

Curriculum Vitae

Ethan L. Miller

July 2022

Computer Science & Engineering Department
University of California, Santa Cruz
1156 High Street, MS SOE3
Santa Cruz, CA 95064

MOBILE: +1 (831) 295-8432

EMAIL: elm@ucsc.edu
<https://www.soe.ucsc.edu/~elm/>

EMPLOYMENT HISTORY

2021–2022 On leave from the University of California, Santa Cruz.
2018–present Professor, Computer Science and Engineering Department, University of California, Santa Cruz
2017–2018 Professor, Computer Engineering Department, University of California, Santa Cruz
2014–2019 Veritas Presidential Chair in Storage [formerly Symantec Presidential Chair in Storage & Security],
Jack Baskin School of Engineering, University of California, Santa Cruz
2013–2019 Director, NSF IUCRC Center for Research in Storage Systems, University of California, Santa Cruz
2012–2013 Pure Storage (on 80% leave from the University of California, Santa Cruz)
2009–2013 Site Director, NSF IUCRC Center for Research in Intelligent Storage, University of California, Santa Cruz
2008–2017 Professor, Computer Science Department, University of California, Santa Cruz
2007–present Associate Director, Storage Systems Research Center, University of California, Santa Cruz
2002–2008 Associate Professor, Computer Science Department, University of California, Santa Cruz
2000–2002 Assistant Professor, Computer Science Department, University of California, Santa Cruz
1999 System Architect, Endeca, Cambridge, MA
1994–2000 Assistant Professor, Computer Science and Electrical Engineering Department, University of Maryland Baltimore County
1988–1994 Research Assistant, Computer Science Division, University of California at Berkeley
1988–1990 Teaching Associate, Computer Science Division, University of California at Berkeley
1987–1988 Software Engineer, BBN Laboratories, Cambridge, MA
1986 Summer intern, GTE Government Systems, Rockville, MD

EDUCATION

1995 Ph. D., University of California at Berkeley, Computer Science (advisor: Randy Katz)
Thesis: *Storage Hierarchy Management for Scientific Computing*
1990 M. S., University of California at Berkeley, Computer Science
1987 Sc. B., Brown University, Computer Science, *magna cum laude*

HONORS & AWARDS

2020 Best presentation award, USENIX ATC 2020.
2016 Best Paper award, MASCOTS 2016.
2015 Best Paper award, MSST 2015.
2015 Best Paper award, SYSTOR 2015.
2015 Named **IEEE Fellow**.
2013 Named **ACM Distinguished Scientist**.
2005 Best Long Paper award, StorageSS, 2005.
2004 Best Paper award, MASCOTS 2004.
2001 Elevated to Senior Member, IEEE.
1987 William Gaston Prize for Academic Excellence (award made to top graduating students at Brown University).
1987 Elected to Sigma Xi, Brown University.

GRANTS

In addition to the grants listed below, I have helped bring in industrial funding for the Storage Systems Research Center, which was funded at \$100,000–\$250,000 per year from 2002–2008. From 2009 onward, industry funding was primarily in the form of membership fees to an NSF Industry/University Cooperative Research Center. The UC Santa Cruz site of this Center has received \$2.88 million in membership fees from 2009 through May 2017. Funding companies have included EMC, Exablox, Hewlett Packard, Hitachi, Honeycomb Data, Huawei, IBM, Intel, Avago / LSI, NetApp, Northrop Grumman, Permabit, Pure Storage, Samsung, Sandisk, Scality, Seagate, SK Hynix, Symantec / Veritas, Toshiba, and Western Digital.

- 2021–2024 UCSC PI, *Collaborative Research: CNS Core: Medium: Secure, Reliable, and Efficient Long-Term Storage*, National Science Foundation, (PI, Erez Zadok [Stony Brook University]; co-PI, Omkant Pandey [Stony Brook University]), \$1,214,649 (UCSC portion \$466,503)
- 2019–2024 PI, *Phase II IUCRC CRSS: Center for Research in Storage Systems*, National Science Foundation, \$500,000 (co-PI: Darrell Long, UC Santa Cruz; co-PI, Heiner Litz, UC Santa Cruz)
- 2018–2020 co-PI, *CSR: Small: A Multi-Layered Deniable Steganographic File System*, National Science Foundation, \$496,897 (PI: Darrell Long, UC Santa Cruz)
- 2016–2019 co-PI, *PFI:BIC:RouteMe2: A Cloud-Integrated Sensor Testbed for Assisted Public Transportation*, \$999,846 (PI: Roberto Manduchi, UC Santa Cruz; co-PI: Sri Kurniawan, UC Santa Cruz; co-PI: Adam Millard-Ball, UC Santa Cruz)
- 2013–2019 PI, *IUCRC: A Single-Site IUCRC Center for Research in Storage Systems (CRSS)* National Science Foundation, \$468,850 (co-PI: Darrell Long, UC Santa Cruz). Industrial membership fees for this center have totaled \$1.86 million, as of May 2017.
- 2010–2013 co-PI, *Dynamic Non-Hierarchical File Systems for Exascale Storage*, Department of Energy (PI: Darrell Long), \$1,462,000.
- 2010–2014 co-PI, *LockBox: Enabling Users to Keep Data Safe*, National Science Foundation, \$496,000. (PI: Darrell Long).
- 2010 Gift from the Academy of Motion Picture Arts and Sciences in support of research on long-term archival storage, \$33,000.
- 2009–2013 PI, *Collaborative Research: A Multi-University IUCRC Center on Intelligent Storage*, National Science Foundation, \$275,000 [UC Santa Cruz portion] (PI: David Du, University of Minnesota; co-PI: Darrell Long, UC Santa Cruz). Industrial membership fees for the UC Santa Cruz site of this center totaled \$1.02 million.
- 2009–2013 PI, *Scalable Data Management Using Metadata and Provenance*, National Science Foundation, \$553,000 (co-PI: Darrell Long; PI Margo Seltzer [Harvard] received an additional \$350,000 for a collaborative grant on this project).
- 2009–2013 PI, *Managing and Indexing Exascale Archival Storage Systems*, National Science Foundation, \$489,000.
- 2009–2010 Co-PI, *Trading Storage for Computation*, NASA Ames University-Affiliated Research Center (PI: Darrell Long), \$61,000.
- 2009–2010 Co-PI, *Development of a Collaborative Project for Remotely Sensed Science and Technology*, NASA Ames University-Affiliated Research Center (PI: Raphe Kudela, co-PIs: Darrell Long, Donald Potts, Eli Silver, Michael Loik, Chris Wilmers, Jeff Myers, Liane Guild, Francis Enomoto), \$100,000.
- 2007 co-PI, *ViewFS: Dynamic Name-Spaces for Metadata-Rich File Systems*, Lawrence Livermore National Laboratory, \$74,994 (PI: Darrell Long).
- 2006–2010 Co-PI, *Petascale Data Storage Institute*, \$9,000,000, Department of Energy (lead PI: Garth Gibson; UCSC PI: Darrell Long; other UC Santa Cruz co-PI is Scott Brandt). The Institute funds research at three universities and five national laboratories, with \$1,200,000 going to UC Santa Cruz.
- 2006–2009 Co-PI, *File System Tracing, Replaying, Profiling, and Analysis on HEC Systems*, National Science Foundation, \$760,252 (PI: Erez Zadok, Stony Brook University; co-PI: Klaus Mueller, Stony Brook University).
- 2006–2007 Co-PI, *Institute for Scalable Scientific Data Management*, Los Alamos National Laboratory, \$750,000 (PI: Darrell Long; co-PI: Scott Brandt; co-PI: Carlos Maltzahn).
- 2005–2006 PI, *Adaptive Workload-Aware Algorithms for Heterogeneous Storage Systems*, UC MICRO, \$80,269 (includes \$45,000 gift from Veritas).

- 2005–2006 Co-PI, *Scalable File Systems For High Performance Computing*, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratory, \$250,000 (PI: Scott Brandt; co-PI: Darrell Long; co-PI: Martín Abadi; co-PI: Carlos Maltzahn).
- 2005–2006 Co-PI, *Institute for Scalable Scientific Data Management*, Los Alamos National Laboratory, \$750,000 (PI: Darrell Long; co-PI: Scott Brandt).
- 2005 Co-PI, *Trustworthy Computing Curriculum Development*, Microsoft, \$50,000 (PI: Ira Pohl; co-PI: Martín Abadi; co-PI: Jim Whitehead).
- 2004 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$25,000.
- 2003 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$35,000.
- 2003–2006 PI, *Building High-performance, Reliable Storage Systems Using Magnetic RAM*, National Science Foundation, \$414,000 (co-PI: Scott Brandt).
- 2002–2005 Co-PI, *Scalable File Systems for High Performance Computing*, Department of Energy, \$900,000 (PI: Darrell Long; co-PI: Scott Brandt; co-PI: Katia Obraczka).
- 2002 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$42,000.
- 2001–2002 Co-PI, *Building a High-Performance Storage System from Commodity Components*, Lawrence Livermore National Laboratory, \$65,000 (PI: Darrell Long; co-PI: Scott Brandt).
- 2001 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$38,000.
- 1998–2001 Faculty Fellowship, University of Maryland Institute for Advanced Computer Studies, \$30,000.
- 1997–2001 Co-PI, *Center for Architectures for Data-Driven Information Processing*, Department of Defense, \$3,000,000 (PI: Charles Nicholas; co-PI: David Ebert).
- 1997–1998 PI, *Scalable Benchmarks for Mass Storage Systems*, NASA Ames Research Center, \$200,000.
- 1995–1997 Co-PI, *Scalability of the TELLTALE Dynamic Hypertext Environment*, Department of Defense, \$100,000 (PI: Charles Nicholas).
- 1995–1998 Faculty Fellowship, University of Maryland Institute for Advanced Computer Studies, \$30,000.
- 1989–1993 Graduate Fellowship, National Science Foundation.
- 1988–1989 Cal MICRO Fellowship, University of California at Berkeley.

SCHOLARLY AND CREATIVE WORK

NOTE: (*) denotes a student co-author, and (‡) denotes a student co-author who was one of my advisees at the time the paper was written.

Edited Books

- EB1. Erik Elmroth, Michael Factor, **Ethan L. Miller**, Margo Seltzer (eds.), “Is the Future of Preservation Cloudy? (Dagstuhl Seminar 12472)”, *Dagstuhl Reports* **2**(11), November 2012, pages 102–134.

