

THE DESIGN OF THE UNIX[®] OPERATING SYSTEM

MAURICE J. BACH

In this timely new book, **Maurice J. Bach** traces the popularity of the UNIX[®] System throughout the computer industry. The author describes the internal algorithms and structures that form the basis of the operating system (the kernel) and their relationship to the programmer interface.

Among its key features, the book:

- describes the outline of the kernel architecture
- introduces the system buffer cache mechanism
- includes data structures and algorithms used internally by the file system
- covers the system calls that provide the user interface to the file system
- defines the context of a process and investigates the internal kernel primitives that manipulate the process context
- presents the system calls that control the process context
- describes process scheduling
- discusses memory management, including swapping and paging systems
- outlines general driver interfaces, with specific discussion of disk drivers and terminal drivers
- presents an overview of streams
- introduces inter-process communication and networking, including System V messages, shared memory, and semaphores
- explains tightly coupled multiprocessor UNIX[®] systems
- investigates distributed UNIX[®] systems

PRENTICE-HALL, INC., Englewood Cliffs, N.J. 07632

ISBN 0-13-201799-7

CH
THE DESIGN OF THE UNIX[®]
OPERATING SYSTEM

PRENTICE-HALL



THE DESIGN OPERATION

Maurice J. Bach

PRENTICE-HALL, INC.

Published by Prentice-Hall, Inc.
A division of Simon & Schuster
Englewood Cliffs, New Jersey 07632

Prentice-Hall Software Series
Brian W. Kernighan, Advisor

Cover Design Consultant: Sarah Bach

UNIX® is a registered trademark of AT&T.
DEC, PDP, and VAX are trademarks of Digital Equipment Corp.
Series 32000 is a trademark of National Semiconductor Corp.
®Ada is a registered trademark of the U.S. Government (Ada Joint Program Office).
UNIVAC is a trademark of Sperry Corp.
This document was set on an AUTOLOGIC, Inc. APS-5 phototypesetter driven by the
TROFF formatter operating under the UNIX system on an AT&T 3B20 computer.

The Publisher offers discounts on this book when ordered in bulk quantities. For more information write:

Special Sales/College Marketing
Prentice-Hall, Inc.
College Technical and Reference Division
Englewood Cliffs, New Jersey 07632

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America

10

ISBN 0-13-201799-7 025

Prentice-Hall International (UK) Limited, *London*
Prentice-Hall of Australia Pty. Limited, *Sydney*
Prentice-Hall Canada Inc., *Toronto*
Prentice-Hall Hispanoamericana, S.A., *Mexico*
Prentice-Hall of India Private Limited, *New Delhi*
Prentice-Hall of Japan, Inc., *Tokyo*
Prentice-Hall of Southeast Asia Pte. Ltd., *Singapore*
Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

003

To my
to my dau
to my son
and to my

CONT

PREFACE .

CHAPTER 1

- 1.1 HI
- 1.2 SY
- 1.3 US
- 1.4 OP
- 1.5 AS
- 1.6 SU

| | |
|--|----|
| CHAPTER 2 INTRODUCTION TO THE KERNEL | 19 |
| 2.1 ARCHITECTURE OF THE UNIX OPERATING SYSTEM | 19 |
| 2.2 INTRODUCTION TO SYSTEM CONCEPTS | 22 |
| 2.3 KERNEL DATA STRUCTURES | 34 |
| 2.4 SYSTEM ADMINISTRATION | 34 |
| 2.5 SUMMARY AND PREVIEW | 36 |
| 2.6 EXERCISES | 37 |
| | |
| CHAPTER 3 THE BUFFER CACHE | 38 |
| 3.1 BUFFER HEADERS | 39 |
| 3.2 STRUCTURE OF THE BUFFER POOL | 40 |
| 3.3 SCENARIOS FOR RETRIEVAL OF A BUFFER | 42 |
| 3.4 READING AND WRITING DISK BLOCKS | 53 |
| 3.5 ADVANTAGES AND DISADVANTAGES OF THE BUFFER CACHE | 56 |
| 3.6 SUMMARY | 57 |
| 3.7 EXERCISES | 58 |
| | |
| CHAPTER 4 INTERNAL REPRESENTATION OF FILES | 60 |
| 4.1 INODES | 61 |
| 4.2 STRUCTURE OF A REGULAR FILE | 67 |
| 4.3 DIRECTORIES | 73 |
| 4.4 CONVERSION OF A PATH NAME TO AN INODE | 74 |
| 4.5 SUPER BLOCK | 76 |
| 4.6 INODE ASSIGNMENT TO A NEW FILE | 77 |
| 4.7 ALLOCATION OF DISK BLOCKS | 84 |
| 4.8 OTHER FILE TYPES | 88 |
| 4.9 SUMMARY | 88 |
| 4.10 EXERCISES | 89 |

| | |
|-----------|---------|
| CHAPTER 5 | 5.1 OI |
| | 5.2 RI |
| | 5.3 W |
| | 5.4 FI |
| | 5.5 AI |
| | 5.6 CI |
| | 5.7 FI |
| | 5.8 CI |
| | 5.9 CI |
| | 5.10 CI |
| | 5.11 ST |
| | 5.12 PI |
| | 5.13 DI |
| | 5.14 M |
| | 5.15 LI |
| | 5.16 U |
| | 5.17 FI |
| | 5.18 FI |
| | 5.19 S |
| | 5.20 E |
| | |
| CHAPTER 6 | |
| | 6.1 PR |
| | 6.2 LA |
| | 6.3 TH |
| | 6.4 SA |
| | 6.5 M. |
| | SP |
| | 6.6 SL |

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.