

18th USENIX Conference on File and Storage Technologies (FAST '20)

Santa Clara, California, USA
25-27 February 2020

ISBN: 978-1-7138-1522-8

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2020) by Usenix Association
All rights reserved.

Printed with permission by Curran Associates, Inc. (2021)

For permission requests, please contact Usenix Association
at the address below.

Usenix Association
2560 Ninth Street, Suite 215
Berkeley, California, 94710

<https://www.usenix.org/>

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

**FAST '20: 18th USENIX Conference on File and Storage
Technologies February 25–27, 2020
Boston, MA, USA**

Cloud Storage

MAPX: Controlled Data Migration in the Expansion of Decentralized Object-Based Storage Systems 1
Li Wang, *Didi Chuxing*; Yiming Zhang, *NiceX Lab, NUDT*; Jiawei Xu and Guangtao Xue, *SJTU*

Lock-Free Collaboration Support for Cloud Storage Services with Operation Inference and Transformation 13
Jian Chen, Minghao Zhao, and Zhenhua Li, *Tsinghua University*; Ennan Zhai, *Alibaba Group Inc.*; Feng Qian, *University of Minnesota - Twin Cities*; Hongyi Chen, *Tsinghua University*; Yunhao Liu, *Michigan State University & Tsinghua University*; Tianyin Xu, *University of Illinois Urbana-Champaign*

POLARDB Meets Computational Storage: Efficiently Support Analytical Workloads in Cloud-Native Relational Database 29
Wei Cao, *Alibaba*; Yang Liu, *ScaleFlux*; Zhushi Cheng, *Alibaba*; Ning Zheng, *ScaleFlux*; Wei Li and Wenjie Wu, *Alibaba*; Linqiang Ouyang, *ScaleFlux*; Peng Wang and Yijing Wang, *Alibaba*; Ray Kuan, *ScaleFlux*; Zhenjun Liu and Feng Zhu, *Alibaba*; Tong Zhang, *ScaleFlux*

File Systems

Carver: Finding Important Parameters for Storage System Tuning 43
Zhen Cao, *Stony Brook University*; Geoff Kuenning, *Harvey Mudd College*; Erez Zadok, *Stony Brook University*

Read as Needed: Building WiSER, a Flash-Optimized Search Engine 59
Jun He and Kan Wu, *University of Wisconsin—Madison*; Sudarsun Kannan, *Rutgers University*; Andrea Arpaci-Dusseau and Remzi Arpaci-Dusseau, *University of Wisconsin—Madison*

How to Copy Files 75
Yang Zhan, *The University of North Carolina at Chapel Hill and Huawei*; Alexander Conway, *Rutgers University*; Yizheng Jiao and Nirjhar Mukherjee, *The University of North Carolina at Chapel Hill*; Ian Groombridge, *Pace University*; Michael A. Bender, *Stony Brook University*; Martin Farach-Colton, *Rutgers University*; William Jannen, *Williams College*; Rob Johnson, *VMWare Research*; Donald E. Porter, *The University of North Carolina at Chapel Hill*; Jun Yuan, *Pace University*

HPC Storage

Uncovering Access, Reuse, and Sharing Characteristics of I/O-Intensive Files on Large-Scale Production HPC Systems 91
Tirthak Patel, *Northeastern University*; Suren Byna, Glenn K. Lockwood, and Nicholas J. Wright, *Lawrence Berkeley National Laboratory*; Philip Carns and Robert Ross, *Argonne National Laboratory*; Devesh Tiwari, *Northeastern University*

GIFT: A Coupon Based Throttle-and-Reward Mechanism for Fair and Efficient I/O Bandwidth Management on Parallel Storage Systems 103
Tirthak Patel, *Northeastern University*; Rohan Garg, *Nutanix*; Devesh Tiwari, *Northeastern University*

SSD and Reliability

Scalable Parallel Flash Firmware for Many-core Architectures 121
Jie Zhang and Miryeong Kwon, *KAIST*; Michael Swift, *University of Wisconsin-Madison*; Myoungsoo Jung, *KAIST*

A Study of SSD Reliability in Large Scale Enterprise Storage Deployments 137
Stathis Maneas and Kaveh Mahdavian, *University of Toronto*; Tim Emami, *NetApp*; Bianca Schroeder, *University of Toronto*

