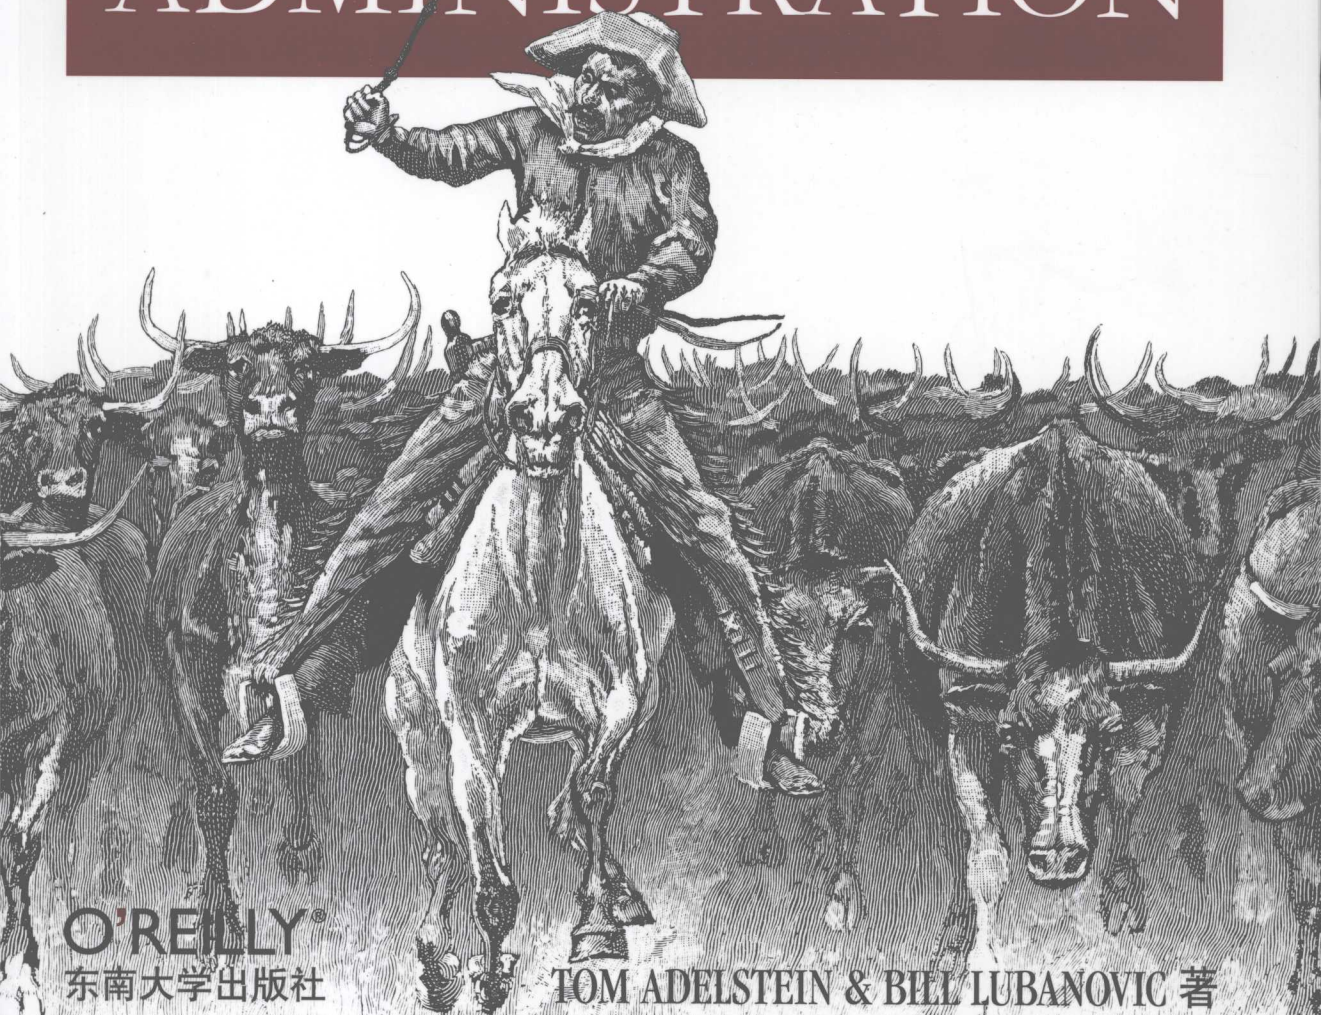


LINUX系统管理 (影印版)

