

使用 *Perl* 实现系统管理自动化 (影印版)

第二版
Revised & Expanded



Automating

System Administration with Perl

O'REILLY®

東南大學出版社

David N. Blank-Edelman 著

第二版

使用Perl 实现系统管理自动化 (影印版)

Automating System Administration with Perl

O'REILLY®

Beijing • Cambridge • Farnham • Köln • Sebastopol • Taipei • Tokyo

O'Reilly Media, Inc. 授权东南大学出版社出版

东南大学出版社

图书在版编目 (CIP) 数据

使用 Perl 实现系统管理自动化: 英文 / (美) 艾德elman (Edelman, D.N.B.) 著. —影印本. —南京: 东南大学出版社, 2010.1

书名原文: Automating System Administration with Perl
ISBN 978-7-5641-1928-7

I . 使… II . 艾… III . PERL 语言—程序设计—英文
IV . TP312

中国版本图书馆 CIP 数据核字 (2009) 第 205664 号

江苏省版权局著作权合同登记

图字: 10-2009-244 号

©2009 by O'Reilly Media, Inc.

Reprint of the English Edition, jointly published by O'Reilly Media, Inc. and Southeast University Press, 2009. Authorized reprint of the original English edition, 2009 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 2009。

英文影印版由东南大学出版社出版 2009。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重制。

使用 Perl 实现系统管理自动化 (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

印 刷: 扬中市印刷有限公司

开 本: 787 毫米 × 980 毫米 16 开本

印 张: 41.75 印张

字 数: 701 千字

版 次: 2010 年 1 月第 1 版

印 次: 2010 年 1 月第 1 次印刷

书 号: ISBN 978-7-5641-1928-7

印 数: 1~1600 册

定 价: 88.00 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

使用Perl 实现系统管理自动化(影印版)

Automating System Administration with Perl

*To Cindy, ever the love of my life, and to Elijah,
a true blessing.*

Preface

Do you need tools for making your system administration work easier and more efficient? You've come to the right place.

Perl is a powerful programming language that grew out of the traditional system administration toolbox. Over the years it has adapted and expanded to meet the challenges of new operating systems and new tasks. If you know a little Perl, and you need to perform system administration tasks, this is the right book for you. Readers with varying levels of both Perl programming experience and system administration experience will all find something of use within these pages.

What's New in This Edition?

A tremendous amount of work went into updating this book so it could be even better than the first edition. Here's some of what has been improved in the second edition:

New title

My editors and I realized that the material in this book was more about how to automate your system administration work in ways that would make your working life more efficient and pleasant than it was about Perl. While Perl is still the toolshed that makes all this possible, it isn't the main focus of the book.

New material

It's hard to know where to begin on this one. The new edition is four chapters and two appendixes bigger (with a total page count that is 50% greater) than the last one. Included in this edition are a cornucopia of new tools and techniques that you are going to love. I tried to add material on the things I wished I had sysadmin-targeted material on, including: XML and YAML best practices (using `XML::LibXML`, `XML::Twig`, and `XPath`); dealing with config files; more advanced LDAP topics (including updated `Net::LDAP` information); email-related topics (including POP3/IMAP, MIME, and spam); new ways of dealing with filesystems; more advanced log file creation and parsing tools; DHCP; mapping/monitoring a network using Nmap and other tools; packet creation and sniffing; information reporting using tools like GraphViz, RRDtool, and Timeline; using SHA-2 instead of MD5;

SNMPv3; Mac OS X; converting VBScript code to Perl; geocoding; MP3 file manipulation; using Google Maps; and so on.

New advice

Part of the value of this book is the advice you can pick up from an experienced system administrator like me who has been doing this stuff for a long time and has compared notes with many other seasoned veterans. This new edition is packed with more sidebars to explain not only the *what*, but also the *why* behind the material.

Operating system and software information updates

All of the text and code has been updated and augmented to work with the latest versions of Unix- (including Linux and Mac OS X) and Windows-based operating systems.

Module and code updates/improvements

The descriptions and code in this book match the latest versions of the modules mentioned in the first edition. In cases where a module is no longer available or a better alternative has emerged, the appropriate replacement modules have been substituted. Also, all example code is now “use strict” friendly.