Chapters in Books

- CH1. Claudia Pearce and **Ethan Miller**, “The TELLTALE Dynamic Hypertext Environment: Approaches to Scalability”, in *Advances in Intelligent Hypertext*, J. Mayfield and C. Nicholas, eds. *Lecture Notes in Computer Science*, Springer-Verlag, October 1997, pages 109–130.

Journal Papers

- J18. Daniel Bittman, Peter Alvaro, Pankaj Mehra, Darrell D. E. Long, **Ethan L. Miller**, “Twizzler: a Data-Centric Operating System for Non-Volatile Memory”, *ACM Transactions on Storage* **17**(2), June 2021.
- J17. Avani Wildani, **Ethan L. Miller**, “Can We Group Storage? Statistical Techniques to Identify Predictive Groupings in Storage System Accesses”, *ACM Transactions on Storage*, **12**(2), February 2016.
- J16. Stephanie Jones*, Ahmed Amer, **Ethan L. Miller**, Darrell D. E. Long, Rekha Pitchumani, Christina Strong*, “Classifying Data to Reduce Long-Term Data Movement in Shingled Write Disks”, *ACM Transactions on Storage* **12**(1), February 2016.

- J15. Alberto Miranda*, Sascha Effert*, Yangwook Kang[‡], **Ethan L. Miller**, Ivan Popov*, André Brinkmann, Tom Friedetzky, Toni Cortes, “Random Slicing: Efficient and Scalable Data Placement for Large-scale Storage Systems”, *ACM Transactions on Storage* **10**(3), July 2014.
- J14. Ian F. Adams[‡], Mark W. Storer, and **Ethan L. Miller**, “Analysis of Workload Behavior in Scientific and Historical Long-Term Data Repositories”, *ACM Transactions on Storage* **8**(2), May 2012. Preliminary version available as Technical Report UCSC-SSRC-11-01.
- J13. Ahmed Amer, JoAnne Holliday, Darrell D. E. Long, **Ethan L. Miller**, Jehan-François Pâris, and Thomas Schwarz, S. J., “Data Management and Layout for Shingled Magnetic Recording”, *IEEE Transactions on Magnetics* **47**(10), October 2011, pages 3691–3697.
- J12. Mark W. Storer[‡], Kevin M. Greenan[‡], **Ethan L. Miller**, and Kaladhar Voruganti, “POTSHARDS—A Secure, Recoverable, Long-Term Archival Storage System”, *ACM Transactions on Storage* **5**(2), June 2009, pages 5:1–5:35.
- J11. Andrew W. Leung[‡], Minglong Shao, Tim Bisson, Shankar Pasupathy, and **Ethan L. Miller**, “High-Performance Metadata Indexing and Search in Petascale Data Storage Systems”, *Journal of Physics: Conference Series* **125** (2008) 012069, July 2008.
- J10. Carlos Maltzahn, Nikhil Bobb*, Mark W. Storer[‡], Damian Eads[‡], Scott A. Brandt, and **Ethan L. Miller**, “Graffiti: A Framework for Testing Collaborative Distributed Metadata”, *Distributed Data & Structures* **7**, *Proceedings in Informatics* **21**, Carleton Scientific, 2007, pages 97–111. Extended version of the WDAS 2006 workshop paper.
- J9. Bo Hong*, Scott A. Brandt, Darrell D. E. Long, **Ethan L. Miller**, and Ying Lin*, “Using MEMS-Based Storage in Computer Systems—MEMS Storage Device Modeling and Management”, *ACM Transactions on Storage* **2**(2), May 2006, pages 139–160.
- J8. Qin Xin[‡], Thomas J. E. Schwarz, and **Ethan L. Miller**, “Availability in Global Peer-To-Peer Storage Systems”, *Distributed Data & Structures* **6**, *Proceedings in Informatics* **20**, Carleton Scientific, 2005, pages 63–77. Extended version of the WDAS 2004 workshop paper.
- J7. Ismail Ari[‡], Ahmed Amer*, Robert Gramacy*, **Ethan L. Miller**, Scott A. Brandt, Darrell D. E. Long, “ACME: Adaptive Caching Using Multiple Experts”, *Distributed Data & Structures* **5**, *Proceedings in Informatics* **14**, Carleton Scientific, 2002, pages 143–158. Extended version of the WDAS 2002 workshop paper.
- J6. **Ethan Miller**, Dan Shen[‡], Junli Liu[‡], and Charles Nicholas, “Performance and Scalability of a Large-Scale N-gram Based Information Retrieval System”, *Journal of Digital Information* **1**(5), January 2000, 25 pages (online refereed journal).
- J5. Christopher Shaw, James Kukla*, Ian Soboroff*, David Ebert, Charles Nicholas, Amen Zwa*, **Ethan Miller**, and D. Aaron Roberts, “Interactive Volumetric Information Visualization for Document Corpus Management”, *International Journal on Digital Libraries* **2**(2–3), 1999, pages 144–156.
- J4. Jeffrey Hollingsworth, **Ethan Miller**, and Kennedy Akala[‡], “Binary Version Management for Computational Grids”, *Parallel Processing Letters* **9**(2), June 1999, pages 215–225.
- J3. **Ethan Miller** and Randy Katz, “RAMA: An Easy-To-Use, High-Performance Parallel File System”, *Parallel Computing* **23**(4), July 1997, pages 419–446.
- J2. David Ebert, Amen Zwa*, **Ethan Miller**, Chris D. Shaw, and D. Aaron Roberts, “Two-handed Volumetric Document Corpus Management”, *IEEE Computer Graphics and Applications* **17**(4), July 1997, pages 60–62.
- J1. Peter Chen, Edward Lee, Ann Drapeau, Ken Lutz, **Ethan Miller**, Srini Seshan, Ken Shirriff, David Patterson, and Randy Katz, “Performance and Design Evaluation of the RAID-II Storage Server”, *Journal of Distributed and Parallel Databases* **2**(3), July 1994, pages 243–260.

Invited Journal Papers

- IJ5. Daniel Bittman[‡], Peter Alvaro, Darrell D. E. Long, **Ethan L. Miller**, “The Flipside: A Bit Flip Saved is Power and Lifetime Earned”, *login*; **44**(2), USENIX Association, Summer 2019.
- IJ4. Andrew W. Leung[‡], Minglong Shao, Tim Bisson, Shankar Pasupathy, and **Ethan L. Miller**, “Spyglass: Fast, Scalable Metadata Search for Large-Scale Storage Systems”, *login*; **34**(3), USENIX Association, June 2009.

- IJ3. Avishay Traeger, Erez Zadok, **Ethan L. Miller**, and Darrell D. E. Long, “Findings from the First Annual Storage and File Systems Benchmarking Workshop”, *login*: **33**(5), USENIX Association, October 2008, pages 113–117.
- IJ2. Mark W. Storer[‡], Kevin M. Greenan[‡], **Ethan L. Miller**, and Kaladhar Voruganti, “Pergamum: energy-efficient archival storage with disk instead of tape”, *login*: **33**(3), USENIX Association, June 2008, pages 15–21.
- IJ1. Randy Katz, John Ousterhout, David Patterson, Peter Chen, Ann Chervenak, Rich Drewes, Garth Gibson, Edward Lee, Ken Lutz, Ethan Miller, and Mendel Rosenblum, “A Project on High-Performance I/O Subsystems”, *Computer Architecture News* **17**(5):24–31, September 1989.

Refereed Conference & Workshop Papers

- C145. Devashish Purandare[‡], Daniel Bittman[‡], **Ethan L. Miller**, “Analysis and Workload Characterization of the CERN EOS Storage System”, *Proceedings of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems (CHEOPS '22)*, April 2022.
- C144. Daniel Bittman[‡], Robert Soulé, **Ethan L. Miller**, Vishal Shrivastav, Pankaj Mehra, Matthew Boisvert, Avi Silberschatz, Peter Alvaro, “Don’t Let RPCs Constrain Your API”, *Proceedings of HotNets '21*, November 2021.
- C143. James Byron[‡], **Ethan L. Miller**, Darrell D. E. Long, “Measuring the Cost of Reliability in Archival Systems”, *Proceedings of the Conference on Mass Storage Systems and Technologies (MSST '20)*, October 2020.
- C142. Austen Barker*, Yash Gupta*, Sabrina Au*, Eugene Chou*, **Ethan L. Miller**, Darrell D. E. Long, “Artifice: Data in Disguise”, *Proceedings of the Conference on Mass Storage Systems and Technologies (MSST '20)*, October 2020.
- C141. Oceane Bel*, Kenneth Chang*, Nathan Tallent, Dirk Duellman, **Ethan L. Miller**, Faisal Nawab, Darrell D. E. Long, “Geomancy: Automated Performance Enhancement through Data Layout Optimization”, *Proceeding of the Conference on Mass Storage Systems and Technologies (MSST '20)*, October 2020.
- C140. Daniel Bittman[‡], Pankaj Mehra, Peter Alvaro, Darrell D. E. Long, **Ethan L. Miller**, “Twizzler: An Operating System Designed for Byte-Addressable Persistent Memory”, *Proceedings of the 2020 USENIX Annual Technical Conference*, Boston, MA, July 2020. Won Best Presentation award.
- C139. Daniel Bittman[‡], Peter Alvaro, **Ethan L. Miller**, “A Persistent Problem: Managing Pointers in NVM”, *Proceedings of the 2019 Workshop on Programming Languages and Operating Systems (PLOS '19)*, held in conjunction with SOSP '19, Huntsville, ON, Canada, October 2019.
- C138. Yuanjiang Ni[‡], Jishen Zhao, Heiner Litz, Daniel Bittman[‡], **Ethan L. Miller**, “SSP: Eliminating Redundant Writes in Failure-Atomic NVRAMs via Shadow Sub-Paging”, *Proceedings of the 52nd IEEE/ACM International Symposium on Microarchitecture (MICRO-52)*, Columbus, OH, October 2019.
- C137. Austen Barker*, Staunton Sample*, Yash Gupta*, Ana McTaggart[‡], **Ethan L. Miller**, Darrell D. E. Long, “Artifice: A Deniable Steganographic File System”, *Proceedings of the 9th USENIX Workshop on Free and Open Communications on the Internet (FOCI '19)*, Santa Clara, CA, August 2019.
- C136. Daniel Bittman[‡], Peter Alvaro, Darrell D. E. Long, **Ethan L. Miller**, “A Tale of Two Abstractions: The Case for Object Space”, *Proceedings of the 11th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage '19)*, Renton, WA, July 2019.
- C135. Daniel Bittman[‡], **Ethan L. Miller**, Peter Alvaro, “Co-evolving Tracing and Fault Injection with Box of Pain”, *Proceedings of the 11th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud '19)*, Renton, WA, July 2019.
- C134. Daniel Bittman[‡], Peter Alvaro, Darrell D. E. Long, **Ethan L. Miller**, “Optimizing Systems for Byte-Addressable NVM by Reducing Bit Flipping”, *Proceedings of the 17th Conference on File and Storage Technologies (FAST '19)*, Boston, MA, February 2019.
- C133. Matheus Ogleari*, Ye Yu, Chen Qian, **Ethan Miller**, Jishen Zhao, “String Figure: A Scalable and Elastic Memory Network Architecture”, *Proceedings of the 25th IEEE International Symposium on High-Performance Computer Architecture (HPCA 2019)*, Washington, DC, February 2019.

