

**USENIX Association**

**Proceedings of the  
2012 USENIX Annual Technical Conference  
(USENIX ATC '12)**

**June 13–15, 2012  
Boston MA, USA**

## Conference Organizers

### Program Co-Chairs

Gernot Heiser, *NICTA and University of New South Wales*  
Wilson Hsieh, *Google Inc.*

### Program Committee

James Anderson, *University of North Carolina, Chapel Hill*  
Katerina Argyraki, *École Polytechnique Fédérale de Lausanne (EPFL)*  
Andrea Arpaci-Dusseau, *University of Wisconsin—Madison*  
Andrew Birrell, *Microsoft Research*  
Pei Cao, *Facebook*  
John Carter, *IBM Research*  
Rohit Chandra, *Yahoo!*  
Haibo Chen, *Shanghai Jiao Tong University*  
Frank Dabek, *Google*  
Angela Demke Brown, *University of Toronto*  
Alexandra Fedorova, *Simon Fraser University*  
Andreas Haeberlen, *University of Pennsylvania*

Jon Howell, *Microsoft Research*  
Sam King, *University of Illinois at Urbana-Champaign*  
Eddie Kohler, *Harvard University*  
Donald Kossmann, *ETH Zurich*  
Hank Levy, *University of Washington*  
Pradeep Padala, *VMware*  
John Regehr, *University of Utah*  
Mahadev Satyanarayanan, *Carnegie Mellon University*  
Wolfgang Schröder-Preikschat, *Friedrich-Alexander University Erlangen-Nuremberg*  
Liuba Shrira, *Brandeis University*  
Emil Sit, *Hadapt*  
Christopher Small, *NetApp*  
Udo Steinberg, *Intel*  
Emmett Witchel, *University of Texas, Austin*

### Poster Session Chair

Emil Sit, *Hadapt*

### The USENIX Association Staff

## External Reviewers

David G. Andersen  
Andrea Bastoni  
Andrew Baumann  
Frank Bellosa  
Tom Bergen  
Jeff Bigham  
George Bissias  
Richard Black  
Marco Canini  
David Cock  
Tobias Distler  
Mihai Dobrescu  
Alan M. Dunn  
Glenn Elliott  
Anja Feldmann  
James Fogarty  
Bryan Ford  
Felix Freiling  
Moises Goldszmidt  
Vishakha Gupta  
Werner Haas

Wanja Hofer  
Anne Holler  
Timo Hönic  
Adam Hupp  
Tomas Isdal  
Shinpei Kato  
Christopher Kenna  
Karl Koscher  
Arvind Krishnamurthy  
Srinivas Krishnan  
Charles Lamb  
Stefan Leue  
Alex Lloyd  
Fabian Monrose  
Toby Murray  
Arjun Narayan  
Hugo Patterson  
Franz J. Rammig  
Charlie Reis  
Luigi Rizzo  
Chris Rossbach

Leonid Ryzhyk  
Fabian Scheler  
Simon Schubert  
Will Scott  
Micah Sherr  
Arrvindh Shriraman  
Mark S. Silberstein  
Michael Sirivianos  
Maxim Smith  
Julian Stecklina  
Isabella Stalkerich  
Mike Swift  
Reinhard Tartler  
Mustafa Uysal  
Matthew Wachs  
Quanqing Xu  
Jie Yu  
Cristian Zamfir  
Mingchen Zhao

**2012 USENIX Annual Technical Conference**  
**June 13–15, 2012**  
**Boston, MA, USA**

Message from the Program Co-Chairs . . . . . vii

**Wednesday, June 13**

**9:00–10:30      Cloud**

Demand Based Hierarchical QoS Using Storage Resource Pools . . . . . 1  
*Ajay Gulati and Ganesh Shanmuganathan, VMware Inc.; Xuechen Zhang, Wayne State University; Peter Varman, Rice University*

Erasure Coding in Windows Azure Storage . . . . . 15  
*Cheng Huang, Huseyin Simitci, Yikang Xu, Aaron Ogus, Brad Calder, Parikshit Gopalan, Jin Li, and Sergey Yekhanin, Microsoft Corporation*

