

# 2017 USENIX Annual Technical Conference

JULY 12–14, 2017 • SANTA CLARA, CA



## Wednesday, July 12

7:30 am–8:45 am Continental Breakfast Grand Ballroom A–D Foyer

8:45 am–9:00 am Grand Ballroom A–D

**Opening Remarks and Awards** Grand Ballroom A–D

Program Co-Chairs: Dilma Da Silva, *Texas A&M University*, and Bryan Ford, *École Polytechnique Fédérale de Lausanne (EPFL)*

9:00 am–10:00 am

**Keynote Address** Grand Ballroom A–D

**Computer Systems Research in the Post-Virtualization Era**

Ed Bugnion, *École Polytechnique Fédérale de Lausanne (EPFL)*

10:00 am–10:30 am Break with Refreshments Grand Ballroom A–D Foyer

10:30 am–12:10 pm

### Track 1 Grand Ballroom AB

#### Kernel

**Lock-in-Pop: Securing Privileged Operating System Kernels by Keeping on the Beaten Path**

Yiwen Li, Brendan Dolan-Gavitt, Sam Weber, and Justin Cappos, *New York University*

**Fast and Precise Retrieval of Forward and Back Porting Information for Linux Device Drivers**

Julia Lawall, Derek Palinski, Lukas Gnirke, and Gilles Muller, *Sorbonne Universités/UPMC/Inria/LIP6*

**Optimizing the TLB Shutdown Algorithm with Page Access Tracking**

Nadav Amit, *VMware Research*

**Falcon: Scaling IO Performance in Multi-SSD Volumes**

Pradeep Kumar and H. Howie Huang, *The George Washington University*

### Track 2 Grand Ballroom CD

#### Datacenters

**deTector: a Topology-aware Monitoring System for Data Center Networks**

Yanghua Peng, *The University of Hong Kong*; Ji Yang, *Xi'an Jiaotong University*; Chuan Wu, *The University of Hong Kong*; Chuanxiong Guo, *Microsoft Research*; Chengchen Hu, *Xi'an Jiaotong University*; Zongpeng Li, *University of Calgary*

**Pricing Intra-Datacenter Networks with Over-Committed Bandwidth Guarantee**

Jian Guo, Fangming Liu, and Tao Wang, *Key Laboratory of Services Computing Technology and System, Ministry of Education, School of Computer Science and Technology, Huazhong University of Science and Technology*; John C.S. Lui, *The Chinese University of Hong Kong*

**Unobtrusive Deferred Update Stabilization for Efficient Geo-Replication**

Chathuri Gunawardhana, Manuel Bravo, and Luis Rodrigues, *University of Lisbon*

**Don't cry over spilled records: Memory elasticity of data-parallel applications and its application to cluster scheduling**

Călin Iorgulescu and Florin Dinu, *EPFL*; Aunn Raza, *NUST Pakistan*; Wajih Ul Hassan, *UIUC*; Willy Zwaenepoel, *EPFL*

12:10 pm–2:00 pm Lunch (on your own)

2:00 pm–3:40 pm

#### Pursuing Efficiency

**Popularity Prediction of Facebook Videos for Higher Quality Streaming**

Linpeng Tang, *Princeton University*; Qi Huang and Amit Puntambekar, *Facebook*; Ymir Vigfusson, *Emory University & Reykjavik University*; Wyatt Lloyd, *University of Southern California & Facebook*; Kai Li, *Princeton University*

**Squeezing out All the Value of Loaded Data: An Out-of-core Graph Processing System with Reduced Disk I/O**

Zhiyuan Ai, Mingxing Zhang, and Yongwei Wu, *Department of Computer Science and Technology, Tsinghua National Laboratory for Information Science and Technology (TNLIST), Tsinghua University and Research Institute of Tsinghua*; Xuehai Qian, *University of Southern California*; Kang Chen and Weimin Zheng, *Department of Computer Science and Technology, Tsinghua National Laboratory for Information Science and Technology (TNLIST), Tsinghua University, and Research Institute of Tsinghua*

**Ending the Anomaly: Achieving Low Latency and Airtime Fairness in WiFi**

Toke Høiland-Jørgensen, *Karlstad University*; Michał Kazior, *Tieto Poland*; Dave Täht, *TekLibre*; Per Hurtig and Anna Brunstrom, *Karlstad University*

**Persona: A High-Performance Bioinformatics Framework**

Stuart Byma and Sam Whitlock, *EPFL*; Laura Flueratoru, *University Politehnica of Bucharest*; Ethan Tseng, *CMU*; Christos Kozyrakis, *Stanford University*; Edouard Bugnion and James Larus, *EPFL*

