



OPERATING SYSTEMS REVIEW

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

Volume 39, Number 5

December 2005

SOSP'05

Proceedings of the
**20th ACM Symposium
on Operating
Systems Principles**

Brighton, United Kingdom
October 23-26, 2005

Sponsored by:



SIGOPS

ACM SIG on Operating Systems

SOSP'05



Proceedings of the 20th ACM Symposium on Operating Systems Principles

October 23-26, 2005 • Brighton, United Kingdom



Sponsored by:



SIGOPS

ACM SIG on Operating Systems

Table of Contents

20 th Symposium on Operating Systems Principles Organization	vii
Additional Reviewers	viii
Sponsor & Supporters	ix

Session 1: Integrity and Isolation

- **Pioneer: Verifying Code Integrity and Enforcing Untampered Code Execution on Legacy Systems** 1
A. Seshadri, M. Luk, E. Shi, A. Perrig (*Carnegie Mellon University*), L. van Doorn (*IBM*),
P. Khosla (*Carnegie Mellon University*)
- **Labels and Event Processes in the Asbestos Operating System** 17
P. Efstathopoulos (*University of California at Los Angeles*),
M. Krohn (*Massachusetts Institute of Technology*),
S. VanDeBogart (*University of California at Los Angeles*),
C. Frey, D. Ziegler (*Massachusetts Institute of Technology*),
E. Kohler (*University of California at Los Angeles*),
D. Mazières (*New York University at Stanford*),
F. Kaashoek, R. Morris (*Massachusetts Institute of Technology*)
- **Mondrix: Memory Isolation for Linux using Modriaan Memory Protection** 31
E. Witchel (*University of Texas at Austin*), J. Rhee (*Purdue University*),
K. Asanović (*Massachusetts Institute of Technology*)

Session 2: Distributed Systems

- **BAR Fault Tolerance for Cooperative Services** 45
A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, C. Porth (*University of Texas at Austin*)
- **Fault-Scalable Byzantine Fault-Tolerant Services** 59
M. Abd-El-Malek, G. R. Ganger (*Carnegie Mellon University*), G. R. Goodson (*Network Appliance, Inc.*),
M. K. Reiter, J. J. Wylie (*Carnegie Mellon University*)
- **Implementing Declarative Overlays** 75
B. T. Loo, T. Condie (*University of California at Berkeley*),
J. M. Hellerstein (*Intel Research Berkeley and University of California at Berkeley*),
P. Maniatis, T. Roscoe (*Intel Research Berkeley*), I. Stoica (*University of California at Berkeley*)

Session 3: History and Context

- **Detecting Past and Present Intrusions through Vulnerability-Specific Predicates** 91
A. Joshi, S. T. King, G. W. Dunlap, P. M. Chen (*University of Michigan*)
- **Capturing, Indexing, Clustering, and Retrieving System History** 105
I. Cohen (*Hewlett-Packard Laboratories*), S. Zhang (*Stanford University*),
M. Goldszmidt, J. Symons, T. Kelly (*Hewlett-Packard Laboratories*), A. Fox (*Stanford University*)
- **Connections: Using Context to Enhance File Search** 119
C. A. N. Soules, G. R. Ganger (*Carnegie Mellon University*)

Session 4: Containment

- **Vigilante: End-to-End Containment of Internet Worms** 133
M. Costa (*University of Cambridge and Microsoft Research*), J. Crowcroft (*University of Cambridge*),
M. Castro, A. Rowstron, L. Zhou, L. Zhang, P. Barham (*Microsoft Research*)

- **Scalability, Fidelity, and Containment in the Potemkin Virtual Honeyfarm**..... 148
M. Vrable, J. Ma, J. Chen, D. Moore, E. Vandekieft, A. C. Snoeren, G. M. Voelker,
S. Savage (*University of California at San Diego*)
- **The Taser Intrusion Recovery System**..... 163
A. Goel, K. Po, K. Farhadi, Z. Li, E. de Lara (*University of Toronto*)

Session 5: Filesystems

- **Hibernator: Helping Disk Arrays Sleep through the Winter** 177
Q. Zhu, Z. Chen, L. Tan, Y. Zhou (*University of Illinois at Urbana-Champaign*),
K. Keeton, J. Wilkes (*Hewlett-Packard Laboratories*)
- **Speculative Execution in a Distributed File System**..... 191
E. B. Nightingale, P. M. Chen, J. Flinn (*University of Michigan*)
- **IRON File Systems**..... 206
V. Prabhakaran, L. N. Bairavasundaram, N. Agrawal, H. S. Gunawi, A. C. Arpaci-Dusseau,
R. H. Arpaci-Dusseau (*University of Wisconsin at Madison*)

Session 6: Bugs

- **RaceTrack: Efficient Detection of Data Race Conditions via Adaptive Tracking** 221
Y. Yu, T. Rodeheffer (*Microsoft Research*), W. Chen (*University of California at Berkeley*)
- **Rx: Treating Bugs As Allergies—A Safe Method to Survive Software Failures**..... 235
F. Qin, J. Tucek, J. Sundaresan, Y. Zhou (*University of Illinois at Urbana Champaign*)

Session 7: Optimization

- **Idletime Scheduling with Preemption Intervals**..... 249
L. Eggert (*NEC Europe Ltd.*), J. D. Touch (*University of Southern California*)
- **FS2: Dynamic Data Replication in Free Disk Space for Improving Disk Performance and Energy Consumption**..... 263
H. Huang, W. Hung, K. G. Shin (*The University of Michigan*)
- **THINC: A Virtual Display Architecture for Thin-Client Computing** 277
R. A. Baratto, L. N. Kim, J. Nieh (*Columbia University*)

- **Author Index**..... 291