LINUX

SYSTEM ADMINISTRATION



O'REILLY®
东南大学出版社

TOM ADELSTEIN & BILL LUBANOVIC 著

LINUX

SYSTEM ADMINISTRATION

LINUX系统管理 (影印版)

江苏工业学院图书馆
藏书章

Tom Adelstein and Bill Lubanovic

O'REILLY®

Beijing · Cambridge · Farnham · Köln · Paris · Sebastopol · Taipei · Tokyo

O'Reilly Media, Inc. 授权东南大学出版社出版

南京 · 东南大学出版社

图书在版编目 (CIP) 数据

Linux 系统管理: 英文 / (美) 阿德斯特 (Adelstein, T.),
(美) 卢布诺维克 (Lubanovic, B.) 著. — 影印本. — 南
京: 东南大学出版社, 2008.2

书名原文: Linux System Administration

ISBN 978-7-5641-1038-3

I . L … II . ①阿… ②卢… III .Linux 操作系统—英
文 IV .TP316.89

中国版本图书馆 CIP 数据核字 (2007) 第 193116 号

江苏省版权局著作权合同登记

图字: 10-2007-218 号

©2007 by O'Reilly Media, Inc.

Reprint of the English Edition, jointly published by O'Reilly Media, Inc. and Southeast University Press, 2008. Authorized reprint of the original English edition, 2007 O'Reilly Media, Inc., the owner of all rights to publish and sell the same.

All rights reserved including the rights of reproduction in whole or in part in any form.

英文原版由 O'Reilly Media, Inc. 出版 2007。

英文影印版由东南大学出版社出版 2008。此影印版的出版和销售得到出版权和销售权的所有者——O'Reilly Media, Inc. 的许可。

版权所有, 未得书面许可, 本书的任何部分和全部不得以任何形式复制。

Linux 系统管理 (影印版)

出版发行: 东南大学出版社

地 址: 南京四牌楼 2 号 邮编: 210096

出版人: 江 汉

网 址: <http://press.seu.edu.cn>

电子邮件: press@seu.edu.cn

经 销: 全国各地新华书店

印 刷: 扬中市印刷有限公司

开 本: 787 毫米 × 980 毫米 16 开本

印 张: 18.75 印张

字 数: 315 千字

版 次: 2008 年 2 月第 1 版

印 次: 2008 年 2 月第 1 次印刷

书 号: ISBN 978-7-5641-1038-3/TP · 174

印 数: 1 ~ 3500 册

定 价: 42.00 元 (册)

本社图书若有印装质量问题, 请直接与读者服务部联系。电话 (传真): 025-83792328

O'Reilly Media, Inc. 介绍

O'Reilly Media, Inc. 是世界上在 UNIX、X、Internet 和其他开放系统图书领域具有领导地位的出版公司，同时是联机出版的先锋。

从最畅销的《The Whole Internet User's Guide & Catalog》（被纽约公共图书馆评为二十世纪最重要的 50 本书之一）到 GNN（最早的 Internet 门户和商业网站），再到 WebSite（第一个桌面 PC 的 Web 服务器软件），O'Reilly Media, Inc. 一直处于 Internet 发展的最前沿。

