Slackware Linux Network Administration Guide

For Slackware Linux 9.1

Daniël de Kok

Slackware Linux Network Administration Guide: For Slackware Linux 9.1
by Daniël de Kok

Table of Contents

List of Tables

List of Figures

Chapter 1. About this book

Introduction

The aim of this book is to provide a good introduction to the administation of network functionality in Slackware Linux. Some basic knowledge about Slackware Linux and GNU/Linux in general is required to read most parts of this book. One of the primary goals is to describe software that is included in Slackware by default. This means that having this book and a Slackware Linux CD set should be enough to build a server, router or networked workstation.

The following tools were used to write this book:

- vi (elvis) for editing the XML files
- xfig for drawing images
- jade for converting DocBook/XML to other formats
- LaTeX and ghostscript to create PS and PDF output

Chapter 2. Basic networking configuration

Hardware

Network cards (NICs)

Drivers for NICs are installed as kernel modules. The module for your NIC has to be loaded during the initialization of Slackware. On most systems the NIC is automatically detected and configured during the installation of Slackware Linux. You can reconfigure your NIC with the **netconfig** command. **netconfig** adds the driver (module) for the detected card to `/etc/rc.d/rc.netdevice`.

It is also possible to manually configure which modules should be loaded during the initialization of the system. This can be done by adding a **modprobe** line to `/etc/rc.d/rc.modules`. For example, if you want to load the module for 3Com 59x NICs (3c59x.o), add the following line to `/etc/rc.d/rc.modules`

```
/sbin/modprobe 3c59x
```

PCMCIA cards

Supported PCMCIA cards are detected automatically by the PCMCIA software. The pcmcia-cs packages from the "a" diskset provides PCMCIA functionality for Slackware Linux.

Configuration of interfaces

Network cards are available under Linux through so-called "interfaces". The **ifconfig**(8) command can be used to display the available interfaces:

```
# ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:20:AF:F6:D4:AD
          inet addr:192.168.1.1  Bcast:192.168.1.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:1301 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1529 errors:0 dropped:0 overruns:0 carrier:0
          collisions:1 txqueuelen:100
          RX bytes:472116 (461.0 Kb)  TX bytes:280355 (273.7 Kb)
          Interrupt:10 Base address:0xdc00

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:77 errors:0 dropped:0 overruns:0 frame:0
          TX packets:77 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:8482 (8.2 Kb)  TX bytes:8482 (8.2 Kb)
```

Network cards get the name ethn, in which n is a number, starting with 0. In the example above, the first network card (eth0) already has an IP address. But unconfigured interfaces have no IP address, the **ifconfig** will not show IP addresses for unconfigured interfaces. Interfaces can be configured in the `/etc/rc.d/rc.inet1.conf` file. You can simply read the comments, and fill in the required information. For example:

```
# Config information for eth0:
IPADDR[0]="192.168.1.1"
NETMASK[0]="255.255.255.0"
USE_DHCP[0]=""
DHCP_HOSTNAME[0]=""
```

In this example the IP address 192.168.1.1 with the 255.255.255.0 netmask is assigned to the first ethernet interface (eth0). If you are using a DHCP server you can change the *USE_DHCP=""* line to *USE_DHP[n]="yes"* (swap "n" with the interface number). Other variables, except *DHCP_HOSTNAME* are ignored when using DHCP. For example:

```
IPADDR[1]=""
NETMASK[1]=""
USE_DHCP[1]="yes"
DHCP_HOSTNAME[1]=""
```

The same applies to other interfaces. You can activate the new settings by rebooting the system or by executing **/etc/rc.d/rc.inet1**.