Composable Reliability for Asynchronous Systems . . . . . 27  
*Sunghwan Yoo, Purdue University and HP Labs; Charles Killian, Purdue University; Terence Kelly, HP Labs; Hyoun Kyu Cho, HP Labs and University of Michigan; Steven Plite, Purdue University*

**11:00–12:30      Multicore**

Managing Large Graphs on Multi-Cores with Graph Awareness . . . . . 41  
*Vijayan Prabhakaran, Ming Wu, Xuetian Weng, Frank McSherry, Lidong Zhou, and Maya Haridasan, Microsoft Research*

MemProf: A Memory Profiler for NUMA Multicore Systems . . . . . 53  
*Renaud Lachaize, UJF; Baptiste Lepers, CNRS; Vivien Quéma, GrenobleINP*

Remote Core Locking: Migrating Critical-Section Execution to Improve the Performance of Multithreaded Applications . . . . . 65  
*Jean-Pierre Lozi, Florian David, Gaël Thomas, Julia Lawall, and Gilles Muller, LIP6/INRIA*

**1:30–3:30      Packet Processing**

The Click2NetFPGA Toolchain . . . . . 77  
*Teemu Rinta-aho and Mika Karlstedt, NomadicLab, Ericsson Research; Madhav P. Desai, Indian Institute of Technology (Bombay)*

Building a Power-Proportional Software Router . . . . . 89  
*Luca Niccolini, University of Pisa; Gianluca Iannaccone, RedBow Labs; Sylvia Ratnasamy, University of California, Berkeley; Jaideep Chandrashekar, Technicolor Labs; Luigi Rizzo, University of Pisa and University of California, Berkeley*

netmap: A Novel Framework for Fast Packet I/O . . . . . 101  
*Luigi Rizzo, Università di Pisa, Italy*

Toward Efficient Querying of Compressed Network Payloads . . . . . 113  
*Teryl Taylor, UNC Chapel Hill; Scott E. Coull, RedJack; Fabian Monrose, UNC Chapel Hill; John McHugh, RedJack*

## Thursday, June 14

### 8:30–10:30 Security

Body Armor for Binaries: Preventing Buffer Overflows Without Recompilation . . . . . 125  
*Asia Slowinska, Vrije Universiteit Amsterdam; Traian Stancescu, Google, Inc.; Herbert Bos, Vrije Universiteit Amsterdam*

Abstractions for Usable Information Flow Control in Aeolus . . . . . 139  
*Winnie Cheng, IBM Research; Dan R.K. Ports and David Schultz, MIT CSAIL; Victoria Popic, Stanford; Aaron Blankstein, Princeton; James Cowling and Dorothy Curtis, MIT CSAIL; Liuba Shrira, Brandeis; Barbara Liskov, MIT CSAIL*

Treehouse: Javascript Sandboxes to Help Web Developers Help Themselves . . . . . 153  
*Lon Ingram, The University of Texas at Austin and Waterfall Mobile; Michael Walfish, The University of Texas at Austin*

Cloud Terminal: Secure Access to Sensitive Applications from Untrusted Systems . . . . . 165  
*Lorenzo Martignoni, University of California, Berkeley; Pongsin Poosankam, University of California, Berkeley, and Carnegie Mellon University; Matei Zaharia, University of California, Berkeley; Jun Han, Carnegie Mellon University; Stephen McCamant, Dawn Song, and Vern Paxson, University of California, Berkeley; Adrian Perrig, Carnegie Mellon University; Scott Shenker and Ion Stoica, University of California, Berkeley*

### 11:00–Noon Short Papers: Tools and Networking

Mosh: An Interactive Remote Shell for Mobile Clients . . . . . 177  
*Keith Winstein and Hari Balakrishnan, M.I.T. Computer Science and Artificial Intelligence Laboratory*

TROPIC: Transactional Resource Orchestration Platform in the Cloud . . . . . 183  
*Changbin Liu, University of Pennsylvania; Yun Mao, Xu Chen, and Mary F. Fernández, AT&T Labs—Research; Boon Thau Loo, University of Pennsylvania; Jacobus E. Van der Merwe, AT&T Labs—Research*

Trickle: Rate Limiting YouTube Video Streaming . . . . . 191  
*Monia Ghobadi, University of Toronto; Yuchung Cheng, Ankur Jain, and Matt Mathis, Google*

