

FAST '10

8th USENIX Conference on File and Storage Technologies

USENIX

FEBRUARY 23–26, 2010 | SAN JOSE, CALIFORNIA
Sponsored by USENIX in cooperation with ACM SIGOPS

The 8th USENIX Conference on File and Storage Technologies (FAST '10) brings together storage system researchers and practitioners to explore new directions in the design, implementation, evaluation, and deployment of storage systems.

Back again for '10, the FAST program is offering tutorials. Taking place on Tuesday, February 23, 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 21 technical papers, as well as two keynote addresses, 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 Monday, February 8, 2010, at www.usenix.org/fast2010 and save up to \$200!

Make your hotel reservation early!

The Fairmont San Jose • 170 S. Market Street • San Jose, CA 95113 • Phone: (408) 998-1900
Call and mention USENIX or FAST or book online via <http://www.usenix.org/fast2010/hotel>

Join us for the premier end-user storage event.

Thanks to Our Sponsors

NetApp
Microsoft Research
HP Labs
IBM

Thanks to Our Media Sponsors

acmqueue
govirtual.org
HPCwire
Linux Journal
LXer.com
Mashable
Mission Critical
Opensourc3
UserFriendly.org

Tutorial Program

Tuesday, February 23, 2009

Half Day Tutorials (a.m.)

T1 Storage Virtualization and the Utility Model **NEW!**

Tom Clark, *Brocade Communications*

Storage virtualization is an integral part of new datacenter strategies aimed at maximizing utilization of assets, automating data processes, and reducing costs. This tutorial covers the underlying mechanics of storage virtualization, the technical implementations in host, appliance, fabric, and array-based solutions, and the applications that can more fully exploit virtualization to streamline storage operations.

T2 Solid-State Storage: Technology, Design, and Application **UPDATED!**

Richard Freitas and Larry Chiu, *IBM Almaden Research Center*

This tutorial will briefly examine the leading solid-state memory technologies and then focus on the impact the introduction of such technologies will have on storage systems. It will include a discussion of SSD design, storage system architecture, applications, and performance assessment.

Half Day Tutorials (p.m.)

T3 Storage and Network Deduplication Technologies **NEW!**

Michael Conduct, *NetApp*

Economic and environmental concerns are currently motivating a push across the computing industry to do more with less: less energy and less money. Deduplication of data is one of the most effective tools to accomplish this. This tutorial will provide a detailed look at the multitude of ways deduplication can be used to improve the efficiency of storage and networking devices.

T4 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 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.

Conference Organizers

Program Co-Chairs

Randal Burns, *Johns Hopkins University*
Kimberly Keeton, *Hewlett-Packard Labs*

Program Committee

Patrick Eaton, *EMC*
Jason Flinn, *University of Michigan*
Gary Grider, *Los Alamos National Lab*
Ajay Gulati, *VMware*
Sudhanva Gurumurthi, *University of Virginia*
Dushyanth Narayanan, *Microsoft Research*
Jason Nieh, *Columbia University*
Christopher Olston, *Yahoo! Research*
Hugo Patterson, *Data Domain*
Beth Plale, *Indiana University*
James Plank, *University of Tennessee*
Erik Riedel, *EMC*
Alma Riska, *Seagate*
Steve Schlosser, *Avere Systems*
Bianca Schroeder, *University of Toronto*
Karsten Schwan, *Georgia Institute of Technology*
Craig Soules, *Hewlett-Packard Labs*
Alan Sussman, *University of Maryland*
Kaladhar Voruganti, *NetApp*
Hakim Weatherspoon, *Cornell University*
Brent Welch, *Panasas*
Ric Wheeler, *Red Hat*
Yuanyuan Zhou, *University of California, San Diego*

