

RUN YOUR OWN **WEB SERVER** USING LINUX & APACHE

BY STUART LANGRIDGE
& TONY STEIDLER-DENNISON



GET STARTED WITH LINUX AND APACHE -- THE EASY WAY!

Summary of Contents

[Introduction](#)

[3. The Command Line](#)

[9. Server Security](#)

[A. Command Line Reference](#)

[B. Troubleshooting](#)

[Index](#)

RUN YOUR OWN WEB SERVER USING LINUX & APACHE

BY STUART LANGRIDGE
& TONY STEIDLER-DENNISON

Run Your Own Web Server Using Linux & Apache

by Stuart Langridge and Tony Steidler-Dennison

Copyright © 2005 SitePoint Pty. Ltd.

Expert Reviewer: Stephen Pierzchala

Managing Editor: Simon Mackie

Technical Editor: Craig Anderson

Technical Director: Kevin Yank

Editor: Georgina Laidlaw

Index Editor: Bill Johncocks

Cover Design: Jess Mason

Cover Layout: Alex Walker

Notice of Rights

All rights reserved. No part of this book may be reproduced, stored in a retrieval system or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical articles or reviews.

Notice of Liability

The author(s) and publisher have made every effort to ensure the accuracy of the information herein. However, the information contained in this book is sold without warranty, either express or implied. Neither the author(s) and SitePoint Pty. Ltd., nor its dealers or distributors will be held liable for any damages to be caused either directly or indirectly by the instructions contained in this book, or by the software or hardware products described herein.

Trademark Notice

Rather than indicating every occurrence of a trademarked name as such, this book uses the names only in an editorial fashion and to the benefit of the trademark owner with no intention of infringement of the trademark.



Published by SitePoint Pty. Ltd.

424 Smith Street Collingwood
VIC Australia 3066.

Web: www.sitepoint.com
Email: business@sitepoint.com

About the Authors

Stuart Langridge has been a Linux user since 1997, and quite possibly is the only person in the world to have a BSc in Computer Science and Philosophy. He's also one-quarter of the team at LugRadio, the world's premiere Free and Open Source Software radio show. He's a keen advocate of Free Software across the board, for its ethics as much as its functionality, and thinks that you should be, too. When he's not fiddling about with computers, he's an information architect, the author of SitePoint's *DHTML Utopia: Modern Web Design Using JavaScript & DOM*, and drinker of decent beers.

Tony Steidler-Dennison is a Systems Engineer with Rockwell Collins, Inc., designing avionics and cabin data servers for commercial airliners. He's also the host of The Roadhouse Podcast, "the finest blues you've never heard."

About the Expert Reviewer

Stephen Pierzchala is currently the Senior Performance Analyst at Gomez, Inc., as well as the Chief

Performance Evangelist for [WebPerformance](#), and the primary developer for the [GrabPERF Performance Monitoring System](#). He has actively worked with, supported and analyzed data from Internet technologies since 1994. A Canadian by birth (and inclination), he has been living in the United States since 1999. Stephen lives in Marlborough, Massachusetts with his wife, Samantha, and two sons, Cameron and Kinnear.

About The Technical Director

As Technical Director for SitePoint, Kevin Yank oversees all of its technical publications—books, articles, newsletters and blogs. He has written over 50 articles for SitePoint, but is best known for his book, *Build Your Own Database Driven Website Using PHP & MySQL*. Kevin lives in Melbourne, Australia, and enjoys performing improvised comedy theatre and flying light aircraft.

About SitePoint

SitePoint specializes in publishing fun, practical, and easy-to-understand content for Web professionals. Visit <http://www.sitepoint.com/> to access our books, newsletters, articles and community forums.

