

# **22nd USENIX Conference on File and Storage Technologies (FAST'24)**

Santa Clara, California, USA  
27 – 29 February 2024

ISBN: 978-1-7138-9651-7

**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© (2024) by Usenix Association  
All rights reserved.

Printed with permission by Curran Associates, Inc. (2024)

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)

**22nd USENIX Conference on File and Storage Technologies (FAST '24)**  
**February 27–29, 2024**  
**Santa Clara, CA, USA**

**Tuesday, February 27**

**Distributed Storage**

- TeRM: Extending RDMA-Attached Memory with SSD** ..... 1  
Zhe Yang, Qing Wang, Xiaojian Liao, and Youyou Lu, *Tsinghua University*; Keji Huang, *Huawei Technologies Co., Ltd.*;  
Jiwu Shu, *Tsinghua University*
- Combining Buffered I/O and Direct I/O in Distributed File Systems** .....17  
Yingjin Qian, *Data Direct Networks*; Marc-André Vef, *Johannes Gutenberg University Mainz*; Patrick Farrell and  
Andreas Dilger, *Whamcloud Inc.*; Xi Li and Shuichi Ihara, *Data Direct Networks*; Yinjin Fu, *Sun Yat-Sen University*;  
Wei Xue, *Tsinghua University and Qinghai University*; André Brinkmann, *Johannes Gutenberg University Mainz*
- OmniCache: Collaborative Caching for Near-storage Accelerators** ..... 35  
Jian Zhang and Yujie Ren, *Rutgers University*; Marie Nguyen, *Samsung*; Changwoo Min, *Igalia*; Sudarsun Kannan,  
*Rutgers University*

**Caching**

- Symbiosis: The Art of Application and Kernel Cache Cooperation** ..... 51  
Yifan Dai, Jing Liu, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau, *University of Wisconsin—Madison*
- Optimizing File Systems on Heterogeneous Memory by Integrating DRAM Cache with Virtual  
Memory Management** ..... 71  
Yubo Liu, Yuxin Ren, Mingrui Liu, Hongbo Li, Hanjun Guo, Xie Miao, and Xinwei Hu, *Huawei Technologies Co., Ltd.*;  
Haibo Chen, *Huawei Technologies Co., Ltd. and Shanghai Jiao Tong University*
- Kosmo: Efficient Online Miss Ratio Curve Generation for Eviction Policy Evaluation** ..... 89  
Kia Shakiba, Sari Sultan, and Michael Stumm, *University of Toronto*

**File Systems**

- I/O Passthru: Upstreaming a flexible and efficient I/O Path in Linux** ..... 107  
Kanchan Joshi, Anuj Gupta, Javier González, Ankit Kumar, Krishna Kanth Reddy, Arun George, and Simon Lund,  
*Samsung Semiconductor*; Jens Axboe, *Meta Platforms Inc.*
- Metis: File System Model Checking via Versatile Input and State Exploration** ..... 123  
Yifei Liu and Manish Adkar, *Stony Brook University*; Gerard Holzmann, *Nimble Research*; Geoff Kuenning,  
*Harvey Mudd College*; Pei Liu, Scott A. Smolka, Wei Su, and Erez Zadok, *Stony Brook University*
- RFUSE: Modernizing Userspace Filesystem Framework through Scalable Kernel-Userspace Communication** ... .141  
Kyu-Jin Cho, Jaewon Choi, Hyungjoon Kwon, and Jin-Soo Kim, *Seoul National University*

**Wednesday, February 28**

**Flash Storage**

- The Design and Implementation of a Capacity-Variant Storage System** ..... 159  
Ziyang Jiao and Xiangqun Zhang, *Syracuse University*; Hojin Shin and Jongmoo Choi, *Dankook University*;  
Bryan S. Kim, *Syracuse University*
- I/O in a Flash: Evolution of ONTAP to Low-Latency SSDs** .....177  
Matthew Curtis-Maury, Ram Kesavan, Bharadwaj V R, Nikhil Mattankot, Vania Fang, Yash Trivedi, Kesari Mishra,  
and Qin Li, *NetApp, Inc*
- We Ain't Afraid of No File Fragmentation: Causes and Prevention of Its Performance Impact  
on Modern Flash SSDs** ..... 193  
Yuhun Jun, *Sungkyunkwan University and Samsung Electronics Co., Ltd.*; Shinhyun Park, *Sungkyunkwan University*;  
Jeong-Uk Kang, *Samsung Electronics Co., Ltd.*; Sang-Hoon Kim, *Ajou University*; Euseong Seo, *Sungkyunkwan University*