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

Hakim Weatherspoon, *Cornell University*

Tutorial Chair

David Pease, *IBM Almaden Research Center*

Steering Committee

Andrea C. Arpaci-Dusseau, *University of Wisconsin—Madison*
Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*
Mary Baker, *Hewlett-Packard Labs*
Greg Ganger, *Carnegie Mellon University*
Garth Gibson, *Carnegie Mellon University and Panasas*
Peter Honeyman, *CITI, University of Michigan, Ann Arbor*
Darrell Long, *University of California, Santa Cruz*
Jai Menon, *IBM Research*
Erik Riedel, *EMC*
Margo Seltzer, *Harvard School of Engineering and Applied Sciences*
Chandu Thekkath, *Microsoft Research*
Ric Wheeler, *Red Hat*
John Wilkes, *Google*
Ellie Young, *USENIX Association*

Register by Monday, February 8, 2010, and save!

www.usenix.org/fast2010

Wednesday, February 24, 2010

9:00 a.m.–10:30 a.m.

Wednesday

OPENING REMARKS AND BEST PAPER AWARDS

Program Co-Chairs: Randal Burns, *Johns Hopkins University*; Kimberly Keeton, *Hewlett-Packard Labs*

KEYNOTE ADDRESS: Technology for Developing Regions

Eric Brewer, *University of California, Berkeley*

11:00 a.m.–12:30 p.m.

Wednesday

BUILD A BETTER FILE SYSTEM AND THE WORLD WILL BEAT A PATH TO YOUR DOOR.

quFiles: The Right File at the Right Time

Kaushik Veeraraghavan and Jason Flinn, *University of Michigan*; Edmund B. Nightingale, *Microsoft Research, Redmond*; Brian Noble, *University of Michigan*

Tracking Back References in a Write-Anywhere File System

Peter Macko and Margo Seltzer, *Harvard University*; Keith A. Smith, *NetApp, Inc.*

End-to-end Data Integrity for File Systems: A ZFS Case Study

Yupu Zhang, Abhishek Rajimwale, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*

12:30 p.m.–2:00 p.m.

Conference Luncheon

2:00 p.m.–3:30 p.m.

Wednesday

LOOKING FOR TROUBLE

Black-Box Problem Diagnosis in Parallel File Systems

Michael P. Kasick, *Carnegie Mellon University*; Jiaqi Tan, *DSO National Labs, Singapore*; Rajeev Gandhi and Priya Narasimhan, *Carnegie Mellon University*

A Clean-Slate Look at Disk Scrubbing

Alina Oprea and Ari Juels, *RSA Laboratories*

Understanding Latent Sector Errors and How to Protect Against Them

Bianca Schroeder, Sotirios Damouras, and Phillipa Gill, *University of Toronto*

4:00 p.m.–5:30 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. See <http://www.usenix.org/fast2010/wips> for more info. Proposals are due by 11:59 p.m. PST on January 28, 2010.

6:00 p.m.–7:30 p.m.

Wednesday

POSTER SESSION & HAPPY HOUR

Held in conjunction with a happy hour, the poster session will allow researchers to present recent and ongoing projects. See <http://www.usenix.org/fast2010/posters> for more info. Proposals are due by 11:59 p.m. PST on January 28, 2010.

Thursday, February 25, 2010

9:00 a.m.–10:30 a.m.

Thursday

KEYNOTE ADDRESS: Enterprise Analytics on Demand

Oliver Ratzesberger, *eBay, Inc.*

11:00 a.m.–12:30 p.m.

Thursday

FLASH: SAVIOR OF THE UNIVERSE?

DFS: A File System for Virtualized Flash Storage

William K. Josephson and Lars A. Bongo, *Princeton University*; David Flynn, *Fusion-io*; Kai Li, *Princeton University*

Extending SSD Lifetimes with Disk-Based Write Caches

Gokul Soundararajan, *University of Toronto*; Vijayan Prabhakaran, Mahesh Balakrishnan, and Ted Wobber, *Microsoft Research Silicon Valley*

Write Endurance in Flash Drives: Measurements and Analysis

Simona Boboila and Peter Desnoyers, *Northeastern University*

12:30 p.m.–2:00 p.m.