- C132. James Byron[‡], **Ethan L. Miller**, Darrell D. E. Long, “Using Simulation to Design Scalable and Cost-Efficient Archival Storage Systems”, *Proceedings of the 26th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2018)*, Milwaukee, WI, September 2018.
- C131. Sinjoni Mukhopadhyay*, Joel Frank*, Daniel Bittman[‡], Darrell D. E. Long, **Ethan L. Miller**, “Efficient Reconstruction Techniques for Disaster Recovery in Secret-Split Datastores”, *Proceedings of the 26th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2018)*, Milwaukee, WI, September 2018.
- C130. Daniel Bittman[‡], Matthew Gray*, Justin Raizes*, Sinjoni Mukhopadhyay*, Matt Bryson[‡], Peter Alvaro, Darrell D. E. Long, **Ethan L. Miller**, “Designing Data Structures to Minimize Bit Flips on NVM”, *Proceedings of the 7th IEEE Non-Volatile Memory Systems and Applications Symposium (NVMSA 2018)*, Hakodate, Japan, August 2018.
- C129. Yuanjiang Ni[‡], Jishen Zhao, Daniel Bittman[‡], **Ethan Miller**, “Reducing NVM Writes with Optimized Shadow Paging”, *Proceedings of HotStorage 2018*, Boston, MA: USENIX, July 2018.
- C128. Veronica Estrada-Galinañes, **Ethan L. Miller**, Pascal Felber, Jehan-François Pâris, “Alpha Entanglement Codes: Practical Erasure Codes to Archive Data in Unreliable Environments”, *Proceedings of the 48th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2018)*, Luxembourg City, Luxembourg, June 2018.
- C127. Oceane Bel*, Kenneth Chang[‡], Daniel Bittman[‡], **Ethan L. Miller**, Darrell D. E. Long, Hiroshi Isozaki, “Inkpack: A Secure, Data-Exposure Resistant Storage System” *Proceedings of SYSTOR 2018: The 11th Annual International Systems and Storage Conference*, Haifa, Israel: ACM, June 2018.
- C126. Matheus Ogleari*, **Ethan L. Miller**, Jishen Zhao, “Steal but no force: Efficient Hardware-driven Undo+Redo Logging for Persistent Memory Systems”, *Proceedings of the 24th IEEE International Symposium on High-Performance Computer Architecture (HPCA 2018)*, Vienna, Austria, February 2018.
- C125. Yan Li*, Kenneth Chang[‡], Oceane Bel*, **Ethan L. Miller**, Darrell D. E. Long, “CAPES: Unsupervised Storage Performance Tuning Using Neural Network-Based Deep Reinforcement Learning”, *Proceedings of SC2017*, Denver, CO: ACM & IEEE, November 2017. Nominated for Best Student Paper.
- C124. Thomas Schwarz, Ahmed Amer, Thomas Kroeger, **Ethan L. Miller**, Darrell D. E. Long, Jehan-François Pâris, “RESAR: Reliable Storage at Exabyte Scale”, *Proceedings of the 24th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2016)*, London, UK: IEEE, September 2016. Won Best Paper Award.
- C123. Yan Li*, **Ethan L. Miller**, Darrell D. E. Long, Yash Gupta*, “Pilot: A Framework that Understands How to Do Performance Benchmarks The Right Way”, *Proceedings of the 24th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2016)*, London, UK: IEEE, September 2016.
- C122. Preeti Gupta[‡], Avani Wildani, Darrell D. E. Long, **Ethan L. Miller**, David S. H. Rosenthal, “Effects of Prolonged Media Usage and Long-term Planning on Archival Systems”, *Proceedings of the 32nd International Conference on Mass Storage Systems and Technologies (MSST 2016)*, Santa Clara, CA, May 2016.
- C121. Yan Li*, Xioayuan Lu[‡], **Ethan L. Miller**, Darrell D. E. Long, “ASCAR: Automating Contention Management for High-Performance Storage Systems”, *Proceedings of the 31st International Conference on Mass Storage Systems and Technologies (MSST 2015)*, Santa Clara, CA, June 2015.
- C120. Joel C. Frank*, Shayna M. Frank[‡], Lincoln A. Thurlow[‡], Thomas M. Kroeger, **Ethan L. Miller**, Darrell D. E. Long, “Percival: A Searchable Secret Split Datastore”, *Proceedings of the 31st International Conference on Mass Storage Systems and Technologies (MSST 2015)*, Santa Clara, CA, June 2015.
- C119. Stephanie Jones*, Ahmed Amer, **Ethan L. Miller**, Darrell D. E. Long, Rekha Pitchumani[‡], Christina Strong*, “Classifying Data to Reduce Long Term Data Movement in Shingled Write Disks”, *Proceedings of the 31st International Conference on Mass Storage Systems and Technologies (MSST 2015)*, Santa Clara, CA, June 2015. Won Best Paper Award.
- C118. Rekha Pitchumani[‡], Shayna Frank[‡], **Ethan L. Miller**, “Realistic Request Arrival Generation In Storage Benchmarks”, *Proceedings of the 31st International Conference on Mass Storage Systems and Technologies (MSST 2015)*, Santa Clara, CA, June 2015.

- C117. John Colgrove, John D. Davis, John Hayes, **Ethan L. Miller**, Cary Sandvig, Russell Sears, Ari Tamches, Neil Vachharajani, Feng Wang, “Purity: Building Fast, Highly-Available Enterprise Flash Storage from Commodity Components”, *Proceedings of SIGMOD 2015: Industrial Track*, Melbourne, Australia, June 2015.
- C116. Rekha Pitchumani[‡], James Hughes, **Ethan L. Miller**, “SMRDB: Key-Value Data Store for Shingled Magnetic Recording Disks”, *Proceedings of SYSTOR 2015: The 8th Annual International Systems and Storage Conference*, Haifa, Israel: ACM, May 2015. Won Best Paper Award.
- C115. Preeti Gupta[‡], Avani Wildani, Ian F. Adams, Christina Strong*, Daniel Rosenthal[‡], **Ethan L. Miller**, Andy Hospodor, “An Economic Perspective of Disk vs. Flash Media in Archival Storage”, *Proceedings of the 22nd IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2014)*, Paris, France: IEEE, September 2014.
- C114. Avani Wildani, **Ethan L. Miller**, Ian F. Adams, Darrell D. E. Long, “PERSES: Data Layout for Low Impact Failures”, *Proceedings of the 22nd IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2014)*, Paris, France: IEEE, September 2014. Preliminary version available as Technical Report UCSC-SSRC-12-06.
- C113. Yangwook Kang[‡], Thomas Marlette[‡], Rekha Pitchumani[‡], **Ethan L. Miller**, “Muninn: A Versioning Flash Key-Value Store Using an Object-based Storage Model”, *Proceedings of SYSTOR 2014: The 7th Annual International Systems and Storage Conference*, Haifa, Israel, June 2014.
- C112. Aleatha Parker-Wood*, Darrell D. E. Long, **Ethan L. Miller**, Philippe Rigaux, and Andy Isaacson, “A File By Any Other Name: Managing File Names with Metadata”, *Proceedings of SYSTOR 2014: The 7th Annual International Systems and Storage Conference*, Haifa, Israel, June 2014.
- C111. Avani Wildani[‡], Ian F. Adams[‡], **Ethan L. Miller**, “Single-Snapshot File System Analysis”, *Proceedings of the 21st IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2013)*, San Francisco, CA: IEEE, August 2013.
- C110. Ian F. Adams[‡], Mark W. Storer, Avani Wildani[‡], **Ethan L. Miller**, Brian A. Madden*, “Validating Storage System Instrumentation”, *Proceedings of the 21st IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2013)*, San Francisco, CA: IEEE, August 2013.
- C109. Yangwook Kang[‡], Yang-Suk Kee, **Ethan L. Miller**, Chanik Park, “Enabling Cost-effective Data Processing with Smart SSD”, *Proceedings of the 29th IEEE Symposium on Mass Storage Systems and Technologies (MSST 2013)*, Long Beach, CA: IEEE, May 2013.
- C108. Avani Wildani[‡], **Ethan L. Miller**, Ohad Rodeh, “HANDS: A Heuristically Arranged Non-Backup In-line Deduplication System”, *Proceedings of the 29th IEEE International Conference on Data Engineering (ICDE 2013)*, Brisbane, Australia: IEEE, April 2013. Preliminary version available as Technical Report UCSC-SSRC-12-03.
- C107. Yan Li*, Nakul Sanjay Dhotre*, Yasuhiro Ohara, Thomas M. Kroeger, Darrell D. E. Long, **Ethan L. Miller**, “Horus: Fine-Grained Encryption-Based Security for Large-Scale Storage”, *Proceedings of the 11th Conference on File and Storage Technologies (FAST ’13)*, San Jose, CA: USENIX, February 2013.
- C106. James S. Plank, Kevin M. Greenan, **Ethan L. Miller**, “Screaming Fast Galois Field Arithmetic Using Intel SIMD Extensions”, *Proceedings of the 11th Conference on File and Storage Technologies (FAST ’13)*, San Jose, CA: USENIX, February 2013.
- C105. Ian F. Adams[‡], Brian A. Madden*, Joel C. Frank[‡], Mark W. Storer, **Ethan L. Miller**, Gene Harano, “Usage Behavior of a Large-Scale Scientific Archive”, *Proceedings of SC2012*, Salt Lake City, UT: ACM, November 2012.
- C104. David S. H. Rosenthal, Daniel C. Rosenthal[‡], **Ethan L. Miller**, Ian F. Adams[‡], Mark W. Storer, Erez Zadok, “The Economics of Long-Term Digital Storage”, *The Memory of the World in the Digital Age: Digitization and Preservation*, Vancouver, BC, September 2012.
- C103. Joel C. Frank[‡], **Ethan L. Miller**, Ian F. Adams[‡], Daniel C. Rosenthal[‡], “Evolutionary Trends in a Supercomputing Tertiary Storage Environment”, *Proceedings of the 20th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2012)*, Washington, DC, August 2012.

