

UNIVERSITY OF CALIFORNIA

Los Angeles

# Stackable Design of File Systems

A dissertation submitted in partial satisfaction  
of the requirements for the degree  
Doctor of Philosophy in Computer Science

by

**John Shelby Heidemann**

1995

© Copyright by  
John Shelby Heidemann  
1995

The dissertation of John Shelby Heidemann is approved.

Kirby A. Baker

Rajive L. Bagrodia

Richard Muntz

D. Stott Parker, Committee Co-chair

Gerald J. Popek, Committee Co-chair

University of California, Los Angeles

1995

2

*Are you dedicated?*

# TABLE OF CONTENTS

<b>1</b>	<b>Introduction . . . . .</b>	<b>1</b>
----------	-------------------------------	----------

# LIST OF FIGURES

# LIST OF TABLES

## ACKNOWLEDGMENTS

Ack! P'tui.

## VITA

1968            Born, Syracuse, New York, USA.

etc.            etc.

1995            Graduated, UCLA

## PUBLICATIONS

Richard G. Guy, John S. Heidemann, Wai Mak, Thomas W. Page, Jr., Gerald J. Popek, and Dieter Rothmeier. Implementation of the Ficus replicated file system. In *USENIX Conference Proceedings*, pages 63–71. USENIX, June 1990.

etc.

# CHAPTER 1

## Introduction

Filing services are one of the most user-visible parts of the operating system, so it is not surprising that many new services are proposed by researchers and that a variety of third parties are interested in providing these solutions. Of the many innovations which have been proposed, very few have become widely available in a timely fashion. We believe this delay results from two deficiencies in practices of current file-system development. First, file systems are large and difficult to implement. This problem is compounded because no good mechanism exists to allow new services to build on those which already exist. Second, file systems today are built around a few fixed interfaces which fail to accommodate the change and evolution inherent in operating systems development. Today's filing interfaces vary from system to system, and even between point releases of a single operating system. These differences greatly complicate and therefore discourage third-party development and adoption of filing extensions.

These problems raise barriers to the widespread development, deployment, and maintenance of new filing services. The thesis of this dissertation is that a layered, *stackable* structure with an *extensible* interface provides a much better methodology for file-system development. We propose construction of filing services from a number of potentially independently developed modules. By stackable, we mean that these modules are bounded by identical, or *symmetric*, interfaces above and below. By extensible, we mean that these interfaces can be independently changed by multiple parties, without invalidating existing or future work.

To validate this thesis we developed a framework supporting stackable file-systems and used that framework to construct several different filing services. This dissertation describes

the design, implementation, and evaluation of this system.