#### Let's Talk about GPUs

**SPIN: Seamless Operating System Integration of Peer-to-Peer DMA Between SSDs and GPUs**

Shai Bergman and Tanya Brokman, *Technion*; Tzachi Cohen, *unaffiliated*; Mark Silberstein, *Technion*

**Poseidon: An Efficient Communication Architecture for Distributed Deep Learning on GPU Clusters**

Hao Zhang, *Carnegie Mellon University*; Zeyu Zheng, *Petuum Inc.*; Shizhen Xu and Wei Dai, *Carnegie Mellon University*; Qirong Ho, *Petuum Inc.*; Xiaodan Liang, Zhiting Hu, Jinliang Wei, and Pengtao Xie, *Carnegie Mellon University*; Eric P. Xing, *Petuum Inc.*

**Garaph: Efficient GPU-accelerated Graph Processing on a Single Machine with Balanced Replication**

Lingxiao Ma, Zhi Yang, and Han Chen, *Computer Science Department, Peking University, Beijing, China*; Jilong Xue, *Microsoft Research, Beijing, China*; Yafei Dai, *Institute of Big Data Technologies Shenzhen Key Lab for Cloud Computing Technology & Applications, School of Electronics and Computer Engineering (SECE), Peking University, Shenzhen, China*

**GPU Taint Tracking**

Ari B. Hayes, *Rutgers University*; Lingda Li, *Brookhaven National Laboratory*; Mohammad Hedayati, *University of Rochester*; Jiahuan He and Eddy Z. Zhang, *Rutgers University*; Kai Shen, *Google*

## Wednesday, July 12 (continued)

3:40 pm–4:00 pm

Break with Refreshments

Grand Ballroom A–D Foyer

4:00 pm–5:30 pm

### Track 1

Grand Ballroom AB

#### Virtualization

##### Optimizing the Design and Implementation of the Linux ARM Hypervisor

Christoffer Dall, Shih-Wei Li, and Jason Nieh, *Columbia University*

##### Multi-Hypervisor Virtual Machines: Enabling an Ecosystem of Hypervisor-level Services

Kartik Gopalan, Rohit Kugve, Hardik Bagdi, and Yaohui Hu, *Binghamton University*; Daniel Williams and Nilton Bila, *IBM TJ. Watson Research Center*

##### Preemptive, Low Latency Datacenter Scheduling via Lightweight Virtualization

Wei Chen, *University of Colorado, Colorado Springs*; Jia Rao, *University of Texas at Arlington*; Xiaobo Zhou, *University of Colorado, Colorado Springs*

##### The RCU-Reader Preemption Problem in VMs

Aravinda Prasad and K Gopinath, *Indian Institute of Science, Bangalore*; Paul E. McKenney, *IBM Linux Technology Center, Beaverton*

### Track 2

Grand Ballroom CD

#### Security and Privacy I

##### Bunshin: Compositing Security Mechanisms through Diversification

Meng Xu, Kangjie Lu, Taesoo Kim, and Wenke Lee, *Georgia Institute of Technology*

##### Glamdring: Automatic Application Partitioning for Intel SGX

Joshua Lind, Christian Priebe, Divya Muthukumaran, Dan O’Keeffe, Pierre-Louis Aublin, and Florian Kelbert, *Imperial College London*; Tobias Reiher, *TU Dresden*; David Goltzsche, *TU Braunschweig*; David Eysers, *University of Otago*; Rudiger Kapitza, *TU Braunschweig*; Christof Fetzer, *TU Dresden*; Peter Pietzuch, *Imperial College London*

##### High-Resolution Side Channels for Untrusted Operating Systems

Marcus Hähnel, *TU Dresden, Operating Systems Group*; Weidong Cui and Marcus Peinado, *Microsoft Research*

##### Understanding Security Implications of Using Containers in the Cloud

Byungchul Tak, *Kyungpook National University*; Canturk Isci, Sastry Duri, Nilton Bila, Shripad Nadgowda, and James Doran, *IBM TJ Watson Research Center*

6:30 pm–8:00 pm

#### USENIX ATC '17 Conference Reception

Terra Courtyard

Mingle with fellow attendees at the the Conference Reception. Enjoy dinner, drinks, and the chance to connect with other attendees, speakers, and conference organizers.