Errata corrected

I have attempted to address all of the errata I received from all of the printings of the first edition. I appreciate the time readers took to report errors to O’Reilly and me so I could fix them at each printing and in this edition. Special thanks go to Andreas Karrer, the German translator for the first edition. Andi pored over every single byte of the original text and submitted almost 200 (mostly layout-related) corrections, all with good cheer.

How This Book Is Structured

Each chapter in this book addresses a different system administration domain and ends with a list of the Perl modules used in that chapter and references to facilitate deeper exploration of the information presented. The chapters are as follows:

Chapter 1, Introduction

This introductory chapter describes the material covered in the book in more detail, explaining how it will serve you and what you need to get the most from it. The material in this book is powerful and is meant to be used by powerful people (e.g., Unix superusers and Windows-based operating system administrators). The introduction provides some important guidelines to help you write more secure Perl programs.

Chapter 2, Filesystems

This chapter is about keeping multiplatform filesystems tidy and ensuring that they are used properly. We’ll start by looking at the salient differences between the native filesystems for each operating system. We’ll then explore the process of

intelligently walking or traversing filesystems from Perl and how that can be useful. Finally, we'll look at manipulating disk quotas from Perl.

Chapter 3, User Accounts

This chapter discusses how user accounts manifest themselves on two different operating systems, including what is stored for each user and how to manipulate the information from Perl. That leads into a discussion of a rudimentary account system written in Perl. In the process of building this system, we'll examine the mechanisms necessary for recording accounts in a simple database, creating these accounts, and deleting them.

Chapter 4, User Activity

Chapter 4 explores ways to automate tasks centered around user activity, introducing a number of ways to track and control process, file, and network operations initiated by users. This chapter also presents various operating system-specific frameworks and tools (e.g., Windows Management Instrumentation, GUI setup tools, *lsOf*, etc.) that are helpful for user-oriented tasks on different platforms.

Chapter 5, TCP/IP Name and Configuration Services

Name and configuration services allow hosts on a TCP/IP network to communicate with each other amicably and to self-configure. This chapter takes a historical perspective by starting with host files, then moving on to the Network Information Service (NIS) and finally to the glue of the Internet, the Domain Name Service (DNS). Each step of the way, it shows how Perl can make professional management of these services easier. We'll also explore how to work with the Dynamic Host Configuration Protocol (DHCP) from Perl in this chapter.

Chapter 6, Working with Configuration Files

Almost every system or software package we touch relies heavily on configuration files to be useful in our environment. This chapter explores the tools that make writing and reading those files from Perl easy. We'll look at various formats, with special attention paid to XML and the current best practices for working with it using Perl.

Chapter 7, SQL Database Administration

Over time, more uses for relational databases are being found in the system administration realm. As a result, system administrators need to become familiar with SQL database administration. This chapter explains DBI, the preeminent SQL database framework for Perl, and provides examples of it in action for database administration.

Chapter 8, Email

This chapter demonstrates how Perl can make better use of email as a system administration tool. After discussing sending via SMTP (including MIME-based HTML messages), receiving via POP3/IMAP, and parsing via Perl, we'll explore several interesting applications, including tools for analyzing unsolicited commercial email (a.k.a. spam) and managing tech support emails.

Chapter 9, Directory Services

As the complexity of the information we deal with increases over time, so does the importance of the directory services we use to access that information. System administrators are increasingly being called upon not only to use these services, but also to build tools for their management. This chapter discusses some of the more popular directory service protocols/frameworks, such as LDAP and ADSI, and shows you how to work with them from Perl.

Chapter 10, Log Files

System administrators are often awash in a sea of log files. Every machine, operating system, and program can (and often does) log information. This chapter looks at the logging systems offered by Unix- and Windows-based operating systems and discusses approaches for analyzing logging information so it can work for you.

Chapter 11, Security

This chapter heads right into the maelstrom called “security,” demonstrating how Perl can make hosts and networks more secure.

Chapter 12, SNMP

This chapter is devoted to the Simple Network Management Protocol (SNMP). It illustrates how to use this protocol to communicate with network devices (both to poll and to receive trap information).

Chapter 13, Network Mapping and Monitoring

Perl offers some excellent tools for the mapping and monitoring of networks. In this chapter, we’ll look at several ways to discover the hosts on the network and the services they offer. We’ll then explore helpful graphical and textual ways to present the information collected, including some of the best tools for graphing and charting the data (such as GraphViz and RRDtool).

