

FAST '09

7th USENIX Conference on File and Storage Technologies

FEBRUARY 24–27, 2009 | SAN FRANCISCO, CALIFORNIA

USENIX

Sponsored by USENIX in cooperation with ACM SIGOPS, IEEE Mass Storage Systems Technical Committee (MSSTC), and IEEE TCOS

The 7th USENIX Conference on File and Storage Technologies (FAST '09) brings together storage system researchers and practitioners to explore new directions in the design, implementation, evaluation, and deployment of storage systems.

Back again for '09, the FAST program is offering tutorials. Taking place on Tuesday, February 24, the four half-day tutorials give you the opportunity to learn from leaders in the storage industry. Take advantage of the special FAST offer: Buy one half-day tutorial and get the second one for free.

This year's innovative technical program includes 23 technical papers, as well as a keynote address, Work-in-Progress Reports (WiPs), and a poster session. See the full program on the reverse side of this page.

Don't miss this opportunity to meet with premier storage system researchers and practitioners for three and one-half days of ground-breaking file and storage information and training. Register by Friday, February 9, 2009, at www.usenix.org/fast2009 and save up to \$200!

Make your hotel reservation early!

The Grand Hyatt San Francisco • Phone: (415) 398-1234 • <http://grandsanfrancisco.hyatt.com/>
Mention USENIX or FAST to get our special group rate.

Join us for the premier end-user storage event.

Co-located with the 1st Workshop on the Theory and Practice of Provenance (TaPP '09),
taking place on February 23, 2009

See <http://www.usenix.org/tapp09> for more information

Tutorial Program

Tuesday, February 24, 2009

Half Day Tutorials (a.m.)

T1 Clustered and Parallel Storage System Technologies **UPDATED!**

Brent Welch and Marc Unangst, Panasas
Cluster-based parallel storage technologies are now capable of delivering performance scaling from 10s to 100s of GB/sec. This tutorial will examine current state-of-the-art high-performance file systems and the underlying technologies employed to deliver scalable performance across a range of scientific and industrial applications.

T2 Security and Usability: What Do We Know? **NEW!**

Simson Garfinkel, Naval Postgraduate School
For years we've heard that security and usability are antagonistic: secure systems aren't usable and usable systems aren't secure. New research in the field of HCI-SEC reveals this myth for what it is. In this tutorial we will review the past few years of research in security and usability and see how to create systems that are both usable and secure. We'll also discuss how to evaluate the usability of a system in the lab, in the field, and with the necessary legal approvals.

Half Day Tutorials (p.m.)

T3 Storage Class Memory, Technology, and Uses **UPDATED!**

Richard Freitas, Winfried Wilcke, Bülent Kurdi, and Geoffrey Burr, IBM Almaden Research Center

This tutorial will give a detailed overview of Storage Class Memory (SCM) device technologies, how they will impact the design of storage controllers and storage systems, and how SCM used as main memory will affect system architecture and software.

T4 Web-Scale Data Management **NEW!** **Christopher Olston and Benjamin Reed, Yahoo! Research**

In this tutorial we survey Web-scale data management technologies, with special focus on open-source instances. We give concrete code examples modeled after real-world use cases at companies like Yahoo!. These technologies have not yet reached maturity; at the end of the tutorial, we discuss some "in-the-works" and "wish-list" features in this space.