# Thursday, July 13

8:00 am–9:00 am	Continental Breakfast	Grand Ballroom A–D Foyer
9:00 am–10:40 am		
<b>Track 1</b> Grand Ballroom AB	<b>Track 2</b> Grand Ballroom CD	
<b>Key-Value Stores and Databases</b> <b>Memshare: a Dynamic Multi-tenant Key-value Cache</b> Asaf Cidon, <i>Stanford University</i> ; Daniel Rushton, <i>University of Utah</i> ; Stephen M. Rumble, <i>Google Inc.</i> ; Ryan Stutsman, <i>University of Utah</i> <b>Replication-driven Live Reconfiguration for Fast Distributed Transaction Processing</b> Xingda Wei, Sijie Shen, Rong Chen, and Haibo Chen, <i>Shanghai Jiao Tong University</i> <b>HiKV: A Hybrid Index Key-Value Store for DRAM-NVM Memory Systems</b> Fei Xia, <i>Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences</i> ; Dejun Jiang, Jin Xiong, and Ninghui Sun, <i>Institute of Computing Technology, Chinese Academy of Sciences</i> <b>TRIAD: Creating Synergies Between Memory, Disk and Log in Log Structured Key-Value Stores</b> Oana Balmau, Diego Didona, Rachid Guerraoui, and Willy Zwaenepoel, <i>EPFL</i> ; Huapeng Yuan, Aashray Arora, Karan Gupta, and Pavan Konka, <i>Nutanix</i>	<b>Invited Talks</b> <b>Visualizing Performance with Flame Graphs</b> Brendan Gregg, <i>Senior Performance Architect, Netflix</i> <b>Performance Superpowers with Enhanced BPF</b> Brendan Gregg, <i>Senior Performance Architect, Netflix</i>	
10:40 am–11:00 am	Break with Refreshments	Grand Ballroom A–D Foyer
11:00 am–12:40 pm		
<b>Help Me Debug</b> <b>Engineering Record And Replay For Deployability</b> Robert O'Callahan and Chris Jones, <i>unaffiliated</i> ; Nathan Froyd, <i>Mozilla Corporation</i> ; Kyle Huey, <i>unaffiliated</i> ; Albert Noll, <i>Swisscom AG</i> ; Nimrod Partush, <i>Technion</i> <b>Proactive error prediction to improve storage system reliability</b> Farzaneh Mahdisoltani, <i>University of Toronto</i> ; Ioan Stefanovici, <i>Microsoft Research</i> ; Bianca Schroeder, <i>University of Toronto</i> <b>Towards Production-Run Heisenbugs Reproduction on Commercial Hardware</b> Shiyong Huang, Bowen Cai, and Jeff Huang, <i>Texas A&amp;M University</i> <b>A DSL Approach to Reconcile Equivalent Divergent Program Executions</b> Luís Pina, Daniel Grumberg, Anastasios Andronidis, and Cristian Cadar, <i>Imperial College London</i>	<b>Networking</b> <b>Titan: Fair Packet Scheduling for Commodity Multiqueue NICs</b> Brent Stephens, Arjun Singhvi, Aditya Akella, and Michael Swift, <i>UW-Madison</i> <b>MopEye: Opportunistic Monitoring of Per-app Mobile Network Performance</b> Daoyuan Wu, <i>Singapore Management University</i> ; Rocky K. C. Chang, Weichao Li, and Eric K. T. Cheng, <i>The Hong Kong Polytechnic University</i> ; Debin Gao, <i>Singapore Management University</i> <b>Emu: Rapid Prototyping of Networking Services</b> Nik Sultana, Salvator Galea, David Greaves, Marcin Wojcik, and Jonny Shipton, <i>University of Cambridge</i> ; Richard Clegg, <i>Queen Mary University of London</i> ; Luo Mai, <i>Imperial College London</i> ; Pietro Bressana and Robert Soule, <i>Università della Svizzera italiana</i> ; Richard Mortier, <i>University of Cambridge</i> ; Paolo Costa, <i>Microsoft Research</i> ; Peter Pietzuch, <i>Imperial College London</i> ; Jon Crowcroft, Andrew W Moore, and Noa Zilberman, <i>University of Cambridge</i> <b>Protego: Cloud-Scale Multitenant IPsec Gateway</b> Jeongseok Son, <i>KAIST and Microsoft Research</i> ; Yongqiang Xiong, <i>Microsoft Research</i> ; Kun Tan, <i>Huawei</i> ; Paul Wang and Ze Gan, <i>Microsoft Research</i> ; Sue Moon, <i>KAIST</i>	
12:40 pm–2:00 pm	USENIX ATC '17 Conference Luncheon	Terra Courtyard