许多书店的反馈表明，O'Reilly Media, Inc. 是最稳定的计算机图书出版商——每一本书都一版再版。与大多数计算机图书出版商相比，O'Reilly Media, Inc. 具有深厚的计算机专业背景，这使得 O'Reilly Media, Inc. 形成了一个非常不同于其他出版商的出版方针。O'Reilly Media, Inc. 所有的编辑人员以前都是程序员，或者是顶尖级的技术专家。O'Reilly Media, Inc. 还有许多固定的作者群体——他们本身是相关领域的技术专家、咨询专家，而现在编写著作，O'Reilly Media, Inc. 依靠他们及时地推出图书。因为 O'Reilly Media, Inc. 紧密地与计算机业界联系着，所以 O'Reilly Media, Inc. 知道市场上真正需要什么图书。

出版说明

随着计算机技术的成熟和广泛应用,人类正在步入一个技术迅猛发展的新时期。计算机技术的发展给人们的工业生产、商业活动和日常生活都带来了巨大的影响。然而,计算机领域的技术更新速度之快也是众所周知的,为了帮助国内技术人员在第一时间了解国外最新的技术,东南大学出版社和美国 O'Reilly Meida, Inc.达成协议,将陆续引进该公司的代表前沿技术或者在某专项领域享有盛名的著作,以影印版或者简体中文版的形式呈献给读者。其中,影印版书籍力求与国外图书“同步”出版,并且“原汁原味”展现给读者。

我们真诚地希望,所引进的书籍能对国内相关行业的技术人员、科研机构的研究人员和高校师生的学习和工作有所帮助,对国内计算机技术的发展有所促进。也衷心期望读者提出宝贵的意见和建议。

最新出版的影印版图书,包括:

- 《深入浅出 PMP》(影印版)
- 《RESTful Web Services》(影印版)
- 《学习 WCF》(影印版)
- 《学习 Ruby》(影印版)
- 《Beautiful Code》(影印版)
- 《ActionScript 3.0 精髓》(影印版)
- 《Linux 系统管理》(影印版)
- 《精通 Perl》(影印版)
- 《深入浅出面向对象分析与设计》(影印版)
- 《Linux Kernel 技术手册》(影印版)
- 《ActionScript 3.0 Cookbook》(影印版)
- 《CSS:The Missing Manual》(影印版)



Preface

As Bill Lubanovic and I were putting the final touches on this book, I overheard a conversation between two coworkers in our Cisco lab discussing Linux. The senior networking guru of the two made an interesting remark. He said that despite all his knowledge, he felt incomplete as a professional because he had never learned Linux. A moment later he and the other gentleman turned to me and looked me square in the eyes. I smiled and went on working.

That evening, our director of Information Technology made an offhand remark to me during a conference that struck me as unusual. He said that he wanted to learn Apache, and when I asked him why he replied, "I just want to learn it," and left it at that.

Later in the conference, our director requested feedback from the group on a solution for patch management, explaining and using the example of *rsync*. He said he wanted something similar, while launching into a detailed technical discussion of incremental and cumulative patch management. I have a good working knowledge of *rsync*, but hadn't heard such a detailed academic explanation of any open source tool in any forum.

In both of those cases and many others, I wished I had this book ready to hand over to highly trained and skilled people who wanted to learn Linux administration. Perhaps you have had similar experiences and wished you had a book like this one at hand. I venture to guess that conversations like the ones I've just described occur many times in many places daily.

When Andy Oram and I began discussing a Linux system administration book, we had a slightly different idea of what we wanted to accomplish. Andy talked about a book in which each chapter took users through the steps of building and deploying application servers without co-mingling detailed discussions. He suggested that the discussion reside in one place in each chapter and the technical steps in another.

Later, I proposed that we make each chapter a module unto itself and let the reader complete the modules he wanted and/or needed. As this book evolved, we felt that we'd accomplished that objective. You do not have to read this book cover to cover to become a Linux system administrator. Simply start where you have the most interest.

When I first started using Linux, the community consisted mostly of programmers and hobbyists. I don't recall any discussion lists that focused on desktops or commercial applications. We logged onto the Internet by starting a daemon. We didn't have dialers or web browsers like the ones available today. The vast majority of people I knew did their own system administration or were in some stage of learning.

Reflecting on the time when we estimated that 30,000 Linux users existed on the planet, I'm amazed at how many people use Linux today and haven't the slightest idea how to write a configuration file. Linux forums seem to be filled with people asking how to get CUPS or Samba to work. On mailing lists, people hold detailed discussions on the technical details of projects like Postfix, JBoss, and Monit.