Tolerating Overload Attacks Against Packet Capturing Systems . . . . . 197  
*Antonis Papadogiannakis, FORTH-ICS; Michalis Polychronakis, Columbia University; Evangelos P. Markatos, FORTH-ICS*

Enforcing Murphy’s Law for Advance Identification of Run-time Failures . . . . . 203  
*Zach Miller, Todd Tannenbaum, and Ben Liblit, University of Wisconsin—Madison*

### 1:30–3:30 Distributed Systems

A Scalable Server for 3D Metaverses. . . . . 209  
*Ewen Cheslack-Postava, Tahir Azim, Behram F.T. Mistree, and Daniel Reiter Horn, Stanford University; Jeff Terrace, Princeton University; Philip Levis, Stanford University; Michael J. Freedman, Princeton University*

Granola: Low-Overhead Distributed Transaction Coordination . . . . . 223  
*James Cowling and Barbara Liskov, MIT CSAIL*

High Performance Vehicular Connectivity with Opportunistic Erasure Coding . . . . . 237  
*Ratul Mahajan, Jitendra Padhye, Sharad Agarwal, and Brian Zill, Microsoft Research*

Server-assisted Latency Management for Wide-area Distributed Systems. . . . . 249  
*Wonho Kim, Princeton University; Kyoungsoo Park, KAIST; Vivek S. Pai, Princeton University*

## Thursday, June 14 (continued)

### 4:00–5:30 Deduplication

- Generating Realistic Datasets for Deduplication Analysis . . . . . 261  
*Vasily Tarasov and Amar Mudrankit, Stony Brook University; Will Buik, Harvey Mudd College; Philip Shilane, EMC Corporation; Geoff Kuenning, Harvey Mudd College; Erez Zadok, Stony Brook University*
- An Empirical Study of Memory Sharing in Virtual Machines . . . . . 273  
*Sean Barker, University of Massachusetts Amherst; Timothy Wood, The George Washington University; Prashant Shenoy and Ramesh Sitaraman, University of Massachusetts Amherst*
- Primary Data Deduplication—Large Scale Study and System Design. . . . . 285  
*Ahmed El-Shimi, Ran Kalach, Ankit Kumar, Adi Oltean, Jin Li, and Sudipta Sengupta, Microsoft Corporation*

## Friday, June 15

### 8:30–10:30 Languages and Tools

- Design and Implementation of an Embedded Python Run-Time System . . . . . 297  
*Thomas W. Barr, Rebecca Smith, and Scott Rixner, Rice University*
- AddressSanitizer: A Fast Address Sanity Checker . . . . . 309  
*Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitriy Vyukov, Google*
- Software Persistent Memory . . . . . 319  
*Jorge Guerra, Leonardo Marmol, Daniel Campello, Carlos Crespo, Raju Rangaswami, and Jinpeng Wei, Florida International University*
- Rivet: Browser-agnostic Remote Debugging for Web Applications . . . . . 333  
*James Mickens, Microsoft Research*

### 11:00–Noon Short Papers: Performance

- Wimpy Nodes with 10GbE: Leveraging One-Sided Operations in Soft-RDMA to Boost Memcached. . . . . 347  
*Patrick Stuedi, Animesh Trivedi, and Bernard Metzler, IBM Research, Zurich*
- Revisiting Software Zero-Copy for Web-caching Applications with Twin Memory Allocation . . . . . 355  
*Xiang Song and Jicheng Shi, Shanghai Jiao Tong University and Fudan University; Haibo Chen, Shanghai Jiao Tong University; Binyu Zang, Shanghai Jiao Tong University and Fudan University*
- Seagull: Intelligent Cloud Bursting for Enterprise Applications. . . . . 361  
*Tian Guo and Upendra Sharma, UMASS Amherst; Timothy Wood, The George Washington University; Sambit Sahu, IBM Watson; Prashant Shenoy, UMASS Amherst*
- The Forgotten ‘Uncore’: On the Energy-Efficiency of Heterogeneous Cores . . . . . 367  
*Vishal Gupta, Georgia Tech; Paul Brett, David Koufaty, Dheeraj Reddy, and Scott Hahn, Intel Labs; Karsten Schwan, Georgia Tech; Ganapati Srinivasa, Intel Corporation*