Configuration of interfaces (IPv6)

Introduction

IPv6 is the next generation internet protocol. One of the advantages is that it has a much larger address space. In IPv4 (the internet protocol that is commonly used today) addresses are 32-bit, this address space is almost completely used right now, and there is a lack of IPv4 addresses. IPv6 uses 128-bit addresses, which provides an unimaginable huge address space (2^128 addresses). IPv6 uses another address notation, first of all hex numbers are used instead of decimal numbers, and the address is noted in pairs of 16-bits, separated by a colon (":"). Let's have a look at an example address:

```
fec0:ffff:a300:2312:0:0:0:1
```

A block of zeroes can be replaced by two colons ("::"). Thus, thee address above can be written as:

```
fec0:ffff:a300:2312::1
```

Each IPv6 address has a prefix. Normally this consists of two elements: 32 bits identifying the address space the provider provides you, and a 16-bit number that specifies the network. These two elements form the prefix, and in this case the prefixlength is 32 + 16 = 48 bits. Thus, if you have a /48 prefix you can make 2^16 subnets and have 2^80 hosts on each subnet. The image below shows the structure of an IPv6 address with a 48-bit prefix.

Figure 2-1. The anatomy of an IPv6 address

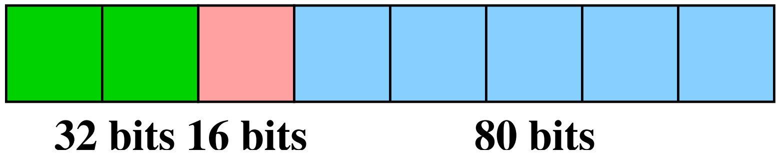

There are a some specially reserved prefixes, most notable include:

Table 2-1. Important IPv6 Prefixes

Prefix	Description
fe80::	Link local addresses, which are not routed.
fec0::	Site local addresses, which are locally routed, but not on or to the internet.
2002::	6to4 addresses, which are used for the transition from IPv4 to IPv6.

Slackware IPv6 support

The Linux kernel binaries included in Slackware Linux do not support IPv6 by default, but support is included as a kernel module. This module can be loaded using **modprobe**:

```
# modprobe ipv6
```

You can verify if IPv6 support is loaded correctly by looking at the kernel output using the **dmesg**:

```
$ dmesg
[..]
IPv6 v0.8 for NET4.0
```

IPv6 support can be enabled permanently by adding the following line to `/etc/rc.d/rc.modules`:

```
/sbin/modprobe ipv6
```

Interfaces can be configured using **ifconfig**. But it is recommended to make IPv6 settings using the **ip** command, which is part of the "iputils" package that can be found in the `extra/` directory of the Slackware tree.

Adding an IPv6 address to an interface

If there are any router advertisers on a network there is a chance that the interfaces on that network already received an IPv6 address when the IPv6 kernel support was loaded. If this is not the case an IPv6 address can be added to an interface using the **ip** utility. Suppose we want to add the address "fec0:0:0:bebe::1" with a prefix length of 64 (meaning "fec0:0:0:bebe" is the prefix). This can be done with the following command syntax:

```
# ip -6 addr add <ip6addr>/<prefixlen> dev <device>
```

For example:

```
# ip -6 addr add fec0:0:0:bebe::1/64 dev eth0
```

Resolving

Hostname

Each computer on the internet has a hostname. If you do not have a hostname that is resolvable with DNS, it is still a good idea to configure your hostname, because some software uses it. You can configure the hostname in `/etc/HOSTNAME`. A single line with the hostname of the machine will suffice. Normally a hostname has the following form: host.domain.tld, for example darkstar.slackfans.org. Be aware that the hostname has to be resolvable, meaning that Linux should be able to convert the hostname to an IP address. You can make sure the hostname is resolvable by adding it to `/etc/hosts`. Read the following section for more information about this file.

/etc/hosts

`/etc/hosts` is a table of IP addresses with associated hostnames. This file can be used to name computers in a small network. Let's look at an example of the `/etc/hosts` file:

```
127.0.0.1        localhost
192.168.1.1             tazzy.slackfans.org tazzy
192.168.1.2             gideon.slackfans.org
```

The localhost line should always be present. It assigns the name "localhost" to a special interface, the loopback. In this example the names "tazzy.slackfans.org" and "tazzy" are assigned to the IP address 192.168.1.1, and the name "gideon.slackfans.org" is assigned to the IP address 192.168.1.2. On the system with this file both computers are available via the mentioned hostnames.