Many people still itch to learn the extensive capabilities of Linux as an application platform. If you use Linux and want to take the next step from a power user to an administrator, this book will help you make the transition. We wrote this book with you in mind.

How This Book Is Organized

Chapter 1, *Requirements for a Linux System Administrator*

Lays out the goals of the book and what you'll gain by reading it.

Chapter 2, *Setting Up a Linux Multifunction Server*

Gets you started with a nearly Internet-ready server.

Chapter 3, *The Domain Name System*

Shows you the basics of setting up primary and secondary DNS servers.

Chapter 4, *An Initial Internet-Ready Environment*

Uses the ISPConfig free software configuration system to get you started with a rich set of services that you can practice while reading the rest of the book.

Chapter 5, *Mail*

Sets up a Postfix mail server with SASL authentication, a POP server, and an IMAP server.

Chapter 6, *Administering Apache*

Gives a quick run-through of the popular Apache, MySQL, and PHP combination (together with Linux, known as a LAMP server), including SSL authentication.

Chapter 7, *Load-Balanced Clusters*

Extends the previous chapter's Apache configuration with IP Virtual Server and *ldirectord* to provide high availability.

Chapter 8, *Local Network Services*

Shows you how to manage users and configure common networking elements such as DHCP and gateway software on local area networks (LANs).

Chapter 9, *Virtualization in the Modern Enterprise*

Shows how to set up Xen, VMware on a Linux host and then add guest operating systems.

Chapter 10, *Scripting*

Shows you some basic techniques for writing robust and powerful *bash* shell scripts that can save you a lot of administration time.

Chapter 11, *Backing Up Data*

Presents a range of techniques for carrying out this crucial function, from basic *rsync* and *tar* to the powerful Amanda system.

Appendix, *bash Script Samples*

Contains a few shell scripts that we've found useful when doing system administration and that might give you tips for how to write your own scripts.

Conventions Used in This Book

The following typographical conventions are used in this book:

Italic

Indicates new terms, URLs, commands and command-line options, email addresses, filenames, file extensions, and directories.

Constant width

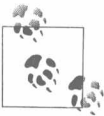
Indicates the contents of files or the output from commands.

Constant width bold

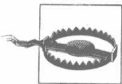
Shows commands or other text that should be typed literally by the user. Also used to highlight key portions of code or files.

Constant width italic

Shows text that should be replaced with user-supplied values.



This icon signifies a tip, suggestion, or general note.



This icon indicates a warning or caution.

Using Code Examples

This book is here to help you get your job done. In general, you may use the code in this book in your programs and documentation. You do not need to contact us for permission unless you're reproducing a significant portion of the code. For example, writing a program that uses several chunks of code from this book does not require permission. Selling or distributing a CD-ROM of examples from O'Reilly books does require permission. Answering a question by citing this book and quoting example code does not require permission. Incorporating a significant amount of example code from this book into your product's documentation does require permission.

We appreciate, but do not require, attribution. An attribution usually includes the title, author, publisher, and ISBN. For example: "*Linux System Administration* by Tom Adelstein and Bill Lubanovic. Copyright 2007 O'Reilly Media, Inc., 978-0-596-00952-6."

If you feel your use of code examples falls outside fair use or the permission given above, feel free to contact us at permissions@oreilly.com.

Safari® Enabled



When you see a Safari® Enabled icon on the cover of your favorite technology book, that means the book is available online through the O'Reilly Network Safari Bookshelf.

Safari offers a solution that's better than e-books. It's a virtual library that lets you easily search thousands of top tech books, cut and paste code samples, download chapters, and find quick answers when you need the most accurate, current information. Try it for free at <http://safari.oreilly.com>.