Chapter 14, Experiential Learning

This is the chapter you don’t want your boss to catch you reading.

Appendixes

Some of the chapters assume basic knowledge about topics with which you may not be familiar. For those who are new to these subjects, this book includes several mini-tutorials to bring you up to speed quickly. The appendixes provide introductions to the eXtensible Markup Language (XML), the XML Path Language (XPath), the Lightweight Directory Access Protocol (LDAP), the Structured Query Language (SQL), the Revision Control System (RCS), translating VBScript to Perl, and SNMP.

Typographical Conventions

This book uses the following typographical conventions:

Italic

Used for file- and pathnames, usernames, directories, program names, hostnames, URLs, and new terms where they are first introduced.

Constant width

Used for Perl module and function names, namespaces, libraries, commands, methods, and variables, and when showing code and computer output.

Constant width bold

Used to indicate user input and for emphasis in code examples.

Constant width italic

Used to indicate parts of a command line that are user-replaceable, and for code annotations.



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



This icon indicates a warning or caution.

Operating System Naming Conventions

This book is steadfastly multiplatform in its thinking. However, reading about “a Microsoft Vista/Microsoft Windows Server 2008/Microsoft Windows Server 2003/Microsoft XP script” or a “Linux/Solaris/Irix/HPUX/Mac OS X/etc. script” gets old fast. Having consulted some style guides, here’s how I’ve chosen to handle discussing the operating system collectives:

- When writing about the Microsoft products—Microsoft Vista, Microsoft Windows Server 2008, Microsoft Windows Server 2003, and Microsoft XP (on which, by and large, all scripts were tested)—I refer to them collectively as “Windows-based operating systems,” at least first time they show up in a chapter or heading. From that point on in the chapter I shorten this to simply “Windows.” If something is particular to a specific Windows-based operating system, I will mention it by name.

- When writing about any of the members of the Unix family (in which I include both Linux and Mac OS X), I refer to them collectively as just “Unix,” “the Unix family,” or sometimes “Unix variants.” If something is particular to a specific Unix vendor or release, I will mention it by name.

Coding Conventions

There are a few points I want to mention about the code in this book:

- All code examples were written and tested with `use strict`; as the first line (I highly recommend you do the same). However, given the number of examples in this book, the repetition would have taken up a significant amount of space, so to save trees and wasted bits I did not include that line in any of the examples. Please just assume that every example uses this convention.
- Almost all of the code is formatted using Steve Hancock’s fabulous *perltidy* utility (<http://search.cpan.org/dist/Perl-Tidy/>) to improve readability.
- Although these examples don’t reach anything like that level of perfection, much of the code has been rewritten with the advice in Damian Conway’s book *Perl Best Practices* (<http://oreilly.com/catalog/9780596001735/>) (O’Reilly) in mind. I highly recommend reading Conway’s book to improve your code and generally reinvigorate your Perl programming. The automated source code analyzer `Perl::Critic` that *Perl Best Practices* inspired was still in heavy development for much of the writing of this book, so I did not use it. You should, though, as it’s another great tool.

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: “*Automating System Administration with Perl*, Second Edition, by David N. Blank-Edelman. Copyright 2009 O’Reilly Media, Inc., 978-0-596-00639-6.”

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.

How to Contact Us

We have tested and verified the information in this book to the best of our ability, but you may find that features have changed (or even that we have made mistakes!). Please let us know of any errors you find, as well as your suggestions for future editions, by writing to:

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

We have a website for the book, where we'll list examples, errata, and any plans for future editions. You can access this page at:

<http://www.oreilly.com/catalog/9780596006396/>

The author has set up a personal website for this book. Please visit it at:

<http://www.otterbook.com>

To ask technical questions or comment on the book, send email to:

bookquestions@oreilly.com

For more information about our books, conferences, software, Resource Centers, and the O'Reilly Network, see the O'Reilly website:

<http://www.oreilly.com>

Safari® Books Online

Safari[®] Books Online When you see a Safari® Books Online 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://my.safaribooksonline.com/>.

Acknowledgments from the First Edition

To keep the preface from becoming too much like an Oscar acceptance speech, here's a condensed version of the acknowledgments from the first edition.