- C102. Rekha Pitchumani[‡], Andy Hospodor, Ahmed Amer, Yangwook Kang[‡], **Ethan L. Miller**, and Darrell D. E. Long, “Emulating a Shingled Write Disk”, *Proceedings of the 20th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2012)*, Washington, DC, August 2012.
- C101. Yan Li*, **Ethan L. Miller**, and Darrell D. E. Long, “Understanding Data Survivability in Archival Storage Systems”, *Proceedings of SYSTOR 2012: The 5th Annual International Systems and Storage Conference*, Haifa, Israel, June 2012.
- C100. Alberto Miranda*, Sascha Effert*, Yangwook Kang[‡], **Ethan L. Miller**, André Brinkmann, and Toni Cortes, “Reliable and Randomized Data Distribution Strategies for Large Scale Storage Systems”, *Proceedings of the 2011 High Performance Computing Conference (HiPC 2011)*, Bengaluru, India, December 2011.
- C99. Ranjana Rajendran*, **Ethan L. Miller**, and Darrell D. E. Long, “Horus: Fine-Grained Encryption-Based Security for High Performance Petascale Storage”, *Proceedings of the 6th Parallel Data Storage Workshop (PDSW11)*, held in conjunction with *SC2011*, Seattle, WA, November 2011.
- C98. Stephanie N. Jones*, Christina R. Strong*, Darrell D. E. Long, and **Ethan L. Miller** “Tracking Emigrant File Data via Transient Provenance”, *Proceedings of the 2011 Workshop on the Theory and Practice of Provenance (TaPP '11)*, Heraklion, Crete, Greece, June 2011.
- C97. Avani Wildani[‡], **Ethan L. Miller**, and Lee Ward, “Efficiently Identifying Working Sets in Block I/O Streams”, *Proceedings of SYSTOR 2011: The 4th Annual International Systems and Storage Conference*, Haifa, Israel, May 2011.
- C96. Yangwook Kang[‡], Jingpei Yang[‡], **Ethan L. Miller**, “Object-based SCM: An Efficient Interface for Storage Class Memories”, *Proceedings of the 27th IEEE Symposium on Mass Storage Systems and Technologies (MSST 2011)*, Denver, CO: IEEE, May 2011.
- C95. Avani Wildani[‡], **Ethan L. Miller**, “Semantic Data Placement for Power Management in Archival Storage”, *Proceedings of the 5th International Workshop on Petascale Data Storage (PDSW10)*, held in conjunction with *SC2010*, New Orleans, LA, November 2010.
- C94. Ian Adams[‡], Mark W. Storer, and **Ethan L. Miller**, “Examining Energy Use in Heterogeneous Archival Storage Systems”, *Proceedings of the 18th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2010)*, Miami, FL: IEEE, September 2010, pages 297–306.
- C93. Yangwook Kang[‡], Jingpei Yang[‡], and **Ethan L. Miller**, “Efficient Storage Management for Object-based Flash Memory”, *Proceedings of the 18th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2010)*, Miami, FL: IEEE, September 2010, pages 407–409.
- C92. Aleatha Parker-Wood[‡], Christina Strong*, **Ethan L. Miller**, and Darrell D. E. Long, “Security Aware Partitioning for Efficient File System Search”, *Proceedings of the 26th IEEE Symposium on Mass Storage Systems and Technologies (MSST 2010)*, Incline Village, NV: IEEE, May 2010, 14 pages.
- C91. Ahmed Amer, Darrell D. E. Long, **Ethan L. Miller**, Jehan-François Pâris, and Thomas Schwarz, “Design Issues for a Shingled Write Disk System”, *Proceedings of the 26th IEEE Symposium on Mass Storage Systems and Technologies (MSST 2010)*, Incline Village, NV: IEEE, May 2010, 12 pages.
- C90. Yangwook Kang[‡] and **Ethan L. Miller**, “Adding Aggressive Error Correction to a High-Performance Flash File System”, *Proceedings of the 9th ACM & IEEE Conference on Embedded Software (EMSOFT '09)*, Grenoble, France : ACM / IEEE, October 2009, pages 305–314.
- C89. Avani Wildani[‡], Thomas J. E. Schwarz, **Ethan L. Miller**, and Darrell D. E. Long, “Protecting Against Rare Event Failures in Archival Systems”, *Proceedings of the 17th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2009)*, London, UK: IEEE, September 2009, 11 pages. Preliminary version available as Technical Report UCSC-SSRC-09-03.
- C88. Kevin Greenan[‡], Darrell D. E. Long, **Ethan L. Miller**, Thomas Schwarz, and Avani Wildani[‡], “Building Flexible, Fault-Tolerant Flash-based Storage Systems”, *Proceedings of the 5th Workshop on Hot Topics in System Dependability (HotDep 2009)*, Estoril, Portugal, June 2009.

- C87. Ian F. Adams[‡], Darrell D. E. Long, **Ethan L. Miller**, Shankar Pasupathy, and Mark W. Storer, “Maximizing Efficiency By Trading Storage for Computation”, *Proceedings of the Workshop on Hot Topics in Cloud Computing (HotCloud '09)*, San Diego, CA: USENIX, June 2009, 5 pages.
- C86. Keren Jin[‡] and **Ethan L. Miller**, “The Effectiveness of Deduplication on Virtual Machine Disk Images”, *Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference*, Haifa, Israel, May 2009, 12 pages.
- C85. Andrew W. Leung[‡], Minglong Shao, Tim Bisson, Shankar Pasupathy, and **Ethan L. Miller**, “Spyglass: Fast, Scalable Metadata Search for Large-Scale Storage Systems”, *Proceedings of the 7th Conference on File and Storage Technologies (FAST '09)*, San Francisco, CA: USENIX, February 2009, pages 153–166.
- C84. Kevin M. Greenan[‡], Darrell D. E. Long, **Ethan L. Miller**, Thomas J. Schwarz, S. J., and Jay J. Wylie, “A Spin-Up Saved is Energy Earned: Achieving Power-Efficient, Erasure-Coded Storage”, *Proceedings of the 4th Workshop on Hot Topics in System Dependability (HotDep'08)*, San Diego, CA: USENIX, December 2008.
- C83. Mark W. Storer[‡], Kevin M. Greenan[‡], Ian Adams[‡], **Ethan L. Miller**, Darrell D. E. Long, and Kaladhar Voruganti, “Logan: Automatic Management for Evolvable, Large-Scale, Archival Storage”, *Proceedings of the 3rd International Workshop on Petascale Data Storage (PDSW08)*, held in conjunction with SC2008, Austin, TX, November 2008.
- C82. Andrew W. Leung[‡] and **Ethan L. Miller**, “Scalable Full-Text Search for Petascale File Systems”, *Proceedings of the 3rd International Workshop on Petascale Data Storage (PDSW08)*, held in conjunction with SC2008, Austin, TX, November 2008.
- C81. Mark W. Storer[‡], Kevin Greenan[‡], Darrell D. E. Long, and **Ethan L. Miller**, “Secure Data Deduplication”, *Proceedings of the 4th Workshop on Storage Security and Survivability (StorageSS 2008)*, held in conjunction with the 15th ACM Conference on Computer and Communications Security (CCS 2008), Alexandria, VA, October 2008.
- C80. Mohammed G. Khatib^{*}, **Ethan L. Miller** and Pieter H. Hartel, “Workload-Based Configuration of MEMS-Based Storage Devices for Mobile Systems”, *Proceedings of the 8th ACM & IEEE Conference on Embedded Software (EMSOFT '08)*, Atlanta, GA: ACM / IEEE, October 2008.
- C79. Kevin M. Greenan[‡], **Ethan L. Miller**, and Thomas Schwarz, “Optimizing Galois Field Arithmetic for Diverse Processor Architectures”, *Proceedings of the 16th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2008)*, Baltimore, MD: IEEE, September 2008, 10 pages (CD-ROM).
- C78. Andrew W. Leung[‡], Shankar Pasupathy, Garth Goodson, and **Ethan L. Miller**, “Measurement and Analysis of Large-Scale Enterprise Network File System Workloads”, *Proceedings of the 2008 USENIX Annual Technical Conference (USENIX '08)*, Boston, MA: USENIX, June 2008, pages 213–226.
- C77. Kevin M. Greenan[‡], **Ethan L. Miller**, and Jay J. Wylie, “Reliability of flat XOR-based erasure codes on heterogeneous devices”, *Proceedings of the 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2008)*, Anchorage, AK: IEEE, June 2008, pages 147–156.
- C76. Neerja Bhatnagar[‡], Kevin M. Greenan[‡], Rosie Wacha^{*}, **Ethan L. Miller**, and Darrell D. E. Long, “Energy-Reliability Tradeoffs in Sensor Networks”, *Proceedings of the 5th Workshop on Embedded Networked Sensors (HotEmNets 2008)*, Charlottesville, VA: ACM, June 2008.
- C75. Mark W. Storer[‡], Kevin M. Greenan[‡], **Ethan L. Miller**, and Kaladhar Voruganti, “Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage”, *Proceedings of the 6th Conference on File and Storage Technologies (FAST '08)*, San Jose, CA: USENIX, February 2008, pages 1–16.
- C74. Jonathan Koren^{*}, Yi Zhang, Sasha Ames[‡], Andrew W. Leung[‡], Carlos Maltzahn, and **Ethan Miller**, “Searching and Navigating Petabyte Scale File Systems Based on Facets”, *Proceedings of the 2nd International Workshop on Petascale Data Storage (PDSW07)*, held in conjunction with SC2007, Reno, NV, November 2007, pages 21–25.
- C73. Andrew Leung[‡], **Ethan L. Miller**, and Stephanie Jones[‡], “Scalable Security for Petascale Parallel File Systems”, *Proceedings of SC2007*, Reno, NV: ACM, November 2007.
- C72. Kevin M. Greenan[‡], **Ethan L. Miller**, Thomas J. E. Schwarz, S. J., and Darrell D. E. Long, “Disaster Recovery Codes: Increasing Reliability with Large-Stripe Error Correction Codes”, *Proceedings of the 3rd International Workshop on Storage Security and Survivability (StorageSS 2007)*, held in conjunction