How to Contact Us

Please address comments and questions concerning this book to the publisher:

O'Reilly Media, Inc.
1005 Gravenstein Highway North
Sebastopol, CA 95472
800-998-9938 (in the United States or Canada)
707-829-0515 (international or local)
707-829-0104 (fax)

We have a web page for this book, where we list errata and any additional information. You can access this page at:

<http://www.oreilly.com/catalog/9780596009526>

Examples, tips, and new procedures will be posted from time to time at the test site set up by the authors for the book:

<http://www.centralsoft.org>

To comment or ask technical questions about this book, send email to:

bookquestions@oreilly.com

For more information about our books, conferences, Resource Centers, and the O'Reilly Network, see our web site at:

<http://www.oreilly.com>

Acknowledgments

Books such as *Linux System Administration* come into existence only with the contribution of many people's efforts. Consider it impossible to list them all here.

First, we would like to thank Andy Oram, whose editing, writing, and management efforts to get this book into shape seem remarkable. Apart from working as the overall editor, Andy contributed materially to the content of this book. Andy functioned like a project manager and demonstrated both patience and discipline.

We could not have asked more from the contributions of Falko Timme, Phil Howard, and Herschel Cohen. Falko lent his time and expertise to Chapters 2 and 4. Phil wrote the bulk of Chapter 11 and provided the framework for Chapter 10 and the accompanying appendix of scripts. Herschel wrote sections of several chapters, including Chapters 8 and 10, and contributed his expertise to Chapter 6. All three contributors also reviewed other parts of the book.

Many thanks are also due to our technical experts, who spent countless hours reviewing, testing, and making suggestions about our work: Markus Amersdorfer, Keith Burgess, Robert Day, Ammar Ibrahim, and Yaman Saqqa.

Special thanks go to Yvonne Adelstein and Mary Lubanovic, our wives, who showed remarkable patience. We could not have done this without your total support.

Table of Contents

Preface	xi
1. Requirements for a Linux System Administrator	1
About This Book	2
How Can We Help?	2
Where Do You Start?	3
Do You Need a Book?	3
Who Needs You?	4
What System Managers Should Know About Linux	6
What's Next	7
2. Setting Up a Linux Multifunction Server	8
Server Requirements	9
Installing Debian	10
Logging in Remotely	12
Configuring the Network	13
Changing the Default Debian Packages	15
Setting Up Quotas	16
Providing Domain Name Services	18
Adding a Relational Database: MySQL	20
Configuring Mail Securely with Postfix, POP3, and IMAP	22
Putting Apache to Work	33
Adding FTP Services with ProFTPD	34
Summarizing Your Web Statistics with Webalizer	35
Synchronizing the System Clock	36
Installing Perl Modules Needed by SpamAssassin	36
What's Next	37

3. The Domain Name System	38
DNS Basics	38
Getting into the BIND	40
Setting Up a DNS Server	41
Configuring an Authoritative DNS Server	44
Editing the Configuration Files	50
BIND Tools	62
Troubleshooting BIND	66
What's Next	71
4. An Initial Internet-Ready Environment	73
Installing ISPConfig	74
Setting Up a Server and Users with ISPConfig	83
Safeguarding a Linux Web Server	96
What's Next	101
5. Mail	102
Key Mail Service Terms	103
Postfix, Sendmail, and Other MTAs	103
The Postfix SMTP Mail Server on Debian	105
Adding Authentication and Encryption	111
Configuring POP3 and IMAP Mail Delivery Agents	119
Email Client Configuration	120
What's Next	121
6. Administering Apache	122
Static and Dynamic Files	122
A Simple LAMP Setup	123
Installation	124
Apache Configuration Files	127
Logfiles	140
SSL/TLS Encryption	142
suEXEC Support	143
Benchmarking	144
Installing and Administering Drupal	145
Troubleshooting	149
Further Reading	153