### 1:00–2:30 OS

- Software Techniques for Avoiding Hardware Virtualization Exits. . . . . 373  
*Ole Agesen, Jim Mattson, Radu Rugina, and Jeffrey Sheldon, VMware*
- AppScope: Application Energy Metering Framework for Android Smartphone Using Kernel Activity Monitoring . . . . . 387  
*Chanmin Yoon, Dongwon Kim, Wonwoo Jung, Chulkoo Kang, and Hojung Cha, Yonsei University, Korea*
- Gdev: First-Class GPU Resource Management in the Operating System. . . . . 401  
*Shinpei Kato, Michael McThrow, Carlos Maltzahn, and Scott Brandt, UC Santa Cruz*

## Friday, June 15 (continued)

### 3:00–5:00      Replication

Gnothi: Separating Data and Metadata for Efficient and Available Storage Replication . . . . .	413
<i>Yang Wang, Lorenzo Alvisi, and Mike Dahlin, The University of Texas at Austin</i>	
Dynamic Reconfiguration of Primary/Backup Clusters . . . . .	425
<i>Alexander Shraer and Benjamin Reed, Yahoo! Research; Dahlia Malkhi, Microsoft Research; Flavio Junqueira, Yahoo! Research</i>	
Surviving Congestion in Geo-Distributed Storage Systems . . . . .	439
<i>Brian Cho, University of Illinois at Urbana-Champaign; Marcos K. Aguilera, Microsoft Research Silicon Valley</i>	
Practical Hardening of Crash-Tolerant Systems. . . . .	453
<i>Miguel Correia, IST-UTL/INESC-ID; Daniel Gómez Ferro, Flavio P. Junqueira, and Marco Serafini, Yahoo! Research</i>	

# Message from the 2012 USENIX Annual Technical Conference Program Co-Chairs

Welcome to the 2012 USENIX Annual Technical Conference!

This year's conference demonstrates that there is a huge amount of interest in USENIX ATC. We received 185 full-paper and 45 short-paper submissions, which is a further 30% increase over the submission record set at last year's conference. Authors were from institutions in over a dozen countries spread over four continents: a truly international community.

From these submissions, the program committee selected 32 full papers and 9 short papers for the conference program. These papers cover the spectrum of systems work, with cloud computing a focus area. We have followed the ATC tradition of publishing both ground-breaking research and practically oriented tools and experience papers. We had hoped to get some papers on experimentally verifying prior results or designs, but that kind of paper is still outside the norm of systems conferences. Maybe next year!

We had a great program committee: a total of 27 members, 11 of whom are from industry and 16 from academic institutions. Program committee members were allowed to submit papers, which were held to a higher standard; the co-chairs were not associated with any submissions. We deeply appreciate the committee's diligence in reviewing papers: 705 reviews in total, for an average of 26 reviews per PC member. We also asked 61 external reviewers to write an additional 69 reviews when it was necessary to get some outside expertise. The review process wound up being three rounds: if no paper had a champion after any round, it did not survive into the next round (although PC members were allowed to initiate discussion at the meeting on any paper). Each paper received two reviews in the first round, and could have received a third review in the second round, and a fourth and possibly fifth review in the third round. The committee met at the Google office in New York City on March 21, 2012, for an all-day meeting; three members had to videoconference into the meeting. After the meeting, every paper was shepherded. We would like to acknowledge the integrity of the authors of two papers that were withdrawn during the shepherding process because those authors could not reproduce their results.

Besides the paper authors and reviewers, who define the research community for USENIX ATC, there are a lot of other people and organizations that deserve mention. First, the USENIX staff has been wonderful to work with; without their support, our jobs would have been much harder. They made it possible for us to focus on getting the conference program defined. Second, Eddie Kohler's HotCRP software was indispensable in running the PC. Third, we thank Google for hosting and helping to sponsor the meeting. Finally, we thank our corporate sponsors: Gold Sponsors Facebook and VMware; Silver Sponsors Google, EMC, and Microsoft Research; and Bronze Sponsor NetApp.

Thank you for participating in the USENIX ATC community, and enjoy the conference!

**Gernot Heiser, *NICTA and University of New South Wales***  
**Wilson Hsieh, *Google***

