经 典 原 版 书 库

计算机取证

（英文版）

Forensic Discovery

Dan Farmer
Wietse Venema

ADDISON-WESLEY PROFESSIONAL COMPUTING SERIES

（美） Dan Farmer
Wietse Venema 著

机械工业出版社
China Machine Press

经 典 原 版 书 库

计算机取证

（英文版）

Forensic Discovery

（美） Dan Farmer
Wietse Venema 著

机械工业出版社
China Machine Press

本书版权登记号：图字：01-2005-4694

图书在版编目（CIP）数据

计算机取证（英文版）/（美）法默（Farmer, D.）等著. －北京：机械工业出版社，2006.1

（经典原版书库）

书名原文：Forensic Discovery

ISBN 7-111-17365-1

Ⅰ.计… Ⅱ.法… Ⅲ.计算机犯罪－证据－调查－英文 Ⅳ.D915.13

中国版本图书馆CIP数据核字（2005）第104939号

机械工业出版社（北京市西城区百万庄大街22号　邮政编码　100037）

责任编辑：迟振春

北京京北制版厂印刷 · 新华书店北京发行所发行

2006年1月第1版第1次印刷

718mm × 1020mm 1/16 · 14.75印张

印数：0 001 - 3 000册

定价：30.00元

凡购本书，如有倒页、脱页、缺页，由本社发行部调换

本社购书热线：（010）68326294

出版者的话

文艺复兴以降，源远流长的科学精神和逐步形成的学术规范，使西方国家在自然科学的各个领域取得了垄断性的优势；也正是这样的传统，使美国在信息技术发展的六十多年间名家辈出、独领风骚。在商业化的进程中，美国的产业界与教育界越来越紧密地结合，计算机学科中的许多泰山北斗同时身处科研和教学的最前线，由此而产生的经典科学著作，不仅擘划了研究的范畴，还揭橥了学术的源变，既遵循学术规范，又自有学者个性，其价值并不会因年月的流逝而减退。

近年，在全球信息化大潮的推动下，我国的计算机产业发展迅猛，对专业人才的需求日益迫切。这对计算机教育界和出版界都既是机遇，也是挑战；而专业教材的建设在教育战略上显得举足轻重。在我国信息技术发展时间较短、从业人员较少的现状下，美国等发达国家在其计算机科学发展的几十年间积淀的经典教材仍有许多值得借鉴之处。因此，引进一批国外优秀计算机教材将对我国计算机教育事业的发展起积极的推动作用，也是与世界接轨、建设真正的世界一流大学的必由之路。

机械工业出版社华章图文信息有限公司较早意识到“出版要为教育服务”。自1998年开始，华章公司就将工作重点放在了遴选、移译国外优秀教材上。经过几年的不懈努力，我们与Prentice Hall，Addison-Wesley，McGraw-Hill，Morgan Kaufmann等世界著名出版公司建立了良好的合作关系，从它们现有的数百种教材中甄选出Tanenbaum，Stroustrup，Kernighan，Jim Gray等大师名家的一批经典作品，以“计算机科学丛书”为总称出版，供读者学习、研究及庋藏。大理石纹理的封面，也正体现了这套丛书的品位和格调。

“计算机科学丛书”的出版工作得到了国内外学者的鼎力襄助，国内的专家不仅提供了中肯的选题指导，还不辞劳苦地担任了翻译和审校的工作；而原书的作者也相当关注其作品在中国的传播，有的还专程为其书的中译本作序。迄今，“计算机科学丛书”已经出版了近百个品种，这些书籍在读者中树立了良好的口碑，并被许多高校采用为正式教材和参考书籍，为进一步推广与发展打下了坚实的基础。

随着学科建设的初步完善和教材改革的逐渐深化，教育界对国外计算机教材的需求和应用都步入一个新的阶段。为此，华章公司将加大引进教材的力度，在“华章教育”的总规划之下出版三个系列的计算机教材：除“计算机科学丛书”之外，对影印版的教材，则单独开辟出“经典原版书库”；同时，引进全美通行的教学辅导书“Schaum's Outlines”系列组成“全美经典学习指导系列”。为了保证这三套丛书的权威性，同时也为了更好地为学校和老师们服务，华章公司聘请了中国科学院、北京大学、清华大学、国防科技大学、复旦大学、上海交通大学、南京大学、浙江大学、中国科技大学、哈尔

滨工业大学、西安交通大学、中国人民大学、北京航空航天大学、北京邮电大学、中山大学、解放军理工大学、郑州大学、湖北工学院、中国国家信息安全测评认证中心等国内重点大学和科研机构在计算机的各个领域的著名学者组成“专家指导委员会”，为我们提供选题意见和出版监督。

