

7th USENIX Conference on File and Storage Technologies

February 24–27, 2009

San Francisco, CA, USA

Message from the Program Co-Chairs	v
Index of Authors	vi

Wednesday, February 25

Augmenting File System Functionality

The Case of the Fake Picasso: Preventing History Forgery with Secure Provenance	1
<i>Ragib Hasan, University of Illinois at Urbana-Champaign; Radu Sion, Stony Brook University; Marianne Winslett, University of Illinois at Urbana-Champaign</i>	
Causality-Based Versioning	15
<i>Kiran-Kumar Muniswamy-Reddy and David A. Holland, Harvard University</i>	
Enabling Transactional File Access via Lightweight Kernel Extensions	29
<i>Richard P. Spillane, Sachin Gaikwad, Manjunath Chinni, and Erez Zadok, Stony Brook University; Charles P. Wright, IBM T.J. Watson Research Center</i>	

Diagnosis

Understanding Customer Problem Troubleshooting from Storage System Logs	43
<i>Weihsiang Jiang and Chongfeng Hu, University of Illinois at Urbana-Champaign; Shankar Pasupathy and Arkady Kanevsky, NetApp, Inc.; Zhenmin Li, Pattern Insight, Inc.; Yuanyuan Zhou, University of Illinois at Urbana-Champaign</i>	
DIADS: Addressing the “My-Problem-or-Yours” Syndrome with Integrated SAN and Database Diagnosis	57
<i>Shivnath Babu and Nedyalko Borisov, Duke University; Sandeep Uttamchandani, Ramani Routray, and Aameek Singh, IBM Almaden Research Center</i>	

Thursday, February 26

Scheduling

Dynamic Resource Allocation for Database Servers Running on Virtual Storage	71
<i>Gokul Soundararajan, Daniel Lupei, Saeed Ghanbari, Adrian Daniel Popescu, Jin Chen, and Cristiana Amza, University of Toronto</i>	
PARDA: Proportional Allocation of Resources for Distributed Storage Access	85
<i>Ajay Gulati, Irfan Ahmad, and Carl A. Waldspurger, VMware Inc.</i>	
CA-NFS: A Congestion-Aware Network File System	99
<i>Alexandros Batsakis, NetApp and Johns Hopkins University; Randal Burns, Johns Hopkins University; Arkady Kanevsky, James Lentini, and Thomas Talpey, NetApp</i>	

Tools You Wish You Had

Sparse Indexing: Large Scale, Inline Deduplication Using Sampling and Locality	111
<i>Mark Lillibridge and Kave Eshghi, HP Labs; Deepavali Bhagwat, University of California, Santa Cruz; Vinay Deolalikar, HP Labs; Greg Trezise and Peter Camble, HP Storage Works Division</i>	
Generating Realistic Impressions for File-System Benchmarking	125
<i>Nitin Agrawal, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, University of Wisconsin, Madison</i>	
Capture, Conversion, and Analysis of an Intense NFS Workload	139
<i>Eric Anderson, HP Labs</i>	

Thursday, February 26 (continued)

Metadata and Optimization

- Spyglass: Fast, Scalable Metadata Search for Large-Scale Storage Systems 153
Andrew W. Leung, University of California, Santa Cruz; Minglong Shao, Timothy Bisson, and Shankar Pasupathy, NetApp; Ethan L. Miller, University of California, Santa Cruz
- Perspective: Semantic Data Management for the Home 167
Brandon Salmon, Carnegie Mellon University; Steven W. Schlosser, Intel Research Pittsburgh; Lorrie Faith Cranor and Gregory R. Ganger, Carnegie Mellon University
- BORG: Block-reORGanization for Self-optimizing Storage Systems. 183
Medha Bhadkamkar, Jorge Guerra, and Luis Useche, Florida International University; Sam Burnett, Carnegie Mellon University; Jason Liptak, Syracuse University; Raju Rangaswami and Vagelis Hristidis, Florida International University

Distributed Storage

- HYDRAsTOR: A Scalable Secondary Storage. 197
Cezary Dubnicki, Leszek Gryz, Lukasz Heldt, Michal Kaczmarczyk, Wojciech Kilian, Przemyslaw Strzelczak, and Jerzy Szczepkowski, 9LivesData, LLC; Cristian Ungureanu, NEC Laboratories America; Michal Welnicki, 9LivesData, LLC
- Smoke and Mirrors: Reflecting Files at a Geographically Remote Location Without Loss of Performance 211
Hakim Weatherspoon, Lakshmi Ganesh, and Tudor Marian, Cornell University; Mahesh Balakrishnan, Microsoft Research, Silicon Valley; Ken Birman, Cornell University
- Cumulus: Filesystem Backup to the Cloud 225
Michael Vrable, Stefan Savage, and Geoffrey M. Voelker, University of California, San Diego

Friday, February 27

Data Integrity

- WorkOut: I/O Workload Outsourcing for Boosting RAID Reconstruction Performance 239
Suzhen Wu, Huazhong University of Science and Technology; Hong Jiang, University of Nebraska–Lincoln; Dan Feng, Huazhong University of Science and Technology; Lei Tian, Huazhong University of Science and Technology and University of Nebraska–Lincoln; Bo Mao, Huazhong University of Science and Technology
- A Performance Evaluation and Examination of Open-Source Erasure Coding Libraries for Storage 253
James S. Plank, University of Tennessee; Jianqiang Luo, Wayne State University; Catherine D. Schuman, University of Tennessee; Lihao Xu, Wayne State University; Zooko Wilcox-O’Hearn, AllMyData, Inc.
- Tiered Fault Tolerance for Long-Term Integrity 267
Byung-Gon Chun and Petros Maniatis, Intel Research Berkeley; Scott Shenker and John Kubiatowicz, University of California, Berkeley

Controllers and Caching

- A Systematic Approach to System State Restoration during Storage Controller Micro-Recovery 283
Sangeetha Seshadri, Georgia Institute of Technology; Lawrence Chiu, IBM Almaden Research Center; Ling Liu, Georgia Institute of Technology
- CLIC: CLient-Informed Caching for Storage Servers 297
Xin Liu, Ashraf Aboulnaga, Kenneth Salem, and Xuhui Li, University of Waterloo
- Minuet: Rethinking Concurrency Control in Storage Area Networks 311
Andrey Ermolinskiy and Daekyeong Moon, University of California, Berkeley; Byung-Gon Chun, Intel Research, Berkeley; Scott Shenker, University of California, Berkeley, and ICSI