7. Load-Balanced Clusters	154
Load Balancing and High Availability	154
Scaling Without LB and HA	162
Further Reading	162
8. Local Network Services	163
Distributed Filesystems	164
Introduction to Samba	164
Configuring the Network	165
DHCP	168
Gateway Services	173
Print Services	181
User Management	186
9. Virtualization in the Modern Enterprise	194
Why Virtualization Is Popular	194
High-Performance Computing	196
Installing Xen on Fedora 5	199
Installing VMware	204
Virtualization: A Passing Fad?	210
10. Scripting	211
bash Beginnings	212
Useful Elements for bash Scripts	218
Scripting Language Shootout	226
Further Reading	235
11. Backing Up Data	236
Backing Up User Data to a Server with rsync	237
tar Archives	242
Saving Files on Optical Media	245
Backing Up and Archiving to Tape with Amanda	251
Backing Up MySQL Data	254
Appendix. bash Script Samples	257
Index	273



CHAPTER 1

Requirements for a Linux System Administrator

We like Linux. Of all the Unix and Unix-like systems we've used, many now forgotten,* Linux is our favorite. It's an excellent server platform, a good desktop, and the center of much innovation in the current computing world.

Linux probably has the broadest reach of any operating system, from tiny systems the size of phone jacks, to cell phones, to supercomputer clusters bigger than your high school. It has infiltrated the fields of telecommunications, embedded systems, satellites, medical equipment, military systems, computer graphics, and—last but not least—desktop computing.

In a relatively short time, Linux progressed from a Finnish hacker's hobby to a top-tier enterprise-level system backed by high rollers such as IBM and Oracle. The user base has grown from about 30,000 people in 1995 to hundreds of millions today. During the Internet boom of the 1990s, many Unix administrators were surprised to find that Linux on PC hardware could outperform more expensive Unix workstations and servers. Many Windows and Novell administrators saw that Linux could handle DNS, email, and file services more reliably and with less support personnel than their current platforms. The growth of the Internet, and especially the Web, fueled a rapid expansion in the use of Linux servers and the need for people to manage them.

This book is for Linux system administrators. However, you may be a grizzled Unix veteran, a brave MCSE, or a stoic mainframer. You're exploring new territory and need a map and compass. Some of the ground will be familiar, but some will be terra incognita. This book covers many topics that have only recently joined the mainstream, for instance load-balanced clusters and virtualization.

The success of the Internet and open source software is changing business. Google, Amazon, eBay, and others have built huge server farms with commodity hardware and relatively few administrators compared to traditional mainframe and PC installations.

* Our favorite name was PNX, pronounced almost like something that would never appear in an O'Reilly book.

The skills needed to develop and maintain such distributed systems and applications are not taught in schools but learned from experience, sometimes bitter and sometimes sweet.



While writing this book we've constantly tested the latest distributions and tools, and we'll keep up our experimentation after the book is released. We invite readers to come to the test site we set up for the book, <http://www.centralsoft.org>, where we'll publish updates to examples, pointers to useful new tools we've discovered, and other tips.

About This Book

System administration books used to be fairly predictable. They showed you how to manage users, filesystems, devices, processes, printers, networks, and so on. They did not tell you what to do when new problems emerged. If your web site became popular, you had to learn quickly about proxy servers, different levels of caching, load balancing, distributed authentication, and other complex issues. If you added a database, you soon needed to scale it and learn to avoid SQL injection attacks. Overnight, sites became mission critical, and you needed the ability to make hot backups on 24×7 systems.

If you've been through these fire drills, you may have become tired of doing everything the hard way, facing new technical challenges nearly every day with few sources of help. Technical documentation—whether for commercial or open source software—rarely keeps up with the technology, and the gap seems to be widening. For example, open source directory servers have become important for managing computers, users, and resources. The original RFC-compliant protocols underlie many commercial products, but good documentation for community projects is surprisingly scarce.

How Can We Help?

Linux people are problem solvers. A typical Linux power user can put together a small server, get a dedicated Internet pipe with static IP addresses into her home, register a domain name, and build a server on the Internet. If you fall into this category, you can simply plow through the other topics in this book and expand your job possibilities.

To some of you, however, all that may sound like the equivalent of rappelling down a 10,000-foot mountain. If you're one of them, just start somewhere. As the saying goes, you eat an elephant one bite at a time, and damn the torpedoes.