这三套丛书是响应教育部提出的使用外版教材的号召，为国内高校的计算机及相关专业的教学度身订造的。其中许多教材均已为M. I. T.，Stanford，U.C. Berkeley，C. M. U. 等世界名牌大学所采用。不仅涵盖了程序设计、数据结构、操作系统、计算机体系结构、数据库、编译原理、软件工程、图形学、通信与网络、离散数学等国内大学计算机专业普遍开设的核心课程，而且各具特色——有的出自语言设计者之手、有的历经三十年而不衰、有的已被全世界的几百所高校采用。在这些圆熟通博的名师大作的指引之下，读者必将在计算机科学的宫殿中由登堂而入室。

权威的作者、经典的教材、一流的译者、严格的审校、精细的编辑，这些因素使我们的图书有了质量的保证，但我们的目标是尽善尽美，而反馈的意见正是我们达到这一终极目标的重要帮助。教材的出版只是我们的后续服务的起点。华章公司欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方法如下：

电子邮件：hzjsj@hzbook.com
联系电话：(010) 68995264
联系地址：北京市西城区百万庄南街1号
邮政编码：100037

专家指导委员会

（按姓氏笔画顺序）

尤晋元	王　珊	冯博琴	史忠植	史美林
石教英	吕　建	孙玉芳	吴世忠	吴时霖
张立昂	李伟琴	李师贤	李建中	杨冬青
邵维忠	陆丽娜	陆鑫达	陈向群	周伯生
周克定	周傲英	孟小峰	岳丽华	范　明
郑国梁	施伯乐	钟玉琢	唐世渭	袁崇义
高传善	梅　宏	程　旭	程时端	谢希仁
裘宗燕	戴　葵			

For feminists
—Dan

For the people who refuse to stop learning
—Wietse

Today, only minutes pass between plugging in to the Internet and being attacked by some other machine—and that's only the background noise level of nontargeted attacks. There was a time when a computer could tick away year after year without coming under attack. For examples of Internet background radiation studies, see CAIDA 2003, Cymru 2004, or IMS 2004.

With this book, we summarize experiences in post-mortem intrusion analysis that we accumulated over a decade. During this period, the Internet grew explosively, from less than a hundred thousand connected hosts to more than a hundred million (ISC 2004). This increase in the number of connected hosts led to an even more dramatic—if less surprising—increase in the frequency of computer and network intrusions. As the network changed character and scope, so did the character and scope of the intrusions that we faced. We're pleased to share some of these learning opportunities with our readers.

In that same decade, however, little changed in the way that computer systems handle information. In fact, we feel that it is safe to claim that computer systems haven't changed fundamentally in the last 35 years—the entire lifetime of the Internet and of many operating systems that are in use today, including Linux, Windows, and many others. Although our observations are derived from today's systems, we optimistically expect that at least some of our insights will remain valid for another decade.

What You Can Expect to Learn from This Book

The premise of the book is that forensic information can be found everywhere you look. With this guiding principle in mind, we develop tools to collect information from obvious and not-so-obvious sources, we walk through analyses of real intrusions in detail, and we discuss the limitations of our approach.

Although we illustrate our approach with particular forensic tools in specific system environments, we do not provide cookbooks for how to use those tools, nor do we offer checklists for step-by-step investigation. Instead, we present a background on how information persists, how information about past events may be recovered, and how the trustworthiness of that information may be affected by deliberate or accidental processes.

In our case studies and examples, we deviate from traditional computer forensics and head toward the study of system dynamics. Volatility and the persistence of file systems and memory are pervasive topics in our book. And while the majority of our examples are from Solaris, FreeBSD, and Linux systems, Microsoft's Windows shows up on occasion as well. Our emphasis is on the underlying principles that these systems have in common: we look for inherent properties of computer systems, rather than accidental differences or superficial features.

Our global themes are problem solving, analysis, and discovery, with a focus on reconstruction of past events. This approach may help you to discover *why* events transpired, but that is generally outside the scope of this work. Knowing *what* happened will leave you better prepared the next time something bad is about to occur, even when that knowledge is not sufficient to prevent future problems. We should note up front, however, that we do not cover the detection or prevention of intrusions. We do show that traces from one intrusion can lead to the discovery of other intrusions, and we point out how forensic information may be affected by system-protection mechanisms, and by the failures of those mechanisms.

Our Intended Audience

We wrote this book for readers who want to deepen their understanding of how computer systems work, as well as for those who are likely to become involved with the technical aspects of computer intrusion or system analysis. System administrators, incident responders, other computer security professionals, and forensic analysts will benefit from reading this book, but so will anyone who is concerned about the impact of computer forensics on privacy.

Although we have worked hard to make the material accessible to non-expert readers, we definitely do not target the novice computer user. As a minimal requirement, we assume strong familiarity with the basic concepts of UNIX or Windows file systems, networking, and processes.

Organization of This Book

The book has three parts: we present foundations first, proceed with analysis of processes, systems, and files, and end the book with discovery. We do not expect you to read everything in the order presented. Nevertheless, we suggest that you start with the first chapter, as it introduces all the major themes that return throughout the book.

In Part I, "Basic Concepts," we introduce general high-level ideas, as well as basic techniques that we rely on in later chapters.