It is also possible to add IPv6 addresses, which will be used if your system is configured for IPv6. This is an example of a `/etc/hosts` file with IPv4 and IPv6 entries:

```
# IPv4 entries
127.0.0.1               localhost
192.168.1.1             tazzy.slackfans.org tazzy
192.168.1.2             gideon.slackfans.org

# IPv6 entries
::1   localhost
fec0:0:0:bebe::2 flux.slackfans.org
```

Please note that "::1" is the default IPv6 loopback.

/etc/resolv.conf

The `/etc/resolv.conf` file is used to specify which nameservers the system should use. A nameserver converts hostnames to IP addresses. Your provider should have given you at least two name name server addresses (DNS servers). You can add these nameservers to `/etc/resolv.conf` by adding the line "nameserver ipaddress" for each nameserver. For example:

```
nameserver 192.168.1.1
nameserver 192.168.1.69
```

You can check wether the hostnames are tranlated correctly or not with the **host hostname** command. Swap "hostname" with an existing hostname, for example the website of your internet service provider.

Chapter 3. Routing

IPv4 Forwarding

IPv4 forwarding connects two or more networks by sending packets which arrive on one interface to another interface. This makes it possible to let a Linux machine act as a router. For example, you can connect multiple networks, or your home network with the internet. Let's have a look at an example:

Figure 3-1. Router example

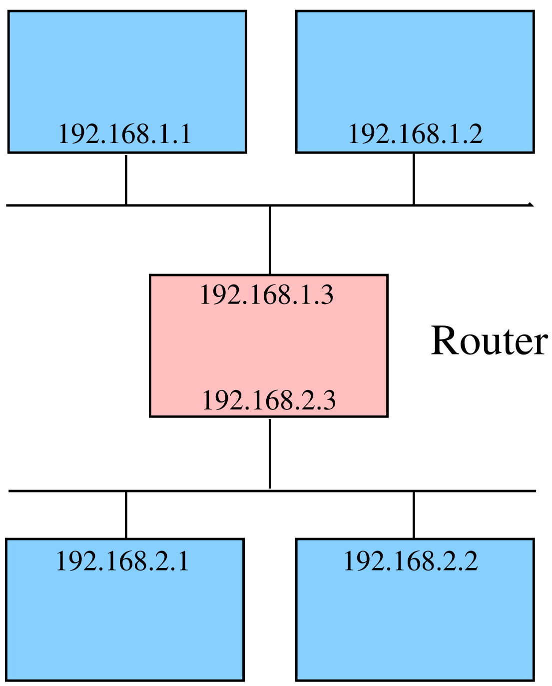

In dit example there are two networks, 192.168.1.0 and 192.168.2.0. Three hosts are connected to both network. One of these hosts is connected to both networks with interfaces. The interface on the 192.168.1.0 network has IP address 192.168.1.3, the interface on the 192.168.2.0 network has IP address 192.168.2.3. If the host acts as a router between both networks it forwards packets from the 192.168.1.0 network to the 192.168.2.0 network and vise versa. Routing of normal IPv4 TCP/IP packages can be enabled by enabling IPv4 forwarding.

IPv4 forwarding can be enabled or disabled under Slackware Linux by changing the IPV4_FORWARD variable in `/etc/rc.d/rc.inet2`. The default setting is as follows:

```
IPV4_FORWARD=1
```

This means that IPv4 forwarding is enabled. You can disable forwarding by changing the IPV4_FORWARD value to 0. This setting can be enabled by rebooting the computer. It is also possible to enable IPv4 forwarding on a running system with the following command (0 disables forwarding, 1 enables forwarding):

```
# echo 0 > /proc/sys/net/ipv4/ip_forward
```

Be cautious! By default there are no active packet filters. This means that anyone can access other networks. Traffic can be filtered and logged with the iptables kernel packet filter. Iptables can be administrated through the **iptables** command. NAT (Network Address Translation) is also a subset of iptables, and can be controlled and enabled through the **iptables** command. NAT makes it possible to "hide" a network behind one IP address. This allows you to use the internet on a complete network with only one IP address.

Chapter 4. The internet super server

Introduction

There are two ways to offer TCP/IP services: by running server applications standalone as a daemon or by using the internet super server, **inetd**(8). **inetd** is a daemon which monitors a range of ports. If a client attempts to connect to a port **inetd** handles the connection and forwards the connection to the server software which handles that kind of connection. The advantage of this approach is that it adds an extra layer of security and it makes it easier to log incoming connections. The disadvantage is that it is somewhat slower than using a standalone daemon. It is thus a good idea to run a standalone daemon on, for example, a heavily loaded FTP server.

