



Introduction

Welcome to *UNIX Filesystems—Evolution, Design, and Implementation*, the first book that is solely dedicated to UNIX internals from a filesystem perspective.

Much has been written about the different UNIX and UNIX-like kernels since Maurice Bach's book *The Design of the UNIX Operating System* [BACH86] first appeared in 1986. At that time, he documented the internals of System V Release 2 (SVR2). However, much had already happened in the UNIX world when SVR2 appeared. The earliest documented kernel was 6th Edition as described in John Lions' work *Lions' Commentary on UNIX 6th Edition—with Source Code* [LION96], which was an underground work until its publication in 1996. In addition to these two books, there have also been a number of others that have described the different UNIX kernel versions.

When writing about operating system internals, there are many different topics to cover from process management to virtual memory management, from device drivers to networking, and hardware management to filesystems. One could fill a book on each of these areas and, in the case of networking and device drivers, specialized books have in fact appeared over the last decade.

Filesystems are a subject of great interest to many although they have typically been poorly documented. This is where this book comes into play.

This book covers the history of UNIX describing how filesystems were implemented in the early research editions of UNIX up to today's highly scalable enterprise class UNIX systems. All of the major changes in the history of UNIX

that pertain to filesystems are covered along with a view of how some of the more well known filesystems are implemented.

Not forgetting the user interface to filesystems, the book also presents the file and filesystem-level system call and library-level APIs that programmers expect to see. By providing this context it is easier to understand the services that filesystems are expected to provide and therefore why they are implemented the way they are.

Wherever possible, this book provides practical examples, either through programmatic means or through analysis. To provide a more practical edge to the material presented, the book provides a complete implementation of a filesystem on Linux together with instructions on how to build the kernel and filesystem, how to install it, and analyze it using appropriate kernel-level debuggers. Examples are then given for readers to experiment further.

Who Should Read This Book?

Rather than reach for the usual group of suspects—kernel engineers and operating system hobbyists—this book is written in such a way that anyone who has an interest in filesystem technology, regardless of whether they understand operating system internals or not, can read the book to gain an understanding of file and filesystem principles, operating system internals, and filesystem implementations.

This book should appeal to anyone interested in UNIX, its history, and the standards that UNIX adheres to. Anyone involved in the storage industry should also benefit from the material presented here.

Because the book has a practical edge, the material should be applicable for undergraduate degree-level computer science courses. As well as a number of examples throughout the text, which are applicable to nearly all versions of UNIX, the chapter covering Linux filesystems provides a number of areas where students can experiment.