To Niamh, who is going to know what this one is
all about one day

—Stuart

For my girls, whom I seldom see when I'm
cloistered away writing, and who have the utmost
patience when I'm surly and behind my deadlines

—Tony

Introduction

More organizations install Linux into their server rooms every day. The reasons for this vary, but those who make the switch to Linux often claim that its reliability, cost, choice, scalability, and the freedom it offers from vendor lock-in, are some of the reasons why they decided to switch. But, whatever your reasons for choosing Linux, as system administrator, you need to know what to do with these new servers.

This book gives you the knowledge you need to build, configure, and maintain servers running the LAMP (Linux, Apache, MySQL, and PHP) open source Web application platform.

In these pages, we'll show you how to build a Linux server, and help you decide which flavour of Linux best suits your situation. You'll learn how to set up Apache to serve Websites, use MySQL to store data, and employ PHP to build Web applications. You'll also discover how to secure your new LAMP server, and how best to access and control it both on site, and remotely.

Everything you'll need to build and maintain your Linux servers, and to deploy Web applications to them, is contained in these chapters. Enjoy!

Who Should Read This Book?

If you know what it's like to be a systems administrator, but don't know about Linux, this book is for you.

If you're currently thinking about introducing Linux to your firm on a trial basis—perhaps to run some Websites, or because your development team keep banging on at you to let them use Apache—this book will give you the grounding you need to successfully build Linux servers and keep them running.

You and your organization can enjoy the stability and ease-of-use of a free operating system and tools that are compliant with open standards. This book will show you how.

What's In This Book?

[*Chapter 1: Building The Linux Environment*](#)

We kick off by discussing what Linux is, and seeing how easy it is to set up a Linux server. I'll walk you through a review of the hardware you'll need before we choose which flavour of Linux we'll install. In making this decision, we'll explore the alternatives, and I'll provide a few tips on how to ensure you make the right decision. By the chapter's close, you'll have installed Linux successfully on your server.

[*Chapter 2: Day-to-day Usage*](#)

This chapter explains how to run and manage your Linux server on a daily basis. In particular,

we'll discuss the key differences between Linux and Windows systems. You'll finish up with a solid grounding in the essentials, including filesystems and layout.

[Chapter 3: The Command Line](#)

The command line is one of the most powerful aspects of Linux. While it's easy to use Linux's graphical tools—tools that allow us to achieve most of our goals—the command line gives us an extra level of control over our systems. In this chapter, we'll identify those extra capabilities, and discuss the command line's advantages over the GUI.

[Chapter 4: System Administration](#)

Linux system administrators must be comfortable with creating new users, and scheduling tasks to run unattended, as well as concepts like services and runlevels. In this chapter, we discuss the lot as we take a tour of the Linux system administrator's toolkit.

[Chapter 5: Building The Server](#)

It's time to turn our Linux server into a LAMP server as we install Apache, MySQL, and PHP. We'll explore some of the basics of Apache itself, including how it works, and how it fits into the Linux environment. Then, you'll see how to configure Apache on setup, how to set up secure (SSL/https) access for your Websites, and how to add MySQL and PHP to the mix.

[Chapter 6: Server Administration](#)

This chapter focuses on a selection of handy tools that will help you to configure your LAMP server and add new packages to it. In particular, we discuss Webmin, which facilitates the Web-based configuration of services, and yum, which helps with package installation.

[Chapter 7: Remote Administration](#)

Remote administration makes the job of the Linux system administrator much easier. In this chapter, we'll get a feel for working with SSH—the secure shell—which allows command line access to a server across the network, and VNC, which enables you to access your LAMP server GUI remotely. We'll discuss which tools are best used in particular situations, and look at some of the extra functionality that SSH offers above and beyond its primary job as a command-line shell.

[Chapter 8: Occasional Administration](#)

“Occasional Administration” encompasses those system elements that you'll likely need to set up once, then tweak only occasionally. After an introduction to backup tools, we set up Web traffic reporting, which will help us understand the nature of the visits the server receives. We also take a close look at the log files that the system creates, and discuss how these can be used to track and manage server usage, identify errors, and more.