Thanks to the Perl Community, especially Larry Wall, Tom Christiansen, and the kerjillions of programmers and hackers who poured countless hours and energy into the language and then chose to share their work with me and the rest of the Perl community.

Thanks to the SysAdmin community: members of Usenix, SAGE, and the people who have contributed to the LISA conferences over the years. Thanks to Rémy Evard for being such a great influence on my professional and personal understanding of this field as a friend, mentor, and role model. He is still one of the system administrators I want to be when I grow up.

Thanks to the reviewers of the first edition: Jerry Carter, Toby Everett, Æleen Frisch, Joe Johnston, Tom Limoncelli, John A. Montgomery, Jr., Chris Nandor, Michael Peppler, Michael Stok, and Nathan Torkington.

Thanks to the O'Reilly staff: to Rhon Porter for his illustrations, to Hanna Dyer and Lorrië LeJeune for the most amazing cover animal, and to the O'Reilly production staff. I am still thankful to Linda Mui, my editor for the first edition, whose incredible skill, finesse, and care allowed me to birth this book and raise it in a good home.

Thanks to my spiritual community: Havurat Shalom in Somerville. Thank you, M' kor HaChayim, for this book and all of the many blessings in my life.

Thanks to the Shona people of Zimbabwe for their incredible mbira music.

Thanks to my friends (Avner, Ellen, Phil Shapiro, Alex Skovronek, Jon Orwant, and Joel Segel), the faculty and staff at the Northeastern University College of Computer and Information Science (especially the folks in the CCIS Systems group), and my boss Larry Finkelstein, the Dean of the College of Computer Science.

Thanks to my nuclear family (Myra, Jason, and Steven Edelman-Blank), my cats Shimmer and Bendir (bye-bye, Bendir, I'll miss you), and my TCM pit crew (Kristen Porter and Thom Donovan).

The first edition was dedicated to Cindy, love of my life.

Acknowledgments for the Second Edition

One of the only things better than having all of these great people and things in your life is to have them remain in your life. I'm still thankful for all of the above from the first edition. Here are some of the changes and additions:

This edition had a much expanded and tremendous group of technical reviewers. I'm very grateful to Æleen Frisch, Aaron Crane, Aleksey Tsalolikhin, Andrew Langmead, Bill Cole, Cat Okita, Chaos Golubitsky, Charles Richmond, Chris Grau, Clifton Royston, Dan Wilson, Dean Wilson, Denny Allain, Derek J. Balling, Earl Gay, Eric Sorenson, Eric Toczek, Federico Lucifredi, Gordon "Fyodor" Lyon, Graham Barr, Grant McLean, Hugh Brown, James Keating, Jan Dubois, Jennifer Davis, Jerry Carter, Jesse Vincent, Joe Morri, John Levine, John Tsangaris, Josh Roberts, Justin Mason, Mark

Bergman, Michel Rodriguez, Mike DeGraw-Bertsch, Mike Stok, Neil Neely, Petr Pajas, Philip J. Hollenback, Randy Dees, Scott Murphy, Shlomi Fish, Stephen Potter, Steve Atkins, Steven Tylock, Terry Zink, Thomas Leyer, Tim Bunce, Tobias Oetiker, Toby Ovod-Everett, and Tom Regner for the time and energy they spent on making this book better. I continue to be amazed by the generosity and kindness shown by the members of the SysAdmin and Perl communities.

The editorial chain was a bit longer than usual on this book, so thanks to *all* of the editors. Starting from the first edition in chronological order: Linda Mui, Paula Ferguson, Nathan Torkington, Allison Randal, Colleen Gorman, Tatiana Apandi, Isabel Kunkle, and Andy Oram. I'm also thankful to the other O'Reilly people who have had a hand in bringing this book to fruition, including Mike Hendrickson, Rachel Head, Sarah Schneider, Rob Romano, Sanders Kleinfeld, and all the others.

I was taken with sea otters even before the first edition was published with one on the front cover, but since then my appreciation for them keeps on growing. They are an amazing species in so many ways. Unfortunately, humans historically haven't been particularly kind to the sea otters. They are still classified as an endangered species, and some of our activities actively threaten their survival. I believe they deserve our protection and our support. One organization that works toward this end is Friends of the Sea Otter (<http://www.seaotters.org>), based in Monterey, California. I'm a member, and I encourage you to join, too.