You may have certifications for operating systems other than Linux. While you're applying patches and hot fixes, your boss may ask you to deploy an Apache server, or handle your own DNS lookups, or replace Exchange with Zimbra.

Whether you just want to learn or actually *have* to learn, you'll likely need some help climbing the Linux power user curve. That's exactly what we're here for: to help you explore the Linux system landscape without all the hardships our forefathers experienced.

Where Do You Start?

This book summarizes the steps you need to follow to build standalone servers. If you need to build a mail server, create a web server and blogging system, or set up a gateway for your LAN, you can jump right into the middle of the book. You don't have to read *Linux System Administration* from cover to cover.

We start you working right away, presenting a step-by-step guide to building a Linux server in Chapter 2. You can choose whatever path works for you, whether it involves creating a highly available cluster for web services, server consolidation through virtualization using Xen or VMware, or setting up a server for local area networks.

Running a modern operating system is incredibly cheap. You can set up a sophisticated learning center for yourself on hardware that many sites would consider obsolete and give away for free. We started with a used box powered by an Intel CPU two generations older than current models, added older versions of hard drives and memory, and went with a no-frills, free version of Linux.

Do You Need a Book?

Technical books have waned in popularity as the Internet has matured. To write a successful book today, the author has to provide significant value to the reader. An interesting story about one of the first e-commerce sites on the Web helps explain the value a book should deliver. A cheesecake company put up an advertisement in the earliest days of the Web. According to the story, several months passed and the company didn't receive a single order. In an unusual move, the president of the company published the company's secret cheesecake recipe. Within hours, he began receiving calls on his toll-free line. People began ordering cheesecakes in large numbers. Consumers looked at the recipe, considered the effort required to make their own cheesecakes, and saw the value in buying them from the company.

Many of the ingredients for this book were scattered across the Internet, in mailing lists, forums, and discussion groups, while others were mined from books, periodicals, and the experiences of colleagues. We solved a number of problems whose solutions were completely undocumented in the course of researching this book, and we pass our lessons on to you.

Many excellent project sites have inadequate documentation. Developers work hard to provide excellent software for free, but prose often trails code for many reasons: lack of time, lack of resources, lack of interest, language barriers, and so on.

Together with our readers, editors, and reviewers, we hope we've decreased entropy slightly in this little corner of the computing world.

Who Needs You?

A few years ago, most Linux system administrators would have told you that they didn't choose their careers—Linux chose them. In the old days, Linux was like an adolescent Unix. Most Linux system administrators learned the ropes on single workstations and very small networks. Linux inherited some servers from Unix (BIND, Sendmail, Apache), but little office software and few applications. Today, Linux system administration involves thousands of packages and interoperability with other operating systems.

Who needs Linux administrators? The NASA Center for Computational Sciences (NCCS) at the Goddard Space Flight Center does. Its Linux-based high-performance computing (HPC) clusters are designed to dramatically increase throughput for applications ranging from studying weather and climate variability to simulating astrophysical phenomena. Linux supplements NCCS architecture designed to scale to as many as 40 trillion floating-point operations per second (TFLOPS) in its full configuration.

Linux runs more of the world's top supercomputers than any other operating system. In fact, as of this writing Linux runs an astonishing 75 percent of the top 500 supercomputers on the planet.* According to department heads at the Lawrence Livermore National Laboratory in Livermore, CA, Linux runs 10 of their massive systems, all of which are on the TOP500 List. Those systems include BlueGene/L, the world's most powerful supercomputer, and Thunder, which currently ranks nineteenth (<http://www.top500.org/list/2006/11/100>).

Help Wanted

Linux administrators are in high demand. To give you an idea of what's expected of them, we looked at a small selection of the tens of thousands of ads for Linux system administrators on a national job listing agency's web site. Here's a tiny snapshot of some of the jobs' responsibilities:

- Administer and manage large Linux server environment, with an emphasis on performance monitoring, tuning, and management.
- Oversee database physical design, administration, and documentation.
- Provide network troubleshooting, escalated service desk support, and proactive monitoring of mission-critical systems.

* See <http://www.top500.org/stats/28/osfam>.