Lunch (on your own)

Thursday, February 25, 2010 (continued)

2:00 p.m.–3:30 p.m.

Thursday

I/O, I/O, TO PARALLEL I/O WE GO

Accelerating Parallel Analysis of Scientific Simulation Data via Zazen

Tiankai Tu, Charles A. Rendleman, Patrick J. Miller, Federico Sacerdoti, and Ron O. Dror, *D.E. Shaw Research*; David E. Shaw, *D.E. Shaw Research and Columbia University*

Efficient Object Storage Journaling in a Distributed Parallel File System

Sarp Oral, Feiyi Wang, David Dillow, Galen Shipman, and Ross Miller, *National Center for Computational Sciences at Oak Ridge National Laboratory*; Oleg Drokin, *Lustre Center of Excellence at Oak Ridge National Laboratory and Sun Microsystems Inc.*

Panache: A Parallel File System Cache for Global File Access

Marc Eshel, Roger Haskin, Dean Hildebrand, Manoj Naik, Frank Schmuck, and Renu Tewari, *IBM Almaden*

4:00 p.m.–5:30 p.m.

Thursday

MAKING MANAGEMENT MORE MANAGEABLE

BASIL: Automated IO Load Balancing Across Storage Devices

Ajay Gulati, Chethan Kumar, and Irfan Ahmad, *VMware, Inc.*; Karan Kumar, *Carnegie Mellon University*

Discovery of Application Workloads from Network File Traces

Neeraja J. Yadwadkar, Chiranjib Bhattacharyya, and K. Gopinath, *Indian Institute of Science*; Thirumale Niranjan and Sai Susarla, *NetApp Advanced Technology Group*

Provenance for the Cloud

Kiran-Kumar Muniswamy-Reddy, Peter Macko, and Margo Seltzer, *Harvard School of Engineering and Applied Sciences*

6:00 p.m.–7:30 p.m.

Thursday

CONFERENCE RECEPTION

Friday, February 26, 2010

9:00 a.m.–10:30 a.m.

Friday

CONCENTRATION: THE DEDUPLICATION GAME

I/O Deduplication: Utilizing Content Similarity to Improve I/O Performance

Ricardo Koller and Raju Rangaswami, *Florida International University*

HydraFS: A High-Throughput File System for the HYDRAsTOR Content-Addressable Storage System

Cristian Ungureanu, *NEC Laboratories America*; Benjamin Atkin, *Google*; Akshat Aranya, Salil Gokhale, and Stephen Rago, *NEC Laboratories America*; Grzegorz Całkowski, *VMware*; Cezary Dubnicki, *9LivesData, LLC*; Aniruddha Bohra, *Akamai*

Bimodal Content Defined Chunking for Backup Streams

Erik Kruus and Cristian Ungureanu, *NEC Laboratories America*; Cezary Dubnicki, *9LivesData, LLC*

11:00 a.m.–12:30 p.m.

Friday

THE POWER BUTTON

Evaluating Performance and Energy in File System Server Workloads

Priya Sehgal, Vasily Tarasov, and Erez Zadok, *Stony Brook University*

SRMap: Energy Proportional Storage Using Dynamic Consolidation

Akshat Verma, *IBM Research, India*; Ricardo Koller, Luis Useche, and Raju Rangaswami, *Florida International University*

Membrane: Operating System Support for Restartable File Systems

Swaminathan Sundararaman, Sriram Subramanian, Abhishek Rajimwale, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, and Michael M. Swift, *University of Wisconsin—Madison*

FAST '10 is co-located with the First USENIX Workshop on Sustainable Information Technology (SustainIT '10) and the 2nd USENIX Workshop on the Theory and Practice of Provenance (TaPP '10)

Both workshops will take place on Monday, February 22, 2010. See <http://www.usenix.org/sustainit10> and <http://www.usenix.org/tapp10> for more information.

Register by Monday, February 8, 2010, and save!

www.usenix.org/fast2010