable if you wish to remove headers such as SA-Exim's X-SA-Exim-Mail-From header.

The best way is to add a system filter rule. Outlined below are methods for if you do or don't already have a system filter. These assume you just want to blanket remove headers from all e-mails; you could use conditions etc. to remove them from only certain mails.

9.1. If you don't already have a system filter

First, you need to create a filter file. You can put this wherever you like, but a good place might be /etc/exim/system.filter. Put the following lines in it (note in particular the '#Exim filter' line which is crucial):

```
# Exim filter
headers remove X-SA-Exim-Scanned:X-SA-Exim-Mail-From
```

The above example removes the X-SA-Exim-Scanned and X-SA-Mail-From headers; adjust to your taste.

Now you need to tell Exim to use your filter as the 'system filter'. Add this to the main (top) part of your Exim configuration file:

```
system_filter = /etc/exim/system.filter
```

You should now find that messages passed through your system will have the headers you defined in the system.filter file removed.

9.2. If you already have a system filter

Just add the following line to the start of your filter (the following example removes X-SA-Exim-Scanned and X-SA-Exim-Mail-From; adjust to your taste).

```
headers remove X-SA-Exim-Scanned:X-SA-Exim-Mail-From
```

10. Appendix C – The GNU Free Documentation License

Version 1.2, November 2002

Copyright (C) 2000,2001,2002 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the

copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- **A.** Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- **B.** List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- **C.** State on the Title page the name of the publisher of the Modified Version, as the publisher.
- **D.** Preserve all the copyright notices of the Document.
- **E.** Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- **F.** Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- **G.** Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- **H.** Include an unaltered copy of this License.
- **I.** Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- **J.** Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- **K.** For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- **L.** Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- **M.** Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- **N.** Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- **O.** Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.