- Chapter 1, "The Spirit of Forensic Discovery," shows how general properties of computer architecture can impact post-mortem analysis. Many of the limitations and surprises that we encounter later in the book can already be anticipated by reading this chapter.
- Chapter 2, "Time Machines," introduces the concept of timelining, using examples of host-based and network-based information, including information from the domain name system. We look at an intrusion that stretches out over an entire year, and we show examples of finding time information in non-obvious places.

In Part II, "Exploring System Abstractions," we delve into the abstractions of file systems, processes, and operating systems. The focus of these chapters is on analysis: making sense of information found on a computer system and judging the trustworthiness of our findings.

- Chapter 3, "File System Basics," introduces fundamental file system concepts, as well as forensic tools and techniques that we will use in subsequent chapters.
- Chapter 4, "File System Analysis," unravels an intrusion by examining the file system of a compromised machine in detail. We look at both existing files and deleted information. As in Chapter 2, we use correlation to connect different observations, and to determine their consistency.
- Chapter 5, "Systems and Subversion," is about the environment in which user processes and operating systems execute. We look at subversion of observations, ranging from straightforward changes to system utilities to almost undetectable malicious kernel modules, and detection of such subversion.
- Chapter 6, "Malware Analysis Basics," presents techniques to discover the purpose of a process or a program file that was left behind after an intrusion. We also discuss safeguards to prevent malware from escaping, and their limitations.

In Part III, "Beyond the Abstractions," we look beyond the constraints of the file, process, and operating system abstractions. The focus of this part is on discovery, as we study the effects of system architecture on the decay of information.

- Chapter 7, "The Persistence of Deleted File Information," shows that large amounts of deleted file information can survive intact for extended periods. We find half-lives on the order of two to four weeks on actively used file systems.

- Chapter 8, "Beyond Processes," shows examples of persistence of information in main memory, including the decrypted contents of encrypted files. We find large variations in persistence, and we correlate these variations to operating system architecture properties.

The appendices present background material: Appendix A is an introduction to the Coroner's Toolkit and related software. Appendix B presents our current insights with respect to the order of volatility and its ramifications when capturing forensic information from a computer system.

Conventions Used in This Book

In the examples, we use `constant-width font` for program code, command names, and command input/output. User input is shown in **`bold constant-width font`**. We use `$` as the shell command prompt for unprivileged users, and we reserve `#` for super-user shells. Capitalized names, such as Argus, are used when we write about a system instead of individual commands.

Whenever we write "UNIX," we implicitly refer to Solaris, FreeBSD, and Linux. In some examples we include the operating system name in the command prompt. For example, we use `solaris$` to indicate that an example is specific to Solaris systems.

As hinted at earlier, many examples in this book are taken from real-life intrusions. To protect privacy, we anonymize information about systems that are not our own. For example, we replace real network addresses with private network addresses such as 10.0.0.1 or 192.168.0.1, and we replace host names or user names. Where appropriate, we even replace the time and time zone.

Web Sites

The examples in this book feature several small programs that were written for the purpose of discovery and analysis. Often we were unable to include the entire code listing because the additional detail would only detract from the purpose of the book. The complete source code for these and other programs is made available online at these Web sites:

http://www.fish.com/forensics/
http://www.porcupine.org/forensics/

On the same Web sites, you will also find bonus material, such as case studies that were not included in the book and pointers to other resources.

Acknowledgments

We owe a great deal of gratitude to Karen Gettman, Brian Kernighan, and the rest of Addison-Wesley for their patience and support over the many years that this book has been under construction.

While we take full responsibility for any mistakes, this book would not be what it is without our review team. In particular, we would like to thank (in alphabetical order): Aleph1, Muffy Barkocy, Brian Carrier, Eoghan Casey, Fred Cohen, Rik Farrow, Gary McGraw, Brad Powell, Steve Romig, Douglas Schales, and Elizabeth Zwicky. Ben Pfaff and Jim Chow helped with a chapter, and Dalya Sachs provided valuable assistance with editing an early version of the text. Tsutomu Shimumura inspired us to do things that we thought were beyond our skills. Wietse would like to thank the FIRST community for the opportunity to use them as a sounding board for many of the ideas that were developed for this book. And contrary to current practice, the manuscript was produced as HTML draft with the vi text editor plus a host of little custom scripts and standard UNIX tools that helped us finish the book.

Dan Farmer
zen@fish.com

Wietse Venema
wietse@porcupine.org

ABOUT THE AUTHORS

Dan Farmer is the author or coauthor of a variety of security programs and papers. He's currently the chief technical officer of Elemental Security, a computer security software company.

Wietse Venema is the author of widely used software such as the TCP Wrapper and the Postfix mail system. Originally from the Netherlands, he is currently a research staff member at IBM Research in the United States.

The cooperation between the authors goes back many years and has resulted in famous and notorious software such as the SATAN network security scanner and the Coroner's Toolkit for forensic analysis.

CONTENTS