## Thursday, July 13 (continued)

2:00 pm–3:40 pm

### Track 1

Grand Ballroom AB

#### Caching along the Way

##### Cache Modeling and Optimization using Miniature Simulations

Carl Waldspurger, Trausti Saemundson, and Irfan Ahmad, *CachePhysics, Inc.*; Nohhyun Park, *Datos IO, Inc.*

##### Hyperbolic Caching: Flexible Caching for Web Applications

Aaron Blankstein, *Princeton University*; Siddhartha Sen, *Microsoft Research*; Michael J. Freedman, *Princeton University*

##### Execution Templates: Caching Control Plane Decisions for Strong Scaling of Data Analytics

Omid Mashayekhi, Hang Qu, Chinmayee Shah, and Philip Levis, *Stanford University*

##### cHash: Detection of Redundant Compilations via AST Hashing

Christian Dietrich and Valentin Rothberg, *Leibniz Universität Hannover*; Ludwig Füracker and Andreas Ziegler, *Friedrich-Alexander Universität Erlangen-Nürnberg*; Daniel Lohmann, *Leibniz Universität Hannover*

### Track 2

Grand Ballroom CD

#### Best of the Rest

##### Application Crash Consistency and Performance with CCFS

Thanumalayan Sankaranarayanan Pillai, Ramnatthan Alagappan, and Lanyue Lu, *University of Wisconsin—Madison*; Vijay Chidambaram, *The University of Texas at Austin*; Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, *University of Wisconsin—Madison*

Best Paper at FAST '17

##### Push-Button Verification of File Systems via Crash Refinement

Helgi Sigurbjarnarson, James Bornholt, Emina Torlak, and Xi Wang, *University of Washington*

Best Paper at OSDI '16

##### Early Detection of Configuration Errors to Reduce Failure Damage

Tianyin Xu, Xinxin Jin, Peng Huang, and Yuanyuan Zhou, *University of California, San Diego*; Shan Lu, *University of Chicago*; Long Jin, *University of California, San Diego*; Shankar Pasupathy, *NetApp, Inc.*

Best Paper at OSDI '16: [Link to Paper](#)

##### Fast, Lean, and Accurate: Modeling Password Guessability Using Neural Networks

William Melicher, Blase Ur, Sean M. Segreti, Saranga Komanduri, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor, *Carnegie Mellon University*

Best Paper at USENIX Security '16

3:40 pm–4:00 pm

Break with Refreshments

Grand Ballroom A–D Foyer

4:00 pm–5:30 pm

#### Storage

##### Giza: Erasure Coding Objects across Global Data Centers

Yu Lin Chen, *NYU & Microsoft Corporation*; Shuai Mu and Jinyang Li, *NYU*; Cheng Huang, Jin Li, Aaron Ogus, and Douglas Phillips, *Microsoft Corporation*

##### SmartCuckoo: A Fast and Cost-Efficient Hashing Index Scheme for Cloud Storage Systems

Yuanyuan Sun and Yu Hua, *Huazhong University of Science and Technology*; Song Jiang, *University of Texas, Arlington*; Qiuyu Li, Shunde Cao, and Pengfei Zuo, *Huazhong University of Science and Technology*

##### Repair Pipelining for Erasure-Coded Storage

Runhui Li, Xiaolu Li, Patrick P. C. Lee, and Qun Huang, *The Chinese University of Hong Kong*

##### PARIX: Speculative Partial Writes in Erasure-Coded Systems

Huiba Li, *mos.meituan.com*; Yiming Zhang, *NUDT*; Zhiming Zhang, *mos.meituan.com*; Shengyun Liu, Dongsheng Li, Xiaohui Liu, and Yuxing Peng, *NUDT*

#### Multicore

##### E-Team: Practical Energy Accounting for Multi-Core Systems

Till Smejkal and Marcus Hähnel, *TU Dresden*; Thomas Illsche, *Center for Information Services and High Performance Computing (ZIH) Technische Universität Dresden*; Michael Roitzsch, *TU Dresden*; Wolfgang E. Nagel, *Center for Information Services and High Performance Computing (ZIH) Technische Universität Dresden*; Hermann Härtig, *TU Dresden*

##### Scalable NUMA-aware Blocking Synchronization Primitives

Sanidhya Kashyap, Changwoo Min, and Taesoo Kim, *Georgia Institute of Technology*

##### StreamBox: Modern Stream Processing on a Multicore Machine

Hongyu Miao and Heejin Park, *Purdue ECE*; Myeongjae Jeon and Gennady Pekhimenko, *Microsoft Research*; Kathryn S. McKinley, *Google*; Felix Xiazhu Lin, *Purdue ECE*

##### Everything you always wanted to know about multicore graph processing but were afraid to ask

Jasmina Malicevic, Baptiste Lepers, and Willy Zwaenepoel, *EPFL*

6:30 pm–8:00 pm

#### USENIX ATC '17 Poster Session and Happy Hour

Magnolia Room

Check out the cool new ideas and the latest preliminary research on display at the Poster Session and Reception. Take part in discussions with your colleagues over complimentary food and drinks. View the list of accepted posters at [www.usenix.org/atc17/posters](http://www.usenix.org/atc17/posters).