- with the 14th ACM Conference on Computer and Communications Security (CCS 2007), Alexandria, VA, October 2007.
- C71. Neerja Bhatnagar[‡] and **Ethan L. Miller**, “A Secure and Reliable File System for Sensor Nodes”, *Proceedings of the 3rd International Workshop on Storage Security and Survivability (StorageSS 2007)*, held in conjunction with the 14th ACM Conference on Computer and Communications Security (CCS 2007), Alexandria, VA, October 2007.
- C70. Kevin Greenan[‡] and **Ethan L. Miller**, “PRIMS: Making NVRAM Suitable for Extremely Reliable Storage”, *Proceedings of the 3rd Workshop on Hot Topics in System Dependability (HotDep’07)*, Edinburgh, UK: USENIX, June 2007, 4 pages.
- C69. Mark W. Storer[‡], Kevin Greenan[‡], **Ethan L. Miller**, and Kaladhar Voruganti, “POTSHARDS: Secure Long-Term Storage Without Encryption”, *Proceedings of the 2007 USENIX Annual Technical Conference (USENIX ’07)*, Santa Clara, CA: USENIX, June 2007, pages 143–156.
- C68. Sage A. Weil*, Scott A. Brandt, **Ethan L. Miller**, and Carlos Maltzahn, “CRUSH: Controlled, Scalable And Decentralized Placement Of Replicated Data”, *Proceedings of SC2006*, Tampa, FL: ACM, November 2006, 12 pages (published on CD-ROM).
- C67. Sage A. Weil*, Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Carlos Maltzahn, “Ceph: A Scalable, High-Performance Distributed File System”, *Proceedings of the 7th Conference on Operating Systems Design and Implementation (OSDI ’06)*, Seattle, WA, November 2006, pages 307–320.
- C66. Mark W. Storer[‡], Kevin M. Greenan[‡], and **Ethan L. Miller**, “Long-Term Threats to Secure Archives”, *Proceedings of the 2nd International Workshop on Storage Security and Survivability (StorageSS 2006)*, held in conjunction with the 13th ACM Conference on Computer and Communications Security (CCS 2006), Virginia, October 2006, pages 9–16.
- C65. Andrew Leung[‡] and **Ethan L. Miller**, “Scalable Security for Large, High Performance Storage Systems”, *Proceedings of the 2nd International Workshop on Storage Security and Survivability (StorageSS 2006)*, held in conjunction with the 13th ACM Conference on Computer and Communications Security (CCS 2006), Virginia, October 2006, pages 29–40.
- C64. Kevin M. Greenan[‡] and **Ethan L. Miller**, “Reliability Mechanisms for File Systems Using Non-Volatile Memory as a Metadata Store”, *Proceedings of the 6th ACM & IEEE Conference on Embedded Software (EMSOFT ’06)*, Seoul, Korea, October 2006, pages 178–187.
- C63. Deepavali Bhagwat*, Kristal Pollack*, Darrell D. E. Long, **Ethan L. Miller**, Jehan-François Pâris, and Thomas Schwarz, S. J., “Providing High Reliability in a Minimum Redundancy Archival Storage System”, *Proceedings of the 14th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2006)*, Monterey, CA, September 2006, pages 413–421.
- C62. Thomas J. E. Schwarz and **Ethan L. Miller**, “Store, Forget, and Check: Using Algebraic Signatures to Check Remotely Administered Storage”, *Proceedings of the 26th International Conference on Distributed Computing Systems (ICDCS ’06)*, Lisboa, Portugal, July 2006, 10 pages (published on CD-ROM).
- C61. Sasha Ames[‡], Nikhil Bobb*, Kevin M. Greenan[‡], Owen S. Hofmann*, Mark W. Storer[‡], Carlos Maltzahn, **Ethan L. Miller**, and Scott A. Brandt, “LiFS: An Attribute-Rich File System for Storage Class Memories”, *Proceedings of the 23rd IEEE / 14th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2006)*, College Park, MD, May 2006, pages 63–76.
- C60. Nikhil Bobb*, Damian Eads[‡], Mark W. Storer[‡], Scott A. Brandt, Carlos Maltzahn, and **Ethan L. Miller**, “Graffiti: A Framework for Testing Collaborative Distributed Metadata”, *7th Workshop on Distributed Data and Structures (WDAS 2006)*, Santa Clara, CA, January 2006.
- C59. Mark Storer[‡], Kevin Greenan[‡], **Ethan L. Miller**, and Carlos Maltzahn, “POTSHARDS: Storing Data for the Long-term Without Encryption”, *Proceedings of the 3rd International IEEE Security in Storage Workshop*, San Francisco, CA, December 2005.
- C58. Christopher Olson[‡] and **Ethan L. Miller**, “Secure Capabilities for a Petabyte-Scale Object-Based Distributed File System”, *Proceedings of the International Workshop on Storage Security and Survivability (StorageSS)*, held in conjunction with the 12th ACM Conference on Computer and Communications Security (CCS 2005), Fairfax, VA, November 2005, pages 64–73. Received Best Paper award.

- C57. Qin Xin[‡], Thomas J. E. Schwarz, S. J., and **Ethan L. Miller**, “Disk Infant Mortality in Large Storage Systems”, *Proceedings of the 13th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2005)*, Atlanta, GA, September 2005, pages 125–134.
- C56. Qin Xin[‡], **Ethan L. Miller**, Thomas J. E. Schwarz, S. J., and Darrell D. E. Long, “Impact Of Failure On Interconnection Networks for Large Storage Systems”, *Proceedings of the 22nd IEEE / 13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2005)*, Monterey, CA, April 2005, pages 189–196.
- C55. Alexander Ames[‡], Nikhil Bobb^{*}, Scott A. Brandt, Adam Hiatt[‡], Carlos Maltzahn, **Ethan L. Miller**, Alisa Neeman[‡], and Deepa Tuteja^{*}, “Richer File System Metadata Using Links and Attributes”, *Proceedings of the 22nd IEEE / 13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2005)*, Monterey, CA, April 2005, pages 49–60.
- C54. Sage A. Weil^{*}, Kristal T. Pollack^{*}, Scott A. Brandt, and **Ethan L. Miller**, “Dynamic Metadata Management for Petabyte-scale File Systems”, *Proceedings of SC2004*, Pittsburgh, PA: ACM, November 2004. Nominee, Best Student Paper.
- C53. Thomas J. E. Schwarz, S. J., Qin Xin[‡], **Ethan L. Miller**, Darrell D. E. Long, Andy Hospodor, and Spencer Ng, “Disk Scrubbing in Large Archival Storage Systems”, *Proceedings of the 12th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2004)*, Volendam, Netherlands, October 2004, pages 409–418. Received Best Paper award.
- C52. Nathan K. Edel[‡], Deepa Tuteja^{*}, **Ethan L. Miller**, and Scott A. Brandt, “MRAMFS: A Compressing File System for Non-Volatile RAM”, *Proceedings of the 12th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2004)*, Volendam, Netherlands, October 2004, pages 596–603.
- C51. Thomas J. E. Schwarz, Qin Xin[‡], and **Ethan L. Miller**, “Availability in Global Peer-To-Peer Storage Systems”, *6th Workshop on Distributed Data and Structures (WDAS 2004)*, Lausanne, Switzerland, July 2004.
- C50. Nathan K. Edel[‡], **Ethan L. Miller**, Karl S. Brandt^{*}, and Scott A. Brandt, “Measuring the Compressibility of Metadata and Small Files for Disk/NVRAM Hybrid Storage Systems”, *Proceedings of the 2004 International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS’04)*, San Jose, CA, July 2004.
- C49. Ismail Ari[‡] and **Ethan L. Miller**, “Caching Support for Push-Pull Data Dissemination using Data-snooping Routers”, *Proceedings of the 10th International Conference on Parallel and Distributed Systems (ICPADS)*, Newport Beach, CA: IEEE, July 2004, pages 101–108.
- C48. Qin Xin[‡], **Ethan L. Miller**, and Thomas J. E. Schwarz, S. J. “Evaluation of Distributed Recovery in Large-Scale Storage Systems”, *Proceedings of the 13th IEEE International Symposium on High Performance Distributed Computing (HPDC-13)*, Honolulu, HI: IEEE, June 2004, pages 172–181.
- C47. Sage A. Weil^{*}, Scott A. Brandt, **Ethan L. Miller**, and Kristal T. Pollack^{*}, “Intelligent Metadata Management for a Petabyte-Scale File System”, 2nd Intelligent Storage Workshop, University of Minnesota, May 2004.
- C46. Emilia Rosti and **Ethan L. Miller**, “Security Threats and Responses for Object-Based Storage Devices”, 2nd Intelligent Storage Workshop, University of Minnesota, May 2004 (poster presentation).
- C45. R. J. Honicky[‡] and **Ethan L. Miller**, “Replication Under Scalable Hashing: A Family of Algorithms for Scalable Decentralized Data Distribution”, *Proceedings of the 18th International Parallel and Distributed Processing Symposium (IPDPS 2004)*, Santa Fe, NM: IEEE, April 2004, 10 pages (published on CD-ROM).
- C44. Feng Wang^{*}, Scott A. Brandt, **Ethan L. Miller**, and Darrell D. E. Long, “OBFS: A File System for Object-Based Storage Devices”, *Proceedings of the 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004)*, College Park, MD, April 2004, pages 283–300.
- C43. Andy Hospodor and **Ethan L. Miller**, “Interconnection Architectures for Petabyte-Scale High-Performance Storage Systems”, *Proceedings of the 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004)*, College Park, MD, April 2004, pages 273–281.

