

# **21st USENIX Conference on File and Storage Technologies (FAST'23)**

Santa Clara, California, USA  
21-23 February 2023

ISBN: 978-1-7138-7232-0

**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© (2023) by Usenix Association  
All rights reserved.

Printed with permission by Curran Associates, Inc. (2023)

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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

**21st USENIX Conference on File and Storage Technologies (FAST '23)**  
**February 21–23, 2023**  
**Santa Clara, CA, USA**

**Tuesday, February 21**

**Coding and Cloud Storage**

**Practical Design Considerations for Wide Locally Recoverable Codes (LRCs)**..... 1  
Saurabh Kadekodi, Shashwat Silas, David Clausen and Arif Merchant, *Google*

**ParaRC: Embracing Sub-Packetization for Repair Parallelization in MSR-Coded Storage** .....17  
Xiaolu Li, *Huazhong University of Science and Technology*; Keyun Cheng, Kaicheng Tang, and Patrick P. C. Lee, *The Chinese University of Hong Kong*; Yuchong Hu and Dan Feng, *Huazhong University of Science and Technology*; Jie Li and Ting-Yi Wu, *Huawei Technologies Co., Ltd., Hong Kong*

**InftyDedup: Scalable and Cost-Effective Cloud Tiering with Deduplication**..... 33  
Iwona Kotlarska, Andrzej Jackowski, Krzysztof Lichota, Michal Welnicki, and Cezary Dubnicki, *9LivesData, LLC*; Konrad Iwanicki, *University of Warsaw*

**PERSEUS: A Fail-Slow Detection Framework for Cloud Storage Systems** ..... 49  
Ruiming Lu, *Shanghai Jiao Tong University*; Erci Xu, *Alibaba Inc. and Shanghai Jiao Tong University*; Yiming Zhang, *Xiamen University*; Fengyi Zhu, Zhaosheng Zhu, Mengtian Wang, and Zongpeng Zhu, *Alibaba Inc.*; Guangtao Xue, *Shanghai Jiao Tong University*; Jiwu Shu, *Xiamen University*; Minglu Li, *Shanghai Jiao Tong University and Zhejiang Normal University*; Jiasheng Wu, *Alibaba Inc.*

**Key-Value Stores**

**ADOC: Automatically Harmonizing Dataflow Between Components in Log-Structured Key-Value Stores for Improved Performance**..... 65  
Jinghuan Yu, *City University of Hong Kong*; Sam H. Noh, *UNIST & Virginia Tech*; Young-ri Choi, *UNIST*; Chun Jason Xue, *City University of Hong Kong*

**FUSEE: A Fully Memory-Disaggregated Key-Value Store** ..... 81  
Jiacheng Shen, *The Chinese University of Hong Kong*; Pengfei Zuo, *Huawei Cloud*; Xuchuan Luo, *Fudan University*; Tianyi Yang, *The Chinese University of Hong Kong*; Yuxin Su, *Sun Yat-sen University*; Yangfan Zhou, *Fudan University*; Michael R. Lyu, *The Chinese University of Hong Kong*

**ROLEX: A Scalable RDMA-oriented Learned Key-Value Store for Disaggregated Memory Systems**..... 99  
Pengfei Li, Yu Hua, Pengfei Zuo, Zhangyu Chen, and Jiajie Sheng, *Huazhong University of Science and Technology*

**AI and Storage**

**GL-Cache: Group-level learning for efficient and high-performance caching** ..... 115  
Juncheng Yang, *Carnegie Mellon University*; Ziming Mao, *Yale University*; Yao Yue, *Pelikan Foundation*; K. V. Rashmi, *Carnegie Mellon University*

**SHADE: Enable Fundamental Cacheability for Distributed Deep Learning Training** ..... 135  
Redwan Ibne Seraj Khan and Ahmad Hossein Yazdani, *Virginia Tech*; Yuqi Fu, *University of Virginia*; Arnab K. Paul, *BITS Pilani*; Bo Ji and Xun Jian, *Virginia Tech*; Yue Cheng, *University of Virginia*; Ali R. Butt, *Virginia Tech*