Configuration

inetd can be configured using the `/etc/inetd.conf` file. Let's have a look at an example line from `inetd.conf`:

```
# File Transfer Protocol (FTP) server:
ftp     stream  tcp     nowait  root    /usr/sbin/tcpd  proftpd
```

This line specifies that **inetd** should accept FTP connections and pass them to **tcpd**. This may seem a bit odd, because **proftpd** normally handles FTP connections. You can also specify to use **proftpd** directly in `inetd.conf`, but Slackware normally passes the connection to **tcpd**. This program passes the connection to **proftpd** in turn, as specified. **tcpd** is used to monitor services and to provide host based access control.

Services can be disabled by adding the comment character (#) at the beginning of the line. It is a good idea to disable all services and enable services you need one at a time. After changing `/etc/inetd.conf` **inetd** needs to be restarted to activate the changes. This can be done by sending the HUP signal to the inetd process:

```
# ps ax | grep 'inetd'
   64 ?        S      0:00 /usr/sbin/inetd
# kill -HUP 64
```

TCP wrappers

As you can see in `/etc/inetd.conf` connections for most protocols are made through **tcpd**, instead of directly passing the connection to a service program. For example:

```
# File Transfer Protocol (FTP) server:
ftp     stream  tcp     nowait  root    /usr/sbin/tcpd  proftpd
```

In this example ftp connections are passed through **tcpd**. **tcpd** logs the connection through syslog and allows for additional checks. One of the most used features of **tcpd** is host-based access control. Hosts that should be denied are controlled via `/etc/hosts.deny`, hosts that should be allowed via `/etc/hosts.allow`. Both files have one rule on each line of the following form:

```
service: hosts
```

Hosts can be specified by hostname or IP address. The ALL keyword specifies all hosts or all services.

Suppose we want to block access to all services managed through **tcpd**, except for host “trusted.example.org”. To do this the following `hosts.deny` and `hosts.allow` files should be created.

`/etc/hosts.deny`:

```
ALL: ALL
```

`/etc/hosts.allow`:

```
ALL: trusted.example.org
```

In the `hosts.deny` access is blocked to all (ALL) services for all (ALL) hosts. But `hosts.allow` specifies that all (ALL) services should be available to “trusted.example.org”.

Chapter 5. Apache

Introduction

Apache is the most popular web server since April 1996. It was originally based on NCSA httpd, and has grown into a full-featured HTTP server. Slackware Linux currently uses the 1.3.x branch of Apache. This chapter is based on Apache 1.3.x.

Installation

Apache can be installed by adding the `apache` pacakge from the "n" diskset. If you also want to use PHP, the `php` ("n" diskset) and `mysql` ("ap" diskset) are also required. MySQL is required, because the precompiled PHP depends on MySQL shared libraries. You do not have to run MySQL itself. After installing Apache it can be started automatically while booting the system by making the `/etc/rc.d/rc.httpd` file executable. You can do this by executing:

```
# chmod a+x /etc/rc.d/rc.httpd
```

The Apache configuration can be altered in the `/etc/apache/httpd.conf` file. Apache can be stopped/started/restarted every moment with the **apachectl** command, and the "stop", "start" and "restart" parameters. For example, execute the following command to restart Apache:

```
# apachectl restart
/usr/sbin/apachectl restart: httpd restarted
```

User directories

Apache provides support for so-call user directories. This means every user gets webspace in the form of *http://host/~user/*. The contents of “~user/” is stored in a subdirectory in the home directory of the user. This directory can be specified using the “UserDir” option in `httpd.conf`, for example:

```
UserDir public_html
```

This specifies that the `public_html` directory should be used for storing the webpages. For example, the webpages at URL *http://host/~snail/* are stored in `/home/snail/public_html`.

Virtual hosts

The default documentroot for apache under Slackware is `/var/www/htdocs`. Without using virtual hosts every client connecting to the Apache server will receive the website in this directory. So, if we have two hostnames pointing to the server, “www.example.org” and “forum.example.org”, both will display the same website. You can make seperate sites for different hostnames by using virtual hosts.