- C42. Feng Wang*, Qin Xin[‡], Bo Hong*, Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Tyce T. McLarty, “File System Workload Analysis for Large Scale Scientific Computing Applications”, *Proceedings of the 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004)*, College Park, MD, April 2004, pages 139–152.
- C41. Bo Hong*, Scott A. Brandt, Darrell D. E. Long, **Ethan L. Miller**, Karen A. Glocer*, and Zachary N. J. Peterson*, “Zone-Based Shortest Positioning Time First Scheduling for MEMS-Based Storage Devices”, *Proceedings of the 11th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2003)*, Orlando, FL, October 2003, pages 104–113.
- C40. Ismail Ari[‡], Bo Hong*, **Ethan L. Miller**, Scott A. Brandt and Darrell D. E. Long, “Managing Flash Crowds On The Internet”, *Proceedings of the 11th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2003)*, Orlando, FL, October 2003, pages 246–249.
- C39. R. J. Honicky[‡] and **Ethan L. Miller**, “A Fast Algorithm for Online Placement and Reorganization of Replicated Data”, *Proceedings of the 17th International Parallel and Distributed Processing Symposium*, Nice, France: IEEE, April 2003, 10 pages (published on CD-ROM). Also available as Technical Report UCSC-CRL-02-36.
- C38. Qin Xin[‡], **Ethan L. Miller**, Thomas Schwarz, Scott A. Brandt, Darrell D. E. Long, and Witold Litwin, “Reliability Mechanisms for Very Large Storage Systems”, *Proceedings of the 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2003)*, San Diego, CA: IEEE, April 2003, pages 146–156.
- C37. Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Lan Xue*, “Efficient Metadata Management in Large Distributed File Systems”, *Proceedings of the 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2003)*, San Diego, CA: IEEE, April 2003, pages 290–298.
- C36. Ying Lin*, Scott A. Brandt, Darrell D. E. Long, and **Ethan L. Miller**, “Power Conservation Strategies for MEMS-based Storage Devices”, *Proceedings of the 10th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2002)*, Fort Worth, TX: IEEE, October 2002, pages 53–62.
- C35. Scott A. Banachowski*, Zachary N. J. Peterson*, **Ethan L. Miller**, and Scott A. Brandt, “Intra-file Security for a Distributed File System”, *Proceedings of the 19th IEEE Symposium on Mass Storage Systems and Technologies*, College Park, MD: IEEE, April 2002, pages 153–163.
- C34. Ismail Ari[‡], Ahmed Amer*, **Ethan Miller**, Scott Brandt, and Darrell Long. “Who is more adaptive? ACME: adaptive caching using multiple experts”, *Workshop on Distributed Data and Structures (WDAS 2002)*, Paris, France, March 2002.
- C33. **Ethan L. Miller**, Darrell D. E. Long, William E. Freeman[‡], and Benjamin C. Reed*, “Strong Security for Network-Attached Storage”, *Proceedings of the First Conference on File and Storage Technologies (FAST)*, Monterey, CA: Usenix, January 2002, pages 1–13.
- C32. **Ethan L. Miller**, Scott A. Brandt and Darrell D. E. Long, “HeRMES: High-Performance Reliable MRAM-Enabled Storage”, *Proceedings of the 8th IEEE Workshop on Hot Topics in Operating Systems (HotOS-VIII)*, Elmau, Germany: IEEE, May 2001, pages 83–87.
- C31. Lee Butler*, Travis Atkison*, and **Ethan Miller**, “Comparing CPU Performance Between and Within Processor Families”, *Proceedings of the 25th Annual International Conference on Computer Measurement and Performance (CMG 2000)*, Orlando, FL, December 2000, pages 421–430.
- C30. **Ethan Miller** and Jon Squire*, “esim: A Structural Design Language and Simulator for Computer Architecture Education”, *2000 Workshop on Computer Architecture Education (WCAE 2000)*, Vancouver, Canada: ACM & IEEE, June 2000, pages 42–48.
- C29. William Freeman[‡] and **Ethan Miller**, “Design for A Decentralized Security System For Network Attached Storage”, *Proceedings of the 8th Goddard Conference on Mass Storage Systems and Technologies / 17th IEEE Symposium on Mass Storage Systems*, College Park, MD, March 2000, pages 361–373.
- C28. Timothy Gibson[‡] and **Ethan Miller**, “An Improved Long-Term File Usage Prediction Algorithm”, *Proceedings of the 25th Annual International Conference on Computer Measurement and Performance (CMG ’99)*, Reno, NV: CMG, December 1999, pages 639–648.

- C27. William E. Freeman[‡] and **Ethan L. Miller**, “An Experimental Analysis of Cryptographic Overhead in Performance-Critical Systems”, *Proceedings of the 7th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS '99)*, College Park, MD: IEEE, October 1999, pages 348–357.
- C26. **Ethan Miller**, Dan Shen[‡], Junli Liu[‡], Charles Nicholas, and Ting Chen[‡], “Techniques for Gigabyte-Scale N-gram Based Information Retrieval on Personal Computers”, *Proceedings of the 1999 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '99)*, Las Vegas, NV, June 1999, pages 1410–1416.
- C25. Michael Shapiro[‡] and **Ethan Miller**, “Managing Databases with Binary Large Objects”, *Proceedings of the 16th IEEE Mass Storage System Symposium*, San Diego, CA: IEEE, March 1999, pages 185–193.
- C24. Timothy Gibson[‡], **Ethan L. Miller** and Darrell D. E. Long. “Long-term File Activity and Inter-reference Patterns”, *Proceedings of the 24th Annual International Conference on Computer Measurement and Performance (CMG '98)*, Anaheim, CA: CMG, December 1998, pages 976–987.
- C23. Kennedy Akala[‡], **Ethan Miller**, and Jeff Hollingsworth, “Using Content-Derived Names for Package Management in Tcl”, *Proceedings of the 6th Annual Tcl/Tk Conference*, San Diego, CA: Usenix, September 1998, pages 171–179.
- C22. Theodore Johnson and **Ethan Miller**, “Performance Measurements of Tertiary Storage Devices”, *Proceedings of the 1998 Conference on Very Large Databases (VLDB '98)*, New York, NY: VLDB Foundation, August 1998, pages 50–61.
- C21. Jem Y. Fan, Xiangjun Zhao, J. P. Zhang, Fow-Sen Choa, Yanjie Chai, Jye-Hong Chen, **Ethan Miller**, Howard Motteler, Pao-Lo Liu, Tawee Tanbun-Ek, Patrick Wisk, Won-Tien Tsang, George J. Zyzdik, and Charles A. Burrus, “Wavelength-division-multiplexed (WDM) data block switching for parallel computing and interconnect”, *SPIE International Conference on Applications of Photonic Technology*, Ottawa, Canada: SPIE, July 1998, vol. 3491, pages 634–638.
- C20. Steven Gribble*, Gurmeet Singh Manku*, Drew Roselli*, Eric Brewer, Timothy Gibson[‡], and **Ethan Miller**, “Self-Similarity in File Systems”, *Proceedings of the SIGMETRICS '98 / PERFORMANCE '98 Joint International Conference on Measurement and Modeling of Computer Systems*, Madison, WI: ACM, June 1998, pages 141–150.
- C19. Timothy J. Gibson[‡] and **Ethan Miller**, “Long-Term File Activity Patterns in a UNIX Workstation Environment”, *Proceedings of the 6th Goddard Conference on Mass Storage Systems and Technologies / 15th IEEE Symposium on Mass Storage Systems*, College Park, MD: IEEE, March 1998, pages 355–372.
- C18. Theodore Johnson and **Ethan Miller**, “Benchmarking Tape System Performance”, *Proceedings of the 6th Goddard Conference on Mass Storage Systems and Technologies / 15th IEEE Symposium on Mass Storage Systems*, College Park, MD: IEEE, March 1998, pages 95–112.
- C17. R. Scott Cost*, Jeegar Lakhani*, Ian Soboroff*, Tim Finin, **Ethan Miller**, and Charles Nicholas, “TKQML: A Scripting Tool for Building Agents”, *Proceedings of the 1997 Conference on Agent Theories and Agent Languages (ATAL97)*, Newport, RI: AAAI, July 1997, pages 339–343.
- C16. R. Scott Cost*, Jeegar Lakhani*, Ian Soboroff*. Tim Finin, **Ethan Miller**, and Charles Nicholas, “Agent Development Support for Tcl”, *5th Annual Tcl/Tk Workshop '97*, Boston, MA: Usenix, July 1997, pages 177–178.
- C15. Jeff Hollingsworth and **Ethan Miller**, “Using Content-Derived Names for Configuration Management”, *Proceedings of the 1997 Symposium on Software Reusability (SSR '97)*, Boston, MA: IEEE, May 1997, pages 104–109.
- C14. David Ebert, Chris Shaw, Amen Zwa*, **Ethan Miller**, and D. A. Roberts, “Interactive Volumetric Information Visualization for Document Corpus Management”, *Proceedings of Graphics Interface*, Kelowna, BC, Canada: Canadian Human-Computer Communications Society, May 1997, pages 121–128.
- C13. Timothy Gibson[‡] and **Ethan Miller**, “The Case for Personal Computers as Workstations”, *Proceedings of the 22nd Annual International Conference on Computer Measurement and Performance (CMG '96)*, San Diego, CA: CMG, December 1996, pages 644–652.
- C12. Amen Zwa[‡], David Ebert, and **Ethan Miller**, “Multiresolution Document Analysis with Wavelets”, *Proceedings of the 1996 Conference on Information and Knowledge Management, Workshop on New Paradigms in Information Visualization and Manipulation*, Rockville, MD: ACM, December 1996, pages 50–53.

- C11. David Ebert, Chris Shaw, Amen Zwa[‡], and **Ethan Miller**, “Minimally-immersive Interactive Volumetric Information Visualization”, *Proceedings of IEEE Information Visualization '96*, San Francisco, CA: IEEE, October 1996, pages 66–68.
- C10. **Ethan Miller**, “Towards Scalable Benchmarks for Mass Storage Systems”, *5th NASA Goddard Space Flight Center Conference on Mass Storage Systems and Technologies*, College Park, MD: IEEE & NASA, September 1996, pages 515–528.
- C9. **Ethan Miller** and Randy Katz, “RAMA: Easy Access to a High-Bandwidth Massively Parallel File System”, *Proceedings of the Winter 1995 USENIX Conference*, New Orleans, LA: Usenix, January 1995, pages 59–70.
- C8. Ann Drapeau, Peter Chen, John Hartman, Edward Lee, **Ethan Miller**, Ken Shirriff, Srinu Seshan, Randy Katz, Garth Gibson, and David Patterson, “RAID-II: A Scalable Storage Architecture for High-Bandwidth Network File Service”, *Proceedings of the 21st International Symposium on Computer Architecture*, Chicago, IL: ACM, April 1994, pages 234–244.
- C7. Peter Chen, Edward Lee, Ann Drapeau, Ken Lutz, **Ethan Miller**, Srinu Seshan, Ken Shirriff, David Patterson, and Randy Katz, “Performance and Design Evaluation of the RAID-II Storage Server”, *Proceedings of the International Parallel Processing Symposium Workshop on I/O in Parallel Computer Systems*, Newport Beach, CA: IEEE, April 1993, pages 110–120.
- C6. **Ethan Miller** and Randy Katz, “RAMA: A File System for Massively Parallel Computers”, *Digest of Papers, 12th IEEE Symposium on Mass Storage Systems*, Monterey, CA: IEEE, April 1993, pages 163–168.
- C5. Randy Katz, Peter Chen, Ann Drapeau, Edward Lee, **Ethan Miller**, Srinu Seshan, and David Patterson, “RAID-II: Design and Implementation of a Large Scale Disk Array Controller”, *Proceedings of the VLSI System Design Conference*, Seattle, WA: IEEE, March 1993.
- C4. **Ethan Miller** and Randy Katz, “An Analysis of File Migration in a UNIX Supercomputing Environment”, *Proceedings of the Winter 1993 USENIX Conference*, San Diego, CA: Usenix, January 1993, pages 421–433.
- C3. Randy H. Katz, David A. Patterson, Ann Chervenak-Drapeau, Joel Fine, and **Ethan Miller**, “An Approach to Cost-Effective Terabyte Memory Systems”, *Digest of Papers, Compcon Spring '92, 37th IEEE Computer Society International Conference*, San Francisco, CA: IEEE, February 1992, pages 395–400.
- C2. **Ethan Miller** and Randy Katz, “Input/Output Behavior of Supercomputing Applications”, *Proceedings of Supercomputing '91*, Albuquerque, NM: IEEE, November 1991, pages 567–576.
- C1. **Ethan Miller** and Randy Katz, “Analyzing the I/O Behavior of Supercomputing Applications”, *Digest of Papers, 11th IEEE Symposium on Mass Storage Systems*, Monterey, CA: IEEE, October 1991, pages 51–55.