[Chapter 9: Server Security](#)

Security is a critical aspect of running any server, but it's particularly important for those that

offer services over the Internet. In this chapter, we set up a firewall on our LAMP server, and ~~install intrusion detection services as a means to identify remote cracking attempts.~~ We'll also meet Tripwire, a security system that protects against malicious users compromising the server they somehow manage to gain access.

[Appendix A: Command Line Reference](#)

As we step through the process of setting up your server, you'll be introduced to a number of powerful command line tools. This appendix lists the more useful tools, and some of the options that can be used to customize their behavior.

[Appendix B: Troubleshooting](#)

Our tour concludes with some troubleshooting and an FAQ section that provides answers to common questions about Linux, Apache, and the other tools you may have installed.

Linux and Distributions

If you're being very technical, Linux is just the operating system kernel: the bit at the very lowest level of your software that talks directly to the hardware. All the other programs—the graphical interface, the Apache Web server, MySQL, the menus at the top of the screen—they're all separate, open-source programs, coded by different teams of developers, and released at different times. It's possible to build your own Linux system out of these disparate parts, but it's a long and complicated job. Instead, various groups and companies have taken on the role of providing a Linux **distribution** (sometimes shortened to **distro**): they collect all the bits of software you need, make sure they all fit together correctly, and give them all to you in one go. There are many, many Linux distributions. Some have specialized purposes: the distributor has made sure that the distribution contains software suited to musicians, for example, or medical personnel, or security analysts, or that the distribution is designed to run directly from CD, or from a USB pen drive, or without a graphical interface. Most, though, are general: they're designed to cover all bases.

General-purpose Distributions

Some of the most popular general distributions are Debian, Canonical's Ubuntu Linux, Novell's SuSE Linux Desktop, Knoppix, Linspire, and Red Hat's Enterprise Linux and Fedora Core distributions. Some distributions contain proprietary software, and require a licence fee; others do not. Each has its merits, and each its proponents.

Debian has a very strong free-software ethos, and an excellent packaging system (apt) which has been emulated by most other distributions. Ubuntu Linux is derived from Debian, but places a much stronger focus on being a good desktop distribution. Novell's SuSE Linux has a commercial edge to its mixing open-source and proprietary tools; Novell is a relatively new player in the Linux sphere, but SuSE Linux has been around for some time. Red Hat Enterprise Linux also has a commercial edge, and comes in flavors tailored to desktops and servers. Development of Red Hat Enterprise Linux is in part driven by the more community spirited Fedora Core. Linspire is heavily focused on home users; like SuSE, it mixes open-source tools with proprietary software, and is oriented towards being a desktop operating system. Knoppix is slightly unusual in that it is not designed to be installed and run

instead, it comes on, and runs entirely from, a so-called “Live CD.” You can simply put the CD in an
boot up to obtain all the benefits of a working Linux computer without losing or overwriting your
existing system or files. It’s perfect for testing out hardware, or getting familiar with the Linux
environment without taking the ultimate plunge and installing the system from scratch.

Fedora Core

In this book, we’ll be focusing on Fedora Core 4, from Red Hat. The Fedora distribution is very
current, so you’ll have all the latest tools at your disposal, and boasts a very wide portfolio of
compatible software. Red Hat employs many notable open-source developers to work on distribution
and Fedora receives the benefits of this work, while still remaining open-source and community-
maintained. Using Fedora, you can enjoy those benefits: you’ll have the most robust, modern tools at
your fingertips, while using the most popular Linux distribution available.

The Book’s Website

Located at <http://www.sitepoint.com/books/linux1/>, the Website that supports this book will give you
access to the following facilities.

The Code Archive

One of the more powerful aspects of Linux is the scriptable command line. This book includes some
scripts to help you get started with shell scripting, which can be downloaded from the book’s web site

Updates and Errata

The Errata page on the book’s Website has the latest information about known typographical and code
errors, and updates necessitated by changes to technologies.