**Intelligent Resource Scheduling for Co-located Latency-critical Services: A Multi-Model Collaborative Learning Approach** ..... 153  
Lei Liu, *Beihang University*; Xinglei Dou and Yuetao Chen, *ICT, CAS; Sys-Inventor Lab*

## Wednesday, February 22

### File Systems

**CJFS: Concurrent Journaling for Better Scalability** ..... 167  
Joontaek Oh, Seung Won Yoo, and Hojin Nam, *KAIST*; Changwoo Min, *Virginia Tech*; Youjip Won, *KAIST*

**Unsafe at Any Copy: Name Collisions from Mixing Case Sensitivities** ..... 183  
Aditya Basu and John Sampson, *The Pennsylvania State University*; Zhiyun Qian, *University of California, Riverside*;  
Trent Jaeger, *The Pennsylvania State University*

**CONF: Analyzing Configuration Dependencies of File Systems for Fun and Profit** ..... 199  
Tabassum Mahmud, Om Rameshwar Gatla, Duo Zhang, Carson Love, Ryan Bumann, and Mai Zheng, *Iowa State University*

**HadaFS: A File System Bridging the Local and Shared Burst Buffer for Exascale Supercomputers** ..... 215  
Xiaobin He, *National Supercomputing Center in Wuxi*; Bin Yang, *Tsinghua University, Dept. of C.S*; *National Supercomputing Center in Wuxi*; Jie Gao and Wei Xiao, *National Supercomputing Center in Wuxi*; Qi Chen, *Tsinghua University, Dept. of C.S*; Shupeng Shi and Dexun Chen, *National Supercomputing Center in Wuxi*; Weiguo Liu, *Shandong University*; Wei Xue, *Tsinghua University, Dept. of C.S*; *Tsinghua University, BNRist.*; *National Supercomputing Center in Wuxi*; Zuo-ning Chen, *Chinese Academy of Engineering*

**Fisc: A Large-scale Cloud-native-oriented File System** ..... 231  
Qiang Li, *Alibaba Group*; Lulu Chen, *Fudan University and Alibaba Group*; Xiaoliang Wang, *Nanjing University*; Shuo Huang, *Alibaba Group*; Qiao Xiang, *Xiamen University*; Yuanyuan Dong, Wenhui Yao, Minfei Huang, Puyuan Yang, Shanyang Liu, Zhaosheng Zhu, Huayong Wang, Haonan Qiu, Derui Liu, Shaozong Liu, Yujie Zhou, Yaohui Wu, Zhiwu Wu, Shang Gao, Chao Han, Zicheng Luo, Yuchao Shao, Gexiao Tian, Zhongjie Wu, Zheng Cao, and Jinbo Wu, *Alibaba Group*; Jiwu Shu, *Xiamen University*; Jie Wu, *Fudan University*; Jiasheng Wu, *Alibaba Group*

### Persistent Memory Systems

**TENET: Memory Safe and Fault Tolerant Persistent Transactional Memory** ..... 247  
R. Madhava Krishnan, *Virginia Tech*; Diyu Zhou, *EPFL*; Wook-Hee Kim, *Konkuk University*; Sudarsun Kannan, *Rutgers University*; Sanidhya Kashyap, *EPFL*; Changwoo Min, *Virginia Tech*

**MadFS: Per-File Virtualization for Userspace Persistent Memory Filesystems** ..... 265  
Shawn Zhong, Chenhao Ye, Guanzhou Hu, Suyan Qu, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau, and Michael Swift, *University of Wisconsin–Madison*

**On Stacking a Persistent Memory File System on Legacy File Systems** ..... 281  
Hobin Woo, *Samsung Electronics*; Daegy Han, *Sungkyunkwan University*; Seungjoon Ha, *Samsung Electronics*; Sam H. Noh, *UNIST & Virginia Tech*; Beomseok Nam, *Sungkyunkwan University*