Invited Conference Papers

- IC3. **Ethan L. Miller** and Darrell D. E. Long, “Including Experimental Methods in Operating Systems Courses”, ACM Workshop on Experimental Computer Science (ExpCS 2007), San Diego, CA, June 2007.
- IC2. **Ethan L. Miller**, “Dealing with Long-Lived Data in High Performance Storage Systems”, *Storage on the Lunatic Fringe: Beyond Peta-Scale Storage Systems*, workshop at Supercomputing 2003, Phoenix, AZ, November 2003.
- IC1. **Ethan Miller**, Darrell Long, William Freeman[‡], and Benjamin Reed*, “Strong Security for Distributed File Systems”, *Proceedings of the 20th IEEE International Performance, Computing, and Communications Conference (IPCCC 2001)*, Phoenix, AZ: IEEE, April 2001, pages 34–40.

Patents

- PA142. **Ethan L. Miller**, John Colgrove. Dynamically resizable structures for approximate membership queries. US Patent 11,341,136, issued May 24, 2022. Assignee: Pure Storage.
- PA141. Jonas R. Irwin, **Ethan L. Miller**, John D. Davis. Storage efficiency of encrypted host system data. US Patent 11,307,998, issued Apr. 19, 2022. Assignee: Pure Storage.

- PA140. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Responding to variable response time behavior in a storage environment. US Patent 11,307,772, issued Apr. 19, 2022. Assignee: Pure Storage.
- PA139. **Ethan L. Miller**, John Colgrove. Placing data within a storage device. US Patent 11,294,588, issued Apr. 5, 2022. Assignee: Pure Storage.
- PA138. John Colgrove, Craig Harmer, John Hayes, Bo Hong, **Ethan Miller**, Feng Wang. Intelligently sizing high latency I/O requests in a storage environment. US Patent 11,275,509, issued Mar. 15, 2022. Assignee: Pure Storage.
- PA137. **Ethan L. Miller** Dynamically resizable structures for approximate membership queries. US Patent 11,269,884, issued Mar. 8, 2022. Assignee: Pure Storage.
- PA136. John Colgrove, **Ethan L. Miller**, Neil Amar Vachharajani, Feng Wang. Memory efficient searching. US Patent 11,249,999, issued Feb. 15, 2022. Assignee: Pure Storage.
- PA135. **Ethan Miller**, John Colgrove. Preparing data for deduplication. US Patent 11,221,778, issued Jan. 11, 2022. Assignee: Pure Storage.
- PA134. **Ethan Miller**, John Colgrove, John Hayes. Exporting an address space in a thin-provisioned storage device. US Patent 11,169,745, issued Nov. 9, 2021. Assignee: Pure Storage.
- PA133. Andrew Bernat, **Ethan Miller**. Data re-encryption in a storage system. US Patent 11,146,396, issued Oct. 12, 2021. Assignee: Pure Storage.
- PA132. Andrew Bernat, **Ethan Miller**. Quorum-aware secret sharing. US Patent 11,128,448, issued Sep. 21, 2021. Assignee: Pure Storage.
- PA131. Robert Lee, Christopher Lumb, **Ethan L. Miller**, Igor Ostrovsky. Reducing data distribution inefficiencies. US Patent 11,119,656, issued Sep. 14, 2021. Assignee: Pure Storage.
- PA130. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao, Grigori Inozemtsev. Copying data without accessing the data. US Patent 11,099,769, issued Aug. 24, 2021. Assignee: Pure Storage.
- PA129. John Colgrove, **Ethan L. Miller**. Intelligent error correction in a storage device. US Patent 11,095,315, issued Aug. 17, 2021. Assignee: Pure Storage.
- PA128. **Ethan L. Miller**, Benjamin Borowiec, Steve Hodgson. Cloud-based disaster recovery of a storage system. US Patent 11,061,786, issued July 13, 2021. Assignee: Pure Storage.
- PA127. **Ethan Miller**, Lydia Do, John Colgrove. Migrating data between volumes using virtual copy operation. US Patent 11,036,393, issued June 15, 2021. Assignee: Pure Storage.
- PA126. Andrew Bernat, Timothy Brennan, **Ethan Miller**, John Colgrove. Data protection in a storage system. US Patent 11,032,259, issued June 8, 2021. Assignee: Pure Storage.
- PA125. John Colgrove, Lydia Do, **Ethan Miller**, Terence Noonan. Utilizing redundant resources in a storage system. US Patent 10,996,859, issued May 4, 2021. Assignee: Pure Storage.
- PA124. John Colgrove, Peter E. Kirkpatrick, Robert Lee, **Ethan L. Miller**. Separate encryption for a solid-state drive. US Patent 10,979,223, issued Apr. 13, 2021. Assignee: Pure Storage.
- PA123. **Ethan Miller**, Robert Lee, Par Botes, Ronald Karr. Providing for increased flexibility for large scale parity. US Patent 10,929,226, issued Feb. 23, 2021. Assignee: Pure Storage.
- PA122. **Ethan L. Miller**, Andrew R. Bernat. Data encryption in a distributed system. US Patent 10,887,099, issued Jan. 5, 2021. Assignee: Pure Storage.
- PA121. John Colgrove, John Hayes, **Ethan L. Miller**. Protecting data in a storage system. US Patent 10,887,086, issued Jan. 5, 2021. Assignee: Pure Storage.
- PA120. John Hayes, **Ethan Miller**, John Colgrove. Key management in a storage device. US Patent 10,846,275, issued Nov. 24, 2020. Assignee: Pure Storage.
- PA119. Christopher Golden, John Colgrove, **Ethan L. Miller**, Malcolm Sharpe, Steve Hodgson. Multi-source data replication. US Patent 10,838,640, issued Nov. 17, 2020. Assignee: Pure Storage.
- PA118. Benjamin P. Borowiec, Jimmy T. Hu, **Ethan L. Miller**, Terence W. Noonan, Constantine P. Saputakis, Neil A. Vachharajani, Daquan Zuo. Hybrid cloud-based authentication for flash storage array access. US Patent 10,834,086, issued Nov. 10, 2020. Assignee: Pure Storage.
- PA117. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Generating protection data in a storage system. US Patent 10,817,375, issued Oct. 27, 2020. Assignee: Pure Storage.

- PA116. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Decreasing parity overhead in a storage system. US Patent 10,810,083, issued Oct. 20, 2020. Assignee: Pure Storage.
- PA115. **Ethan L. Miller**, Marco Sanvido. Feature-based deduplication. US Patent 10,789,211, issued Sep. 29, 2020. Assignee: Pure Storage.
- PA114. Ronald S. Karr, **Ethan L. Miller**. Deduplication of data in a storage device. US Patent 10,761,759, issued Sep. 1, 2020. Assignee: Pure Storage.
- PA113. **Ethan L. Miller**, Ronald Karr. Disaster recovery for high-bandwidth distributed archives. US Patent 10,733,053, issued Aug. 4, 2020. Assignee: Pure Storage.
- PA112. Par Botes, John Hayes, **Ethan Miller**. Secure data replication in a storage grid. US Patent 10,691,812, issued June 23, 2020. Assignee: Pure Storage.
- PA111. John Colgrove, Rajesh Kanungo, Ronald Karr, **Ethan L. Miller**. Per-tenant deduplication for shared storage. US Patent 10,678,754, issued June 9, 2020. Assignee: Pure Storage.
- PA110. Andrew R. Bernat, Ganesh Ramanarayanan, Malcom Sharpe, Steve Hodgson, **Ethan Miller**, Alan Driscoll, Christopher Golden, John Colgrove. Data replication within a flash storage array. US Patent 10,656,864, issued May 19, 2020. Assignee: Pure Storage.
- PA109. John Colgrove, Lydia Do, **Ethan Miller**. Upgrading a write buffer in a storage system that includes a plurality of write storage devices and a plurality of write buffer devices. US Patent 10,642,524, issued May 5, 2020. Assignee: Pure Storage.
- PA108. Andrew Bernat, Damian Yurzola, Timothy Brennan, **Ethan Miller**, John Colgrove. Secret Sharing Data Protection in a Storage System. US Patent 10,623,386, issued Apr. 14, 2020. Assignee: Pure Storage.
- PA107. Ronald Karr, **Ethan Miller**, John Colgrove. Improving the accuracy of in-line data deduplication. US Patent 10,620,864, issued Apr. 14, 2020. Assignee: Pure Storage.
- PA106. John Colgrove, Mark L. McAuliffe, **Ethan L. Miller**, Naveen Neelakantam, Marco Sanvido, Neil A. Vachharajani, Taher Vohra. Utilizing an address-independent, non-repeating encryption key to encrypt data. US Patent 10,607,034, issued March 31, 2020. Assignee: Pure Storage.
- PA105. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao, Grigori Inozemtsev. Buffering copy requests in a storage system. US Patent 10,585,617, issued March 10, 2020. Assignee: Pure Storage.
- PA104. Ronald Karr, **Ethan Miller**, Constantine Sapuntzakis. Solid state drives with multiple types of addressable memory. US Patent 10,522,090, issued Feb. 4, 2020. Assignee: Pure Storage.
- PA103. Benjamin Borowiec, **Ethan L. Miller**, Steve Hodgson, Andrew R. Bernat, Ganesh Ramanarayanan, Malcolm Sharpe, Alan S. Driscoll. Replication to the cloud. US Patent 10,545,987, issued Jan. 28, 2020. Assignee: Pure Storage.
- PA102. John Colgrove, Lydia Do, **Ethan Miller**, Terence Noonan. Providing high availability to client-specific applications executing in a storage system. US Patent 10,545,676, issued Jan. 28, 2020. Assignee: Pure Storage.
- PA101. John Colgrove, Joseph S. Hasbani, John Martin Hayes, **Ethan L. Miller**, Cary A. Sandvig. Data object attribute based event detection in a storage system. US Patent 10,540,343, issued Jan. 21, 2020. Assignee: Pure Storage.
- PA100. **Ethan Miller**, John Colgrove, John Hayes, Cary Sandvig. Intelligently mapping virtual blocks to physical blocks in a storage system. US Patent 10,521,120, issued Dec. 31, 2019. Assignee: Pure Storage.
- PA99. **Ethan L. Miller**, Andrew R. Bernat. Periodically re-encrypting user data stored on a storage device. US Patent 10,489,307, issued Nov. 26, 2019. Assignee: Pure Storage.
- PA98. Feng Wang, **Ethan L. Miller**, Wei Zhang, Ronald Karr, Cary A. Sandvig. Generating and optimizing summary index levels in a deduplication storage system. US Patent 10,452,297, issued Oct. 22, 2019. Assignee: Pure Storage.
- PA97. John Colgrove, John Hayes, Bo Hong, **Ethan L. Miller**. Dynamically adjusting an amount of protection data stored in a storage storage system. US Patent 10,452,289, issued Oct. 22, 2019. Assignee: Pure Storage.
- PA96. **Ethan L. Miller**, John Colgrove. Error correction processing in a storage system. US Patent 10,432,233, issued Oct. 1, 2019. Assignee: Pure Storage.