The SitePoint Forums

While I’ve made every attempt to anticipate any questions you may have, and answer them in this
book, there is no way that *any* book could cover everything there is to know about establishing,
running, and maintaining a Linux server. If you have a question about anything in this book, the best
place to go for a quick answer is <http://www.sitepoint.com/forums/>—SitePoint’s vibrant and
knowledgeable community.

The SitePoint Newsletters

In addition to books like this one, SitePoint offers free email newsletters.

The SitePoint Tech Times covers the latest news, product releases, trends, tips, and techniques for all
technical aspects of Web development. The long-running *SitePoint Tribune* is a biweekly digest of the
business and money making aspects of the Web. Whether you’re a freelance developer looking for tips
to score that dream contract, or a marketing major striving to keep abreast of changes to the major

search engines, this is the newsletter for you. *The SitePoint Design View* is a monthly compilation of the best in Web design. From new CSS layout methods to subtle PhotoShop techniques, SitePoint's chief designer shares his years of experience in its pages.

Browse the archives or sign up to any of SitePoint's free newsletters at <http://www.sitepoint.com/newsletter/>.

Your Feedback

If you can't find your answer through the forums, or you wish to contact me for any other reason, the best place to write is books@sitepoint.com. We have a well-manned email support system set up to track your inquiries, and if our support staff is unable to answer your question, they send it straight to me. Suggestions for improvement as well as notices of any mistakes you may find are especially welcome.

Acknowledgements

This book would not have been what it is without the SitePoint team, particularly Stephen, Craig, and Simon. A big round of applause also goes out to Ade, for coping in his typically composed and bald style with the bombardment of questions. Much thanks, bald man.

Inspiration is theirs; mistakes are mine alone.

Building The Linux Environment

Installing a Linux distribution can be both exhilarating and frustrating. My first two attempts at Linux installs—the first in 1996, the second in 1997—were unsuccessful. Installation routines and hardware support in Linux at the time were much less advanced than they are today; Red Hat was still at a relatively early stage in its evolution, Mandriva had yet to be created, and SuSE was just coming out from under the shadow of Slackware. After two failures, I simply decided that I wasn't going to be beaten by a Linux distribution. I set my machine up in a dual-boot configuration (including both Linux and Windows partitions) with the commitment to use Windows as little as possible. Within a year, the only reason Windows remained on the machine was my wife's lack of familiarity with Linux. Given that her computing needs were to surf the Web and read email, she, too, eventually made a smooth transition to Linux as the full-time computing platform.

We'll talk about the dual-boot option at length in this chapter. But first, it's important to undertake some preliminary research that will help you solve the issues you might experience during installation, whether you're using a pure Linux system, or a dual-boot configuration.

The Necessary Research

Few things are more frustrating than a lack of hardware support, especially when you've become used to the quick driver installs offered by Windows. In fact, Windows comes complete with a basic set of drivers that are intended to anticipate the hardware attached to your machine. Hardware manufacturers also release driver discs for devices such as video cards, network cards and scanners for Windows machines. Developing these drivers costs the hardware manufacturers a great deal of money, so for a long time it didn't make economic sense for hardware developers to supply drivers for Linux.

As Linux has gained market share within the server market, Linux driver development has improved markedly. Storage devices, RAID arrays, Ethernet cards—all have enjoyed increasing Linux driver development in the past few years.

In order to avoid the headache of missing drivers, it's important to do a little research before installing your Linux distribution. While it's unlikely that you'll have a problem with modern distributions, you'll still want to do the research just to avoid any hardware issues.

Most of the major distributions release **hardware compatibility lists**. These lists itemize the hardware that's known to work with the drivers included in the distributions. Red Hat/Fedora, Mandriva, and SuSE also provide hardware mailing lists for distributions from their Websites. These lists, though, tend to rely on users to help solve hardware compatibility issues after the fact, rather than providing information for users before an installation.