Mbira kept me sane through the arduous process of writing the first edition of this book. For this edition, I have yoga to thank for my current health and sanity. I'd like to express my profound gratitude to my teacher, Karin Stephan, and her teacher, B.K.S. Iyengar, for sharing such a wonderful union of mind and body with me.

I've tried to cut down the florid prose of the first edition's acknowledgments, but I hope you'll indulge me just one more time. The biggest change for me between these editions was the birth of our first child, Elijah. He's been a constant blessing to us, both in the noun and verb senses of the word.

Table of Contents

Preface	xv
1. Introduction	1
Automation Is a Must	1
How Perl Can Help You	2
This Book Will Show You How	3
What You Need	5
Some Notes About the Perl Versions Used for This Book	6
What About Perl 5.10?	6
What About Strawberry Perl?	6
What About Perl 6?	6
Some Notes About Using Vista with the Code in This Book	7
Locating and Installing Modules	8
Installing Modules on Unix	9
Installing Modules on Win32	9
It's Not Easy Being Omnipotent	10
Don't Do It	10
Drop Your Privileges As Soon As Possible	10
Be Careful When Reading Data	11
Be Careful When Writing Data	12
Avoid Race Conditions	12
Enjoy	13
References for More Information	13
2. Filesystems	15
Perl to the Rescue	15
Filesystem Differences	16
Unix	16
Windows-Based Operating Systems	16
Mac OS X	18
Filesystem Differences Summary	19
Dealing with Filesystem Differences from Perl	19

Walking or Traversing the Filesystem by Hand	21
Walking the Filesystem Using the File::Find Module	26
Walking the Filesystem Using the File::Find::Rule Module	36
Manipulating Disk Quotas	38
Editing Quotas with edquota Trickery	40
Editing Quotas Using the Quota Module	44
Editing NTFS Quotas Under Windows	45
Querying Filesystem Usage	46
Module Information for This Chapter	48
References for More Information	48
3. User Accounts	49
Unix User Identities	50
The Classic Unix Password File	50
Changes to the Password File in BSD 4.4 Systems	57
Shadow Passwords	58
Windows-Based Operating System User Identities	59
Windows User Identity Storage and Access	59
Windows User ID Numbers	61
Windows Passwords Don't Play Nice with Unix Passwords	63
Windows Groups	63
Windows User Rights	68
Building an Account System to Manage Users	71
The Backend Database	73
The Low-Level Component Library	78
The Process Scripts	89
Account System Wrap-Up	94
Module Information for This Chapter	97
References for More Information	97
Unix Password Files	97
Windows User Administration	98
4. User Activity	99
Process Management	100
Windows-Based Operating System Process Control	100
Unix Process Control	119
File and Network Operations	125
Tracking File Operations on Windows	125
Tracking Network Operations on Windows	128
Tracking File and Network Operations in Unix	129
Module Information for This Chapter	135
Installing Win32::Setupsup	135
References for More Information	136

5. TCP/IP Name and Configuration Services	137
Host Files	137
Generating Host Files	140
Error-Checking the Host File Generation Process	143
Improving the Host File Output	144
Incorporating a Source Code Control System	148
NIS, NIS+, and WINS	151
NIS+	154
Windows Internet Name Server (WINS)	154
Domain Name Service (DNS)	155
Generating DNS (BIND) Configuration Files	156
DNS Checking: An Iterative Approach	165
DHCP	174
Active Probing for Rogue DHCP Servers	176
Monitoring Legitimate DHCP Servers	181
Module Information for This Chapter	183
References for More Information	184
 6. Working with Configuration Files	 185
Configuration File Formats	188
Binary	188
Naked Delimited Data	189
Key/Value Pairs	190
Markup Languages	192
All-in-One Modules	235
Advanced Configuration Storage Mechanisms	236
Module Information for This Chapter	236
References for More Information	237
XML and YAML	237
 7. SQL Database Administration	 239
Interacting with a SQL Server from Perl	240
Using the DBI Framework	243
Using ODBC from Within DBI	249
Server Documentation	251
MySQL Server via DBI	252
Oracle Server via DBI	254
Microsoft SQL Server via ODBC	255
Database Logins	258
Monitoring Space Usage on a Database Server	260
Module Information for This Chapter	263
References for More Information	263
DBI	263