In this example we are going to look how you can make two virtual hosts, one for "www.example.org", with the documentroot `/var/www/htdocs-www`, and "forum.example.org", with the documentroot `/var/www/htdocs-forum`. First of all we have to specify which IP addresses Apache should listen to. Somewhere in the `/etc/apache/httpd.conf` configuration file you will find the following line:

```
#NameVirtualHost *:80
```

This line has to be commented out to use name-based virtual hosts. Remove the comment character (#) and change the parameter to "BindAddress IP:port", or "BindAddress *:port" if you want Apache to bind to all IP addresses the host has. Suppose we want to provide virtual hosts for IP address 192.168.1.201 port 80 (the default Apache port), we would change the line to:

```
NameVirtualHost 192.168.1.201:80
```

Somewhere below the NameVirtualHost line you can find a commented example of a virtualhost:

```
#<VirtualHost *:80>#    ServerAdmin webmaster@dummy-host.example.com
#    DocumentRoot /www/docs/dummy-host.example.com
#    ServerName dummy-host.example.com
#    ErrorLog logs/dummy-host.example.com-error_log
#    CustomLog logs/dummy-host.example.com-access_log common
#</VirtualHost>
```

You can use this example as a guideline. For example, if we want to use all the default values, and we want to write the logs for both virtual hosts to the default Apache logs, we would add these lines:

```
<VirtualHost 192.168.1.201:80>
 DocumentRoot /var/www/htdocs-www
 ServerName www.example.org
</VirtualHost>

<VirtualHost 192.168.1.201:80>
 DocumentRoot /var/www/htdocs-forum
 ServerName forum.example.org
</VirtualHost>
```

Chapter 6. BIND

Introduction

The domain name system (DNS) is used to convert human-friendly host names (for example www.slackware.com) to IP addresses. BIND is the most widely used DNS daemon, and will be covered in this chapter.

Delegation

One of the main features is that DNS requests can be delegated. For example, suppose that you own the “linuxcorp.com” domain. You can provide the authorized nameservers for this domain, you nameservers are authorative for the “linuxcorp.com”. Suppose that there are different branches within your company, and you want to give each branch authority over their own zone, that is no problem with DNS. You can delegate DNS for e.g. “sales.linuxcorp.com” to another nameserver within the DNS configuration for the “linuxcorp.com” zone.

The DNS system has so-called root servers, which delegate the DNS for millions of domain names and extensions (for example, country specific extensions, like “.nl” or “.uk”) to authorized DNS servers. This system allows a branched tree of delegation, which is very flexible, and distributes DNS traffic.

DNS records

The following types are common DNS records:

Table 6-1. DNS records

Prefix	Description
A	An A records points to an IP address.
CNAME	A CNAME record points to another DNS entry.
MX	A MX record specifies which should handle mail for the domain.

Masters and slaves

Two kinds of nameservers can be provided for a domain name: a master and slaves. The master server DNS records are authorative. Slave servers download their DNS record from the master servers. Using slave servers besides a master server is recommended for high availability and can be used for load-balancing.

Making a caching nameserver

A caching nameserver provides DNS services for a machine or a network, but does not provide DNS for a domain. That means it can only be used to convert hostnames to IP addresses. Setting up a nameserver with Slackware Linux is fairly easy, because BIND is configured as a caching nameserver by default. Enabling the caching nameserver takes just two steps: you have to install BIND and alter the initialization scripts. BIND can be installed by adding the bind package from the "n" disk set. After that bind can be started by executing the **named(8)** command. If want to start BIND by default, make the `/etc/rc.d/rc.bind` file executable. This can be done by executing the following command as root:

```
# chmod a+x /etc/rc.d/rc.bind
```

If you want to use the nameserver on the machine that runs BIND, you also have to alter `/etc/resolv.conf`.

Appendix A. GNU Free Documentation License

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. Secondarily, 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.

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.

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.

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.

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.

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".

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.

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.

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.

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.

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.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

> Copyright (c) YEAR YOUR NAME. 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 "GNU Free Documentation License".

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

> with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.