## Friday, July 14

8:00 am–9:00 am

Continental Breakfast

Grand Ballroom A–D Foyer

9:00 am–10:40 am

### Track 1

#### Security and Privacy II

##### Graphene-SGX: A Practical Library OS for Unmodified Applications on SGX

Chia-Che Tsai, *Stony Brook University*; Donald E. Porter, *University of North Carolina at Chapel Hill and Fortanix*; Mona Vij, *Intel Corporation*

##### PrivApprox: Privacy-Preserving Stream Analytics

Do Le Quoc and Martin Beck, *TU Dresden*; Pramod Bhatotia, *University of Edinburgh*; Ruichuan Chen, *Nokia Bell Labs*; Christof Fetzer and Thorsten Strufe, *TU Dresden*

##### Mercury: Bandwidth-Effective Prevention of Rollback Attacks Against Community Repositories

Trishank Karthik Kuppusamy, Vladimir Diaz, and Justin Cappos, *New York University*

##### CAB-Fuzz: Practical Concolic Testing Techniques for COTS Operating Systems

Su Yong Kim, *The Affiliated Institute of ETRI*; Sangho Lee, Insu Yun, and Wen Xu, *Georgia Tech*; Byoungyoung Lee, *Purdue University*; Youngtae Yun, *The Affiliated Institute of ETRI*; Taesoo Kim, *Georgia Tech*

### Track 2

#### Don't Forget the Memory

##### Log-Structured Non-Volatile Main Memory

Qingda Hu, *Tsinghua University*; Jinglei Ren and Anirudh Badam, *Microsoft Research*; Jiwu Shu, *Tsinghua University*; Thomas Moscibroda, *Microsoft Research*

##### Soft Updates Made Simple and Fast on Non-volatile Memory

Mingkai Dong and Haibo Chen, *Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University*

##### SmartMD: A High Performance Deduplication Engine with Mixed Pages

Fan Guo, *University of Science and Technology of China*; Yongkun Li, *University of Science and Technology of China, Collaborative Innovation Center of High Performance Computing, and NUDT*; Yinlong Xu, *University of Science and Technology of China, Anhui Province Key Laboratory of High Performance Computing, and USTC*; Song Jiang, *University of Texas, Arlington*; John C. S. Lui, *The Chinese University of Hong Kong*

##### Elastic Memory Management for Cloud Data Analytics

Jingjing Wang and Magdalena Balazinska, *University of Washington*

10:40 am–11:00 am

Break with Refreshments

Grand Ballroom A–D Foyer

11:00 am–12:40 pm

#### File Systems

##### Improving File System Performance of Mobile Storage Systems Using a Decoupled Defragmenter

Sangwook Shane Hahn, *Seoul National University*; Sungjin Lee, *Daegu Gyeongbuk Institute of Science and Technology*; Cheng Ji, *City University of Hong Kong*; Li-Pin Chang, *National Chiao-Tung University*; Inhyuk Yee, *Seoul National University*; Liang Shi, *Chongqing University*; Chun Jason Xue, *City University of Hong Kong*; Jihong Kim, *Seoul National University*

##### Octopus: an RDMA-enabled Distributed Persistent Memory File System

Youyou Lu, Jiwu Shu, and Youmin Chen, *Tsinghua University*; Tao Li, *University of Florida*

##### iJournaling: Fine-Grained Journaling for Improving the Latency of Fsync System Call

Daejun Park and Dongkun Shin, *Sungkyunkwan University, Korea*

##### Scaling Distributed File Systems in Resource-Harvesting Datacenters

Pulkit A. Misra, *Duke University*; Íñigo Goiri, Jason Kace, and Ricardo Bianchini, *Microsoft Research*

12:40 pm–12:50pm

#### Closing Remarks

Program Co-Chairs: Dilma Da Silva, *Texas A&M University*, and Bryan Ford, *École Polytechnique Fédérale de Lausanne (EPFL)*