### Remote Memory

**CITRON: Distributed Range Lock Management with One-sided RDMA** ..... 297  
Jian Gao, Youyou Lu, Minhui Xie, Qing Wang, and Jiwu Shu, *Tsinghua University*

**Patronus: High-Performance and Protective Remote Memory** ..... 315  
Bin Yan, Youyou Lu, Qing Wang, Minhui Xie, and Jiwu Shu, *Tsinghua University*

**More Than Capacity: Performance-oriented Evolution of Pangu in Alibaba** ..... 331  
Qiang Li, *Alibaba Group*; Qiao Xiang, *Xiamen University*; Yuxin Wang, Haohao Song, and Ridi Wen, *Xiamen University and Alibaba Group*; Wenhui Yao, Yuanyuan Dong, Shuqi Zhao, Shuo Huang, Zhaosheng Zhu, Huayong Wang, and Shanyang Liu, Lulu Chen, Zhiwu Wu, Haonan Qiu, Derui Liu, Gexiao Tian, Chao Han, Shaozong Liu, Yaohui Wu, Zicheng Luo, Yuchao Shao, Junping Wu, Zheng Cao, Zhongjie Wu, Jiaji Zhu, and Jinbo Wu, *Alibaba Group*; Jiwu Shu, *Xiamen University*; Jiasheng Wu, *Alibaba Group*

## Thursday, February 23

### IO Stacks

**$\lambda$ -IO: A Unified IO Stack for Computational Storage** ..... 347  
Zhe Yang, Youyou Lu, Xiaojian Liao, Youmin Chen, Junru Li, Siyu He, and Jiwu Shu, *Tsinghua University*

**Revitalizing the Forgotten On-Chip DMA to Expedite Data Movement in NVM-based Storage Systems** ..... 363  
Jingbo Su, Jiahao Li, and Luofan Chen, *University of Science and Technology of China*; Cheng Li, *University of Science and Technology of China and Anhui Province Key Laboratory of High Performance Computing*; Kai Zhang and Liang Yang, *SmartX*; Sam H. Noh, *UNIST & Virginia Tech*; Yinlong Xu, *University of Science and Technology of China and Anhui Province Key Laboratory of High Performance Computing*

**NVMeVirt: A Versatile Software-defined Virtual NVMe Device** ..... 379  
Sang-Hoon Kim, *Ajou University*; Jaehoon Shim, Euidong Lee, Seongyeop Jeong, Ilkueon Kang, and Jin-Soo Kim, *Seoul National University*

**SMRSTORE: A Storage Engine for Cloud Object Storage on HM-SMR Drives** ..... 395  
Su Zhou, Erci Xu, Hao Wu, Yu Du, Jiacheng Cui, Wanyu Fu, Chang Liu, Yingni Wang, Wenbo Wang, Shouqu Sun, Xianfei Wang, Bo Feng, Biyun Zhu, Xin Tong, Weikang Kong, Linyan Liu, Zhongjie Wu, Jinbo Wu, Qingchao Luo, and Jiesheng Wu, *Alibaba Group*

### SSDs and Smartphones

**Multi-view Feature-based SSD Failure Prediction: What, When, and Why** ..... 409  
Yuqi Zhang and Wenwen Hao, *Samsung R&D Institute China Xi'an, Samsung Electronics*; Ben Niu and Kangkang Liu, *Tencent*; Shuyang Wang, Na Liu, and Xing He, *Samsung R&D Institute China Xi'an, Samsung Electronics*; Yongwong Gwon and Chanky Koh, *Samsung Electronics*

**Fast Application Launch on Personal Computing/Communication Devices** ..... 425  
Junhee Ryu, *SK hynix*; Dongeun Lee, *Texas A&M University - Commerce*; Kang G. Shin, *University of Michigan*; Kyungtae Kang, *Hanyang University*

**Integrated Host-SSD Mapping Table Management for Improving User Experience of Smartphones** ..... 441  
Yoona Kim and Inhyuk Choi, *Seoul National University*; Juhyung Park, Jaeheon Lee, and Sungjin Lee, *DGIST*; Jihong Kim, *Seoul National University*