Additionally, there's an excellent compatibility list for Linux in general. It doesn't provide quite the

degree of granularity you'll find in the manufacturer-specific lists, so it should be used as a fallback, rather than your primary source of information.

Hardware Compatibility Lists

Red Hat/Fedora

Red Hat's major product line is Red Hat Enterprise Linux (RHEL), which is mostly based on Red Hat's free software distribution, Fedora. Fedora is not actually maintained by Red Hat; it's maintained by the community of Fedora developers. However, Red Hat does a lot of work on Fedora, because that work flows into RHEL.

[Red Hat's Hardware Catalog](#) doesn't extend beyond RHEL to the Fedora releases, which is something that you'll need to remember when looking to the Red Hat site for Fedora support. The list provides information on CPUs, video cards, SCSI controllers, IDE controllers, network cards, modems, and sound cards.

SuSE

SuSE offers two lists: the [Express Search](#) and [Extended Search](#). The difference between the two is that the Extended Search offers fields beyond Vendor, Device, and Category. In practice, you're likely only to need the Express Search.

Mandriva Linux

The [Mandriva Linux Hardware Compatibility Database](#) is a very comprehensive list of hardware that has been tested by the Mandriva Linux community.

General Linux

The [Linux Hardware Compatibility HOWTO](#) is perhaps the most comprehensive of the high-level Linux links. It was begun in 1997 and is updated as often as twice annually. It provides information on all device types and all major manufacturers.

Aside from providing interesting and useful user forums, LinuxQuestions.org also provides [an outstanding list of Linux-compatible hardware](#). This is the most up-to-date of the high-level Linux lists, with updates appearing daily where applicable. While it's not as comprehensive as the HOWTO, the LinuxQuestions list is easily as important because of this timeliness.

[Linux Compatible](#) provides both updated lists, and forums in which users can help other users resolve existing hardware issues.

Installing the Distribution

Once you have completed your preliminary hardware research, it's time to walk through the

installation process. We'll take a look at both the graphical and text-based installers, the second of which can be useful when you're installing Linux on a machine with limited resources. Don't forget that, if your situation demands it, you can install Fedora Core on your server without a desktop. In any event, it's a good idea to read through the following sections before putting the installation CD into your computer.

The Dual-Boot Option

We've already mentioned the dual-boot option for your server: running both Windows and Linux on the system. As I've mentioned, this provides a great set of "technical training wheels" as you adjust to the new capabilities and options in your Linux server. The following installation instructions will work equally well with a dual-boot configuration. However, there are a few important points to keep in mind when choosing this option.

- If you're building your dual-boot server on a fresh box, be sure to install and configure Windows first. By default, Windows doesn't recognize any of the native Linux filesystems.^[1] If Linux is installed first, the Windows boot loader will take over and load Windows; Linux will be there, but you won't be able to boot into it. A Linux installation will cooperate with Windows and allow you to boot into both.
- Linux provides a means to read the FAT32 (typically used by Windows 98 and ME) or NTFS (usually used by Windows NT, 2000, and XP) filesystems. In the case of FAT32, you'll also be able to write to the Windows partitions. If you're using an NTFS-based Windows installation, the files on the Windows partition will be read-only.
- If you're installing Linux on a system that already contains a Windows operating system, it may be useful to purchase a nondestructive partition management tool, such as [Partition Magic](#). This will allow you to move the partitions on your Windows system, creating room on the drive for the Linux installation, and preserving the data that already exists on the drive.

With the exception of these important points, the process of installing a dual-boot system is the same as a single OS installation.

Graphical Installation

Some would argue that the real rise of Linux began with the advent of graphical installers. Prior to that time, installation was a "mouseless" affair, using the keyboard arrow keys and space bar. Red Hat—the distribution upon which Fedora is based—was a pioneer in graphical Linux installation routines. Since that time, the creators have continued to refine and improve upon the process, the result being a very clean and easy-to-follow installation procedure. As you'll see in the screen shots I'll present throughout the rest of this chapter, installing Fedora on your new server is nearly painless!

