



OPERATING SYSTEMS REVIEW

A Publication of the
Association for Computing Machinery
Special Interest Group on Operating Systems

OSR Special Issue — Winter 1998

USENIX Association

**Proceedings of the
Third Symposium on Operating Systems
Design and Implementation
(OSDI '99)**

Co-sponsored by IEEE TCOS and ACM SIGOPS

**February 22-25, 1999
New Orleans, Louisiana**

Contents

Third Symposium on Operating Systems Design and Implementation

February 22-25, 1999 New Orleans, Louisiana

Index of Authors	vii
Message from the Program Chairs	ix

Tuesday, February 23

I/O

Session Chair: Sean O'Malley, Network Appliance

Automatic I/O Hint Generation Through Speculative Execution	1
<i>Fay Chang, Garth A. Gibson, Carnegie Mellon University</i>	

IO-Lite: A Unified I/O Buffering and Caching System	15
<i>Vivek S. Pai, Peter Druschel, Willy Zwaenepoel, Rice University</i>	

Virtual Log Based File Systems for a Programmable Disk	29
<i>Randolph Y. Wang, University of California, Berkeley; Thomas E. Anderson, University of Washington, Seattle; David A. Patterson, University of California, Berkeley</i>	

Resource Management

Session Chair: Greg Minshall, Siara Systems

Resource Containers: A New Facility for Resource Management in Server Systems	45
<i>Gaurav Banga, Peter Druschel, Rice University; Jeffrey C. Mogul, Western Research Laboratory, Compaq Computer Corp.</i>	

Defending Against Denial of Service Attacks in Scout	59
<i>Oliver Spatscheck, University of Arizona; Larry L. Peterson, Princeton University</i>	

Self-Paging in the Nemesis Operating System	73
<i>Steven M. Hand, University of Cambridge Computer Laboratory</i>	

Wednesday, February 24

Kernels

Session Chair: Rob Pike, Lucent Technologies

Tornado: Maximizing Locality and Concurrency in a Shared Memory Multiprocessor Operating System	87
<i>Ben Gamsa, University of Toronto; Orran Krieger, IBM T.J. Watson Research Center; Jonathan Appavoo, Michael Stumm, University of Toronto</i>	

Interface and Execution Models in the Fluke Kernel	101
<i>Bryan Ford, Mike Hibler, Jay Lepreau, Roland McGrath, Patrick Tullmann, University of Utah</i>	

Fine-Grained Dynamic Instrumentation of Commodity Operating System Kernels	117
<i>Ariel Tamches, Barton P. Miller, University of Wisconsin, Madison</i>	

Real-Time

Session Chair: Mike Jones, Microsoft Corporation

ETI Resource Distributor: Guaranteed Resource Allocation and Scheduling in Multimedia Systems131
Miche Baker-Harvey, Equator Technologies, Inc.

A Feedback-driven Proportion Allocator for Real-Rate Scheduling145
David C. Steere, Ashvin Goel, Joshua Gruenberg, Dylan McNamee, Calton Pu, Jonathan Walpole, Oregon Graduate Institute

A Comparison of Windows Driver Model Latency Performance on Windows NT and Windows 98159
Erik Cota-Robles, James P. Held, Intel Architecture Labs

Distributed Systems

Session Chair: Tom Anderson, University of Washington

Practical Byzantine Fault Tolerance173
Miguel Castro, Barbara Liskov, Massachusetts Institute of Technology

The Coign Automatic Distributed Partitioning System187
Galen C. Hunt, Microsoft Research; Michael L. Scott, University of Rochester

Thursday, February 25

Virtual Memory

Session Chair: Kai Li, Princeton University

Tapeworm: High-Level Abstractions of Shared Accesses201
Peter J. Keleher, University of Maryland

MultiView and Millipage—Fine-Grain Sharing in Page-Based DSMs215
Ayal Itzkovitz, Assaf Schuster, Technion—Israel Institute of Technology

Optimizing the Idle Task and Other MMU Tricks229
Cort Dougan, Paul Mackerras, Victor Yodaiken, New Mexico Institute of Technology

Filesystems

Session Chair: Bruce Lindsay, IBM Almaden Research Center

Logical vs. Physical File System Backup239
Norman C. Hutchinson, University of British Columbia; Stephen Manley, Mike Federwisch, Guy Harris, Dave Hitz, Steven Kleiman, Sean O'Malley, Network Appliance, Inc.

The Design of a Multicast-based Distributed File System251
Björn Grönvall, Assar Westerlund, Stephen Pink, Swedish Institute of Computer Science and Luleå University of Technology

Integrating Content-based Access Mechanisms with Hierarchical File Systems265
Burra Gopal, Microsoft Corp; Udi Manber, University of Arizona