Thanks to our Sponsors

Sun
Microsystems

SNIA
VMware

NetApp

HP Labs

EMC²

Microsoft
Research

Media Sponsors

ACM Queue
BetaNews
InfoSec News
Linux Gazette
Linux Journal

Linux Pro
Magazine
LXer.com
SNIA
Toolbox.com

Conference Organizers

Program Co-Chairs

Margo Seltzer, *Harvard University*
Ric Wheeler, *Red Hat*

Program Committee

Sameer Ajmani, *Google*
Remzi Arpaci-Dusseau, *University of Wisconsin, Madison*
David L. Black, *EMC*
Bill Bolosky, *Microsoft Research, Redmond*
James Bottomley, *Hansen Partnership*
David DeWitt, *Microsoft, Madison*
Daniel Ellard, *BBN Technologies*
Greg Ganger, *Carnegie Mellon University*
Valerie Henson, *Red Hat*
Ethan Miller, *University of California, Santa Cruz*
Alina Oprea, *RSA Security/EMC*
James S. Plank, *University of Tennessee*
Calton Pu, *Georgia Institute of Technology*
Raju Rangaswami, *Florida International University*
Narasimha Reddy, *Texas A&M University*
Ohad Rodeh, *IBM Research, Haifa*
Ken Salem, *University of Waterloo*
Jiri Schindler, *NetApp*
Bianca Schroeder, *University of Toronto*
Liuba Shrira, *Brandeis University*
Niraj Tolia, *HP Labs*
Hakim Weatherspoon, *Cornell University*

Work-in-Progress Reports (WiPs) and Poster Session Chair

Geoff Kuenning, *Harvey Mudd College*

Register by Monday, February 9, 2009, and save!

www.usenix.org/fast2009

Wednesday, February 25

9:00 a.m.–10:30 a.m.

Wednesday

OPENING REMARKS AND BEST PAPER AWARDS

Program Co-Chairs: Margo Seltzer, *Harvard University*; Ric Wheeler, *Red Hat*

KEYNOTE ADDRESS: TBA

11:00 a.m.–12:30 p.m.

Wednesday

AUGMENTING FILE SYSTEM FUNCTIONALITY

The Case of the Fake Picasso: Preventing History Forgery with Secure Provenance

Ragib Hasan, *University of Illinois at Urbana-Champaign*; Radu Sion, *Stony Brook University*; Marianne Winslett, *University of Illinois at Urbana-Champaign*

Causality-Based Versioning

Kiran-Kumar Muniswamy-Reddy and David A. Holland, *Harvard University*

Enabling Transactional File Access via Lightweight Kernel Extensions

Richard P. Spillane, Sachin Gaikwad, and Erez Zadok, *Stony Brook University*; Charles P. Wright, *IBM T.J. Watson Research Center*; Manju Chinni, *Stony Brook University*

12:30 p.m.–2:00 p.m.

Lunch (on your own)

2:00 p.m.–3:00 p.m.

Wednesday

DIAGNOSIS

Understanding Customer Problem Troubleshooting from Storage System Logs

Weihang Jiang and Chongfeng Hu, *University of Illinois at Urbana-Champaign*; Shankar Pasupathy and Arkady Kanevsky, *NetApp*; Zhenmin Li, *PatternInsight*; Yuanyuan Zhou, *University of Illinois at Urbana-Champaign*

DiaDS: Addressing the “My-Problem-or-Yours” Syndrome with Integrated SAN and Database Diagnosis

Shivnath Babu and Nedyalko Borisov, *Duke University*; Sandeep Uttamchandani, Ramani Routray, and Aameek Singh, *IBM Almaden Research Center*

3:30 p.m.–5:00 p.m.

Wednesday

WORK-IN-PROGRESS REPORTS (WIPS)

The FAST technical sessions will include a session for Work-in-Progress reports, preliminary results, and “outrageous” opinion statements. We are particularly interested in presentations of student work. Submit a one-page abstract in PDF format to fast09wips@usenix.org by 11:59 p.m. PST, January 29, 2009.

5:00 p.m.–7:00 p.m.

Wednesday

POSTER SESSION & HAPPY HOUR

Held in conjunction with a happy hour, the poster and demo session will allow researchers to present recent and ongoing projects. See <http://www.usenix.org/fast09/posters> for guidelines. Proposals are due by 11:59 p.m. PST, January 29, 2009.

Thursday, February 26

9:00 a.m.–10:30 a.m.

