



# FAST<sup>↑</sup>'12

FEBRUARY 14–17, 2012  
SAN JOSE, CA

sponsored by USENIX  
in cooperation with ACM SIGOPS

## 10th USENIX Conference on File and Storage Technologies

usenix

**The 10th USENIX Conference on File and Storage Technologies (FAST '12) brings together storage system researchers and practitioners to explore new directions in the design, implementation, evaluation, and deployment of storage systems.**

The FAST program will again be offering tutorials. Taking place on Tuesday, February 14, 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 tutorial and get the second for free!

Beginning on Wednesday, February 15, the technical program includes 26 papers, Work-in-Progress reports (WIPs), and two poster sessions. 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 days of ground-breaking file and storage information and training. Register by Monday, January 23, 2012, at [www.usenix.org/fast12](http://www.usenix.org/fast12) for the greatest savings.

### Make your hotel reservation early!

Fairmont San Jose • 170 South Market Street • San Jose, CA 95113 • Phone: (408) 998-1900  
Call and mention USENIX or FAST or book online via <http://www.usenix.org/fast12/hotel>.

### Thanks to Our Sponsors

NetApp  
EMC  
Google  
Microsoft Research  
VMware  
HP  
IBM Research

### Thanks to Our Media Sponsors

The Data Center  
Journal  
Free Software  
Magazine  
Girls in Tech  
Toolbox.com  
UserFriendly.org

## Tutorial Program

Tuesday, February 14, 2012

### Half Day Tutorials (Morning)

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

**Brent Welch, Panasas**

This tutorial is aimed at administrators and developers who use HPC systems, especially those using storage systems in these environments. Storage is often a critical part of the HPC infrastructure. This tutorial will show you how to get the most out of your HPC storage environment, based on a solid understanding of the fundamentals and the use of cluster-based performance tools and programming techniques.

#### T2 Building a Cloud Storage System **NEW!**

**Jeff Darcy, Red Hat**

The trend toward moving computation into the cloud has resulted in new expectations for storage in the cloud. This tutorial will provide information necessary to build your own cloud-appropriate storage system, including new requirements and new constraints, techniques needed, and case studies.

### Half Day Tutorials (Afternoon)

#### T3 Storage Class Memory: Technologies, Systems, and Applications **NEW!**

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

Over the next few years, inexpensive solid-state storage based on flash SSDs or, eventually, storage class memory technology will have a profound impact on the design and use of storage systems. This tutorial is intended for those interested in the design of storage systems for the latter part of this decade. It will focus on the impact the introduction of solid-state memory technologies will have on storage systems.

#### T4 Understanding the I/O of Columnar and NoSQL Databases **NEW!**

**Jiri Schindler, NetApp**

Structured data management systems have always been great consumers of data storage. The tutorial is aimed at those who are familiar with the basics of storage and database technologies and want to gain a better understanding of how columnar and NoSQL databases are organized and how to best use storage systems.

## Conference Organizers

### Program Co-Chairs

William J. Bolosky, *Microsoft Research*  
Jason Flinn, *University of Michigan*

### Program Committee

Atul Adya, *Google, Inc.*  
Andrea Arpaci-Dusseau, *University of Wisconsin—Madison*  
Lakshmi N. Bairavasundaram, *NetApp*  
John Bent, *EMC*  
Randall Burns, *Johns Hopkins University*  
Peter Desnoyers, *Northeastern University*  
Cezary Dubnicki, *9LivesData, LLC*  
Arkady Kanevsky, *Dell*  
Kimberly Keeton, *HP Labs*  
Mark Lillibridge, *HP Labs*  
Darrell Long, *University of California, Santa Cruz*  
James Mickens, *Microsoft Research*  
Dushyanth Narayanan, *Microsoft Research*  
David Patterson, *University of California, Berkeley*  
Daniel Peek, *Facebook*  
James S. Plank, *University of Tennessee*  
Florentina Popovici, *Google, Inc.*  
Raju Rangaswami, *Florida International University*  
Benjamin Reed, *Yahoo! Research*  
Jiri Schindler, *NetApp*  
Margo Seltzer, *Harvard School of Engineering and Applied Sciences and Oracle*  
Keith A. Smith, *NetApp*  
Theodore Wong, *IBM Research*  
Junfeng Yang, *Columbia University*