## Key-Value Systems

- In-Memory Key-Value Store Live Migration with NetMigrate** ..... 209  
Zeyang Zhu, *University of Maryland*; Yibo Zhao, *Boston University*; Zaoxing Liu, *University of Maryland*
- IONIA: High-Performance Replication for Modern Disk-based KV Stores** ..... 225  
Yi Xu, *University of California, Berkeley*; Henry Zhu, *University of Illinois Urbana-Champaign*; Prashant Pandey, *University of Utah*; Alex Conway, *Cornell Tech and VMware Research*; Rob Johnson, *VMware Research*;  
Aishwarya Ganesan and Ramnathan Alagappan, *University of Illinois Urbana-Champaign and VMware Research*
- Physical vs. Logical Indexing with IDEA: Inverted Deduplication-Aware Index** ..... 243  
Asaf Levi, *Technion - Israel Institute of Technology*; Philip Shilane, *Dell Technologies*; Sarai Sheinvald, *Braude College of Engineering*; Gala Yadgar, *Technion - Israel Institute of Technology*
- MI-DAS: Minimizing Write Amplification in Log-Structured Systems through Adaptive Group Number and Size Configuration** ..... 259  
Seonggyun Oh, Jeeyun Kim, and Soyoung Han, *DGIST*; Jaeho Kim, *Gyeongsang National University*; Sungjin Lee, *DGIST*;  
Sam H. Noh, *Virginia Tech*

## Thursday, February 29

### Cloud Storage

- What's the Story in EBS Glory: Evolutions and Lessons in Building Cloud Block Store** ..... 277  
Weidong Zhang, Erci Xu, Qiuping Wang, Xiaolu Zhang, Yuesheng Gu, Zhenwei Lu, Tao Ouyang, Guanqun Dai, Wenwen Peng, Zhe Xu, Shuo Zhang, Dong Wu, Yilei Peng, Tianyun Wang, Haoran Zhang, Jiasheng Wang, Wenyuan Yan, Yuanyuan Dong, Wenhui Yao, Zhongjie Wu, Lingjun Zhu, Chao Shi, Yinhu Wang, Rong Liu, Junping Wu, Jiaji Zhu, and Jiesheng Wu, *Alibaba Group*
- ELECT: Enabling Erasure Coding Tiering for LSM-tree-based Storage** ..... 293  
Yanjing Ren and Yuanming Ren, *The Chinese University of Hong Kong*; Xiaolu Li and Yuchong Hu, *Huazhong University of Science and Technology*; Jingwei Li, *University of Electronic Science and Technology of China*;  
Patrick P. C. Lee, *The Chinese University of Hong Kong*
- MinFlow: High-performance and Cost-efficient Data Passing for I/O-intensive Stateful Serverless Analytics** .... 311  
Tao Li, Yongkun Li, and Wenzhe Zhu, *University of Science and Technology of China*; Yinlong Xu, *Anhui Province Key Laboratory of High Performance Computing, University of Science and Technology of China*, John C. S. Lui, *The Chinese University of Hong Kong*

### AI and Storage

- COLE: A Column-based Learned Storage for Blockchain Systems** ..... 329  
Ce Zhang and Cheng Xu, *Hong Kong Baptist University*; Haibo Hu, *Hong Kong Polytechnic University*;  
Jianliang Xu, *Hong Kong Baptist University*
- Baleen: ML Admission & Prefetching for Flash Caches** ..... 347  
Daniel Lin-Kit Wong, *Carnegie Mellon University*; Hao Wu, *Meta*; Carson Molder, *UT Austin*; Sathya Gunasekar, Jimmy Lu, Snehal Khandkar, and Abhinav Sharma, *Meta*; Daniel S. Berger, *Microsoft and University of Washington*;  
Nathan Beckmann and Gregory R. Ganger, *Carnegie Mellon University*
- Seraph: Towards Scalable and Efficient Fully-external Graph Computation via On-demand Processing** ..... 373  
Tsun-Yu Yang, Yizou Chen, Yuhong Liang, and Ming-Chang Yang, *The Chinese University of Hong Kong*