Thursday

SCHEDULING

Dynamic Resource Allocation for Databases Running on Virtual Storage

Gokul Soundararajan, Daniel Lupei, Saeed Ghanbari, Adrian Popescu, Jin Chen, and Cristiana Amza, *University of Toronto*

PARDA: Proportional Allocation of Resources for Distributed Storage

Ajay Gulati, Irfan Ahmad, and Carl Waldspurger, *VMware Inc.*

A Congestion-Aware Network File System

Alexandros Batsakis, and Randal Burns, *Johns Hopkins University*; Arkady Kanevsky, James Lentini, and Thomas Talpey, *NetApp*

11:00 a.m.–12:30 p.m.

Thursday

TOOLS YOU WISH YOU HAD

Sparse Indexing: Large Scale, Inline Deduplication Using Sampling and Locality

Mark Lillibridge, Kave Eshghi, Deepavali Bhagwat, and Vinay Deolalikar, *HP Labs*; Greg Trezise and Peter Cambell, *HP StorageWorks Division*

Generating Realistic Impressions for File-System Benchmarking

Nitin Agrawal, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin, Madison*

Capture, Conversion, Storage, and Analysis of an Intense NFS Workload

Eric Anderson, *HP Labs*

12:30 p.m.–2:00 p.m.

Conference Luncheon

2:00 p.m.–3:30 p.m.

Thursday

METADATA AND OPTIMIZATION

Spynlass: Fast, Scalable Metadata Search for Large-Scale Storage Systems

Andrew W. Leung, *University of California, Santa Cruz*; Minglong Shao, Tim Bisson, and Shankar Pasupathy, *NetApp*; Ethan L. Miller, *University of California, Santa Cruz*

Perspective: Semantic Data Management for the Home

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

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*

4:00 p.m.–5:30 p.m.

Thursday

DISTRIBUTED STORAGE

Hydrastor: A Scalable Secondary Storage

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*

Mirrors: Reflecting Files at a Geographically Remote Location Without Loss of Performance

Hakim Weatherspoon, Lakshmi Ganesh, Tudor Marian, Mahesh Balakrishnan, and Ken Birman, *Cornell University*

Cumulus: Filesystem Backup to the Cloud

Michael Vrable, Stefan Savage, and Geoffrey M. Voelker, *University of California, San Diego*

6:00 p.m.–7:30 p.m.

Thursday

CONFERENCE RECEPTION *Sponsored by NetApp*

Friday, February 27

8:30 a.m.–10:00 a.m.

Friday

DATA INTEGRITY

WorkOut: I/O Workload Outsourcing for Boosting RAID Reconstruction Performance

Suzhen Wu, *Huazhong University of Science and Technology*; Hong Jiang, *University of Nebraska, Lincoln*; Dan Feng, Lei Tian, and Bo Mao, *Huazhong University of Science and Technology*

A Performance Evaluation and Examination of Open-Source Erasure Coding Libraries for Storage

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

Byung-Gon Chun and Petros Maniatis, *Intel Research Berkeley*; Scott Shenker and John Kubiatowicz, *University of California, Berkeley*

10:30 a.m.–noon

Friday

CONTROLLERS AND CACHING

A Systematic Approach to System State Restoration During Storage Controller Micro-Recovery

Sangeetha Seshadri, *Georgia Tech*; Lawrence Chiu, *IBM Almaden Research Center*; Ling Liu, *Georgia Tech*

CLIC: CLient-Informed Caching for Storage Servers

Xin Liu, Ashraf Aboulnaga, Kenneth Salem, and Xuhui Li, *University of Waterloo*

Minuet: Rethinking Concurrency Control in Storage Area Networks

Andrey Ermolinskiy and Daekyeong Moon, *University of California, Berkeley*; Byung-Gon Chun, *Intel Research, Berkeley*; Scott Shenker, *University of California, Berkeley*, *International Computer Science Institute*

Early Bird Registration Deadline: Monday, February 9, 2009 • Register Online: <http://www.usenix.org/fast2009>