### Posters and Work-in-Progress Reports (WIPs) Committee

James Mickens, *Microsoft Research*  
Florentina Popovici, *Google, Inc.*  
Jiri Schindler, *NetApp*

### Tutorial Chair

John Strunk, *NetApp*

### Steering Committee

Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*  
Randal Burns, *Johns Hopkins University*  
Greg Ganger, *Carnegie Mellon University*  
Garth Gibson, *Carnegie Mellon University and Panasas*  
Kimberly Keeton, *HP Labs*  
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*

Register by Monday, January 23, 2012, and save!

[www.usenix.org/fast12](http://www.usenix.org/fast12)

## Technical Sessions

### Wednesday, February 15, 2012

9:00 a.m.–10:30 a.m.

Wednesday

#### OPENING REMARKS AND BEST PAPER AWARDS

#### IMPLICATIONS OF NEW STORAGE TECHNOLOGY

##### De-indirection for Flash-based SSDs with Nameless Writes

Yiyi Zhang, Leo Prasath Arulraj, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau, *University of Wisconsin—Madison*

##### The Bleak Future of NAND Flash Memory

Laura M. Grupp, *University of California, San Diego*; John D. Davis, *Microsoft Research*; Steven Swanson, *University of California, San Diego*

##### When Poll Is Better Than Interrupt

Isoo Yang, Dave B. Minturn, and Frank Hady, *Intel Corporation*

11:00 a.m.–12:20 p.m.

Wednesday

#### BACK IT UP

##### WAN Optimized Replication of Backup Datasets Using Stream-Informed Delta Compression

Philip Shilane, Mark Huang, Grant Wallace, and Windsor Hsu, *EMC*

##### Power Consumption in Enterprise-Scale Backup Storage Systems

Zhichao Li, *Stony Brook University*; Kevin M. Greenan and Andrew W. Leung, *EMC*; Erez Zadok, *Stony Brook University*

##### Characteristics of Backup Workloads in Production Systems

Grant Wallace and Fred Dougliis, *EMC*; Hangwei Qian, *Case Western Reserve University*; Philip Shilane, Stephen Smaldone, Mark Chamness, and Windsor Hsu, *EMC*

2:00 p.m.–3:30 p.m.

Wednesday

#### FILE SYSTEM DESIGN AND CORRECTNESS

##### Recon: Verifying File System Consistency at Runtime

Daniel Fryer, Kuei Sun, Rahat Mahmood, Tinghao Cheng, Shaun Benjamin, Ashvin Goel, and Angela Demke Brown, *University of Toronto*

##### Understanding Performance Implications of Nested File Systems in a Virtualized Environment

Duy Le, *College of William and Mary*; Hai Huang, *IBM T.J. Watson Research Center*; Haining Wang, *College of William and Mary*

##### Consistency Without Ordering

Vijay Chidambaram, Tushar Sharma, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*

4:00 p.m.–5:20 p.m.

Wednesday

#### FLASH AND SSDS, PART I

##### Reducing SSD Read Latency via NAND Flash Program and Erase Suspension

Guanying Wu and Xubin He, *Virginia Commonwealth University*

##### Optimizing NAND Flash-Based SSDs via Retention Relaxation

Ren-Shuo Liu and Chia-Lin Yang, *National Taiwan University*; Wei Wu, *Intel Corporation*

##### SFS: Random Write Considered Harmful in Solid State Drives

Changwoo Min, *Samsung Electronics*; Kangnyeon Kim, *Sungkyun-kwan University*; Hyunjin Cho, *Samsung Electronics*; Sang-Won Lee and Young Ik Eom, *Sungkyun-kwan University*

5:30 p.m.–7:30 p.m.

Wednesday

#### POSTER SESSION & RECEPTION

### Thursday, February 16, 2012

9:00 a.m.–10:20 a.m.