I've provided screen shots for nearly every step of the process. While the procedure is easy, there are a few steps that are particularly important to a successful installation. Hopefully, the abundance of screen shots in the following discussion will help you to more easily understand the installation process.

Obtaining Installation CDs

There are two main ways to obtain Fedora Core installation CDs: you can download the CDs from <http://fedora.redhat.com/download/> and burn them yourself, or you can buy them.

The installation CDs are downloaded as a series of ISO images, named something like FC4-i386-disc1.iso (FC4 means Fedora Core 4, i386 means it's for Intel x86 processors, and disc1 means that it's the first CD). ISO images are direct copies of an entire CD, stored in a single file. Once you've downloaded the images, you'll need to burn each of them to a CD.^[2] Most CD burning programs offer a menu option to burn an ISO image; a list of [instructions for the use of various popular Windows CD burning tools](#) is also available online. If in doubt, the help files, or Websites, associated with your CD burning tool are likely to explain how to burn an ISO image onto a CD.^[3]

Buying Fedora on CD will cost you a little, but it's quicker and easier than downloading the images if you don't have a fast broadband connection (the four CD images total almost 2.5GB). You can buy Fedora Installation CDs from any number of vendors, most of whom will charge you little more than the cost of the blank CDs, plus postage and packing; the easiest way to find these vendors is to search the Web for "cheap Linux CDs" in your country, or ask a local Linux User Group. This may well be the best way to get hold of the CDs if this is your first time running Linux.

The Installation

To begin the installation, put the first installation CD in the CD-ROM drive and reboot the machine. If your machine is configured to boot from the CD-ROM, you'll see the screen shown in [Figure 1.1](#) when the machine starts.

The initial installation offers several options. You can choose to install in graphical mode by hitting **Enter**, or in text mode by typing **linux text** at the boot : prompt. Either way, the first thing the installer will do is offer to check the installation media for you. This is a good way to determine if your installation CDs have been tampered with, or have become corrupted. The process will take a little while, but I'd recommend that you do run this test.

Like any operating system, Linux requires a minimal set of hardware drivers during the installation. After testing the installation media, you'll see lots of text scrolling down the screen—this is the initial hardware probing process in action. Red Hat helped pioneer the development of graphical Linux installers with **Anaconda**, Red Hat's installation program. It includes a highly accurate probing and testing mechanism that makes the rest of the installation routine quite painless.

Once all this media testing and hardware probing is done, you'll finally see the Welcome to Fedora Core screen. Click the Next button to get started.




Figure 1.1. The initial Fedora installation screen.

Selecting your Language




Figure 1.2. Choosing an installation language.

Fedora is truly an international operating system: the installation screens are available in more than 3

- [read Edible Forest Gardens - Volume 2: Ecological Design and Practice For Temperate Climate Permaculture](#)
- [Molly Fyde and the Land of Light \(The Bern Saga, Book 2\) pdf, azw \(kindle\), epub](#)
- [download online The Infinite Sea \(The 5th Wave, Book 2\)](#)
- **[download Philosophy of Language: The Key Thinkers online](#)**
- [Cake Balls: Amazingly Delicious Bite-Size Treats pdf, azw \(kindle\), epub, doc, mobi](#)

- <http://anvilpr.com/library/Schlacht-um-Ferrol--Perry-Rhodan-Neo--Band-11--Expedition-Wega--Band-3-.pdf>
- <http://fortune-touko.com/library/Molly-Fyde-and-the-Land-of-Light--The-Bern-Saga--Book-2-.pdf>
- <http://www.1973vision.com/?library/The-Infinite-Sea--The-5th-Wave--Book-2-.pdf>
- <http://weddingcellist.com/lib/Philosophy-of-Language--The-Key-Thinkers.pdf>
- <http://test.markblaustein.com/library/Ballerina--Sex--Scandal--and-Suffering-Behind-the-Symbol-of-Perfection.pdf>