Making Disk Failure Predictions SMARTer! 151
Sidi Lu and Bing Luo, *Wayne State University*; Tirthak Patel, *Northeastern University*; Yongtao Yao, *Wayne State University*; Devesh Tiwari, *Northeastern University*; Weisong Shi, *Wayne State University*

Performance

An Empirical Guide to the Behavior and Use of Scalable Persistent Memory169
Jian Yang, Juno Kim, and Morteza Hoseinzadeh, *UC San Diego*; Joseph Izraelevitz, *University of Colorado, Boulder*;
Steve Swanson, *UC San Diego*

DC-Store: Eliminating Noisy Neighbor Containers using Deterministic I/O Performance and Resource Isolation ..183
Miryeong Kwon, Donghyun Gouk, and Changrim Lee, *KAIST*; Byounggeun Kim and Jooyoung Hwang, *Samsung*;
Myoungsoo Jung, *KAIST*

GoSeed: Generating an Optimal Seeding Plan for Deduplicated Storage 193
Aviv Nachman and Gala Yadgar, *Technion - Israel Institute of Technology*; Sarai Sheinvald, *Braude College of Engineering*

Key Value Storage

Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebook..... 209
Zhichao Cao, *University of Minnesota, Twin Cities, and Facebook*; Siying Dong and Sagar Vemuri, *Facebook*; David
H.C. Du, *University of Minnesota, Twin Cities*

FPGA-Accelerated Compactions for LSM-based Key-Value Store 225
Teng Zhang, *Alibaba Group, Alibaba-Zhejiang University Joint Institute of Frontier Technologies, Zhejiang University*;
Jianying Wang, Xuntao Cheng, and Hao Xu, *Alibaba Group*; Nanlong Yu, *Alibaba-Zhejiang University Joint Institute of
Frontier Technologies, Zhejiang University*; Gui Huang, Tieying Zhang, Dengcheng He, Feifei Li, and Wei Cao, *Alibaba
Group*; Zhongdong Huang and Jianling Sun, *Alibaba-Zhejiang University Joint Institute of Frontier Technologies,
Zhejiang University*

HotRing: A Hotspot-Aware In-Memory Key-Value Store..... 239
Jiqiang Chen, Liang Chen, Sheng Wang, Guoyun Zhu, Yuanyuan Sun, Huan Liu, and Feifei Li, *Alibaba Group*

Caching

BCW: Buffer-Controlled Writes to HDDs for SSD-HDD Hybrid Storage Server..... 253
Shucheng Wang, Ziyi Lu, and Qiang Cao, *Wuhan National Laboratory for Optoelectronics, Key Laboratory of
Information Storage System*; Hong Jiang, *Department of Computer Science and Engineering, University of Texas at
Arlington*; Jie Yao, *School of Computer Science and Technology, Huazhong University of Science and Technology*;
Yuanyuan Dong and Puyuan Yang, *Alibaba Group*

INFINICACHE: Exploiting Ephemeral Serverless Functions to Build a Cost-Effective Memory Cache 267
Ao Wang and Jingyuan Zhang, *George Mason University*; Xiaolong Ma, *University of Nevada, Reno*; Ali Anwar, Lukas
Rupprecht, Dimitrios Skourtis, and Vasily Tarasov, *IBM Research–Almaden*; Feng Yan, *University of Nevada, Reno*; Yue
Cheng, *George Mason University*

Quiver: An Informed Storage Cache for Deep Learning 283
Abhishek Vijaya Kumar and Muthian Sivathanu, *Microsoft Research India*

Consistency and Reliability

CRaft: An Erasure-coding-supported Version of Raft for Reducing Storage Cost and Network Cost..... 297
Zizhong Wang, Tongliang Li, Haixia Wang, Airan Shao, Yunren Bai, Shangming Cai, Zihan Xu, and Dongsheng Wang,
Tsinghua University

Hybrid Data Reliability for Emerging Key-Value Storage Devices 309
Rekha Pitchumani and Yang-suk Kee, *Memory Solutions Lab, Samsung Semiconductor Inc.*

Strong and Efficient Consistency with Consistency-Aware Durability 323
Aishwarya Ganesan, Ramnatthan Alagappan, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau, *University of
Wisconsin—Madison*