Thursday

#### OS TECHNIQUES

##### FIOS: A Fair, Efficient Flash I/O Scheduler

Stan Park and Kai Shen, *University of Rochester*

##### Shredder: GPU-Accelerated Incremental Storage and Computation

Pramod Bhatotia and Rodrigo Rodrigues, *MPI-SWS*; Akshat Verma, *IBM Research—India*

##### Adding Advanced Storage Controller Functionality via Low-Overhead Virtualization

Muli Ben-Yehuda, *IBM Research—Haifa*; Eran Borovik; Michael Factor, Eran Rom, and Avishay Traeger, *IBM Research—Haifa*; Ben-Ami Yassour

10:50 a.m.–12:20 p.m.

Thursday

#### MOBILE AND SOCIAL

##### ZZFS: A Hybrid Device and Cloud File System for Spontaneous Users

Michelle Mazurek, *Carnegie Mellon University*; Eno Thereska, Dinan Gunawardena, Richard Harper, and James Scott, *Microsoft Research*

##### Revisiting Storage for Smartphones

Hyojun Kim, *NEC Labs and Georgia Tech*; Nitin Agrawal and Cristian Ungureanu, *NEC Labs*

##### Serving Large-scale Batch Computed Data with Project Spock

Roshan Sumbaly, Jay Kreps, Lei Gao, Alex Feinberg, Chinmay Soman, and Sam Shah, *LinkedIn Corp*

2:00 p.m.–3:30 p.m.

Thursday

#### WORK-IN-PROGRESS REPORTS (WIPS)

WIPs are the traditional short talks that expose a new idea or project.

4:00 p.m.–5:20 p.m.

Thursday

#### CLOUD

##### BlueSky: A Cloud-Backed File System for the Enterprise

Michael Vrbale, *Google*; Stefan Savage and Geoffrey M. Voelker, *University of California, San Diego*

##### Rethinking Erasure Codes for Cloud File Systems: Minimizing I/O for Recovery and Degraded Reads

Osama Khan and Randal Burns, *Johns Hopkins University*; James Plank and William Pierce, *University of Tennessee*; Cheng Huang, *Microsoft Research*

##### NCCloud: Applying Network Coding for the Storage Repair in a Cloud-of-Clouds

Yuchong Hu, Henry C.H. Chen, and Patrick P.C. Lee, *The Chinese University of Hong Kong*; Yang Tang, *Columbia University*

5:30 p.m.–7:30 p.m.

Thursday

#### POSTER SESSION & RECEPTION

### Friday, February 17, 2012

9:00 a.m.–10:20 a.m.

Friday

#### A LITTLE BIT OF EVERYTHING

##### Extracting Flexible, Replayable Models from Large Block Traces

Vasily Tarasov and Santhosh Kumar, *Stony Brook University*; Jack Ma, *Harvey Mudd College*; Dean Hildebrand and Anna Povzner, *IBM Research—Almaden*; Geoff Kuenning, *Harvey Mudd College*; Erez Zadok, *Stony Brook University*

##### scc: Cluster Storage Provisioning Informed by Application Characteristics and SLAs

Harsha V. Madhyastha, *University of California, Riverside*; John C. McCullough, George Porter, Rishi Kapoor, Stefan Savage, Alex C. Snoeren, and Amin Vahdat, *University of California, San Diego*

##### iDedup: Latency-aware, Inline Data Deduplication for Primary Storage

Kiran Srinivasan, Tim Bisson, Garth Goodson, and Kaladhar Voruganti, *NetApp*

11:00 a.m.–Noon

Friday

#### FLASH AND SSDS, PART II

##### Caching Less for Better Performance: Balancing Cache Size and Update Cost of Flash Memory Cache in Hybrid Storage Systems

Yongseok Oh, *University of Seoul*; Jongmoo Choi, *Dankook University*; Donghee Lee, *University of Seoul*; Sam H. Noh, *Hongik University*

##### Lifetime Management of Flash-Based SSDs Using Recovery-Aware Dynamic Throttling

Sungjin Lee and Taejin Kim, *Seoul National University*; Kyungho Kim, *Samsung Electronics*; Jihong Kim, *Seoul National University*

Register by Monday, January 23, 2012, and save!

[www.usenix.org/fast12](http://www.usenix.org/fast12)