- PA95. John Davis, Jonas R. Irwin, **Ethan L. Miller**. Data reduction with end-to-end security. US Patent 10,387,661, issued Aug. 20, 2019. Assignee: Pure Storage.
- PA94. **Ethan Miller**, John Colgrove, John Hayes. Thin provisioning in a storage device. US Patent 10,365,858, issued July 30, 2019. Assignee: Pure Storage.
- PA93. Robert Lee, Christopher Lumb, **Ethan L. Miller**, Igor Ostrovsky. Deduplication aware scalable content placement. US Patent 10,359,942, issued July 23, 2019. Assignee: Pure Storage.
- PA92. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Simultaneously servicing high latency operations in a storage system. US Patent 10,353,630, issued July 16, 2019. Assignee: Pure Storage.
- PA91. John Colgrove, Ronald Karr, **Ethan Miller**. Resolving fingerprint collisions in flash storage system. US Patent 10,303,390, issued May 28, 2019. Assignee: Pure Storage.
- PA90. **Ethan Miller**, John Colgrove, John Hayes. Encrypting data in a storage system using a plurality of encryption keys. US Patent 10,284,367, issued May 7, 2019. Assignee: Pure Storage.
- PA89. John Colgrove, **Ethan Miller**. Dynamic error processing in a storage device. US Patent 10,284,232, issued May 7, 2019. Assignee: Pure Storage.
- PA88. **Ethan Miller**, John Colgrove, John Hayes. Data protection in a storage system using external secrets. US Patent 10,263,770, issued April 16, 2019. Assignee: Pure Storage.
- PA87. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao. Restoring snapshots in a storage system. US Patent 10,235,093, issued March 19, 2019. Assignee: Pure Storage.
- PA86. **Ethan L. Miller**, Benjamin Borowiec, Steve Hodgson. Dataset replication in a cloud computing environment. US Patent 10,235,065, issued March 19, 2019. Assignee: Pure Storage.
- PA85. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Maintaining a target number of storage devices for variable i/o response times in a storage system. US Patent 10,228,865, issued March 12, 2019. Assignee: Pure Storage.
- PA84. Andrew R. Bernat, **Ethan L. Miller**. Resharing of a split secret. US Patent 10,211,983, issued Feb. 19, 2019. Assignee: Pure Storage.
- PA83. John Colgrove, **Ethan Miller**. Placing Data within a storage device of a flash array. US Patent 10,198,194, issued Feb. 5, 2019. Assignee: Pure Storage.
- PA82. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Inter-device and intra-device protection data. US Patent 10,180,879, issued Jan. 15, 2019. Assignee: Pure Storage.
- PA81. **Ethan Miller**, Lydia Do, John Colgrove. Migrating data between volumes using virtual copy operation. US Patent 10,162,523, issued Dec. 25, 2018. Assignee: Pure Storage.
- PA80. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer, Reducing a number of storage devices that are exhibiting variable I/O response times. US Patent 10,156,998, issued Dec. 18, 2018. Assignee: Pure Storage.
- PA79. John Colgrove, Ronald Karr, **Ethan L. Miller**. Selecting a deduplication process based on a difference between performance metrics. US Patent 10,133,503, issued Nov. 20, 2018. Assignee: Pure Storage.
- PA78. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Adjusting a number of storage devices in a storage system that may be utilized to simultaneously serve high latency operations. US Patent 10,126,982, issued Nov. 13, 2018. Assignee: Pure Storage.
- PA77. Christopher Golden, John Colgrove, **Ethan L. Miller**, Malcolm Sharpe, Steve Hodgson. Optimizing storage allocation in a storage system. US Patent 10,114,574, issued Oct. 30, 2018. Assignee: Pure Storage.
- PA76. **Ethan Miller**, John Colgrove, John Hayes, Cary Sandvig. Identifying fractal regions across multiple storage devices. US Patent 10,089,010, issued Oct. 2, 2018. Assignee: Pure Storage.
- PA75. John Colgrove, Lydia Do, **Ethan Miller**. Preparing for cache upgrade in a storage array that includes a plurality of storage devices and a plurality of write buffer devices. US Patent 10,078,469, issued Sep. 18, 2018. Assignee: Pure Storage.
- PA74. John Colgrove, Joseph S. Hasbani, John Hayes, **Ethan Miller**, Cary Sandvig. Method for maintaining multiple fingerprint tables in a deduplicating storage system. US Patent 10,061,798, issued Aug. 28, 2018. Assignee: Pure Storage.

- PA73. John Colgrove, Mark L. McAuliffe, **Ethan L. Miller**, Naveen Neelakantam, Marco Sanvido, Neil A. Vachharajani, Taher Vohra. Generating a Unique Encryption Key. US Patent 10,037,440, issued July 31, 2018. Assignee: Pure Storage.
- PA72. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao. Restoring a Volume in a Storage System. US Patent 10,013,317, issued July 3, 2018. Assignee: Pure Storage.
- PA71. Ronald Karr, **Ethan L. Miller**, Cary A. Sandvig, Feng Wang, Wei Zhang. Generating and Optimizing Summary Index Levels in a Deduplication Storage System. US Patent 9,983,822, issued May 29, 2018. Assignee: Pure Storage.
- PA70. John Colgrove, Ronald Karr, **Ethan L. Miller**, Vinay K. Perneti, Cary A. Sandvig, Feng Wang, Wei Zhang. Memory use and eviction in a deduplication storage system. US Patent 9,940,060, issued April 10, 2018. Assignee: Pure Storage.
- PA69. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao. "Deduplication of regions with a storage system". US Patent 9,891,858, issued Feb. 13, 2018. Assignee: Pure Storage.
- PA68. Benjamin P. Borowiec, Jimmy T. Hu, **Ethan L. Miller**, Terence W. Noonan, Constantine P. Sapuntzakis, Neil A. Vachharajani, Daquan Zuo. Delivering authorization and authentication for a user of a storage array from a cloud. US Patent 9,882,913, issued Jan. 30, 2018. Assignee: Pure Storage.
- PA67. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao, Grigori Inozemtsev. "Processing copy offload requests in a storage system". US Patent 9,880,779, issued Jan. 30, 2018. Assignee: Pure Storage.
- PA66. John Colgrove, Lydia Do, **Ethan Miller**. Data migration in a storage array that includes a plurality of storage devices. US Patent 9,817,603, issued Nov. 14, 2017. Assignee: Pure Storage.
- PA65. John Hayes, Par Botes, **Ethan Miller**. Secure data replication in a storage grid. US Patent 9,811,677, issued Nov. 7, 2017. Assignee: Pure Storage.
- PA64. John Colgrove, John Hayes, **Ethan Miller**, Joseph S. Hasbani, Cary Sandvig. Utilizing multiple fingerprint tables in a deduplicating storage system. US Patent 9,811,551, issued Nov. 7, 2017. Assignee: Pure Storage.
- PA63. John Davis, **Ethan Miller**, Brian Gold, John Colgrove, Peter Vajgel, John Hayes, Alex Ho. Data placement based on data retention in a tiered storage device system. US Patent 9,811,457, issued Nov. 7, 2017. Assignee: Pure Storage.
- PA62. John Colgrove, Lydia Do, Ethan Miller, Terence Noonan. Deploying client-specific applications in a storage system utilizing redundant system resources. US Patent 9,811,264, issued Nov. 7, 2017. Assignee: Pure Storage.
- PA61. **Ethan Miller**, John Colgrove, John Hayes, Cary Sandvig. Distributing data blocks across a plurality of storage devices. US Patent 9,792,045, issued Oct 17, 2017. Assignee: Pure Storage.
- PA60. John Colgrove, Mark L. McAuliffe, **Ethan L. Miller**, Naveen Neelakantam, Marco Sanvido, Neil A. Vachharajani, Taher Vohra. Utilizing a non-repeating identifier to encrypt data. US Patent 9,779,268, issued Oct. 3, 2017. Assignee: Pure Storage.
- PA59. Andrew R. Bernat, **Ethan L. Miller**. Resharing of a split secret. US Patent 9,768,953, issued Sep. 19, 2017. Assignee: Pure Storage.
- PA58. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao, Grigori Inozemtsev. Performing copies in a storage system. US Patent 9,760,313, issued Sep. 12, 2017. Assignee: Pure Storage.
- PA57. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Proactively correcting behavior that may affect I/O performance in a non-volatile semiconductor storage device. US Patent 9,684,460, issued June 20, 2017. Assignee: Pure Storage.
- PA56. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao. Snapshots in a storage system. US Patent 9,646,039, issued May 9, 2017. Assignee: Pure Storage.
- PA55. John Davis, **Ethan Miller**, Brian Gold, John Colgrove, Peter Vajgel, John Hayes, Alex Ho. Data placement based on data properties in a tiered storage device system. US Patent 9,612,953, issued April 4, 2017. Assignee: Pure Storage.

- PA54. Ronald S. Karr, **Ethan L. Miller**. Preventing duplicate entries of identical data in a storage device. US Patent 9,594,678, issued March 14, 2017. Assignee: Pure Storage.
- PA53. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Adaptive RAID for an SSD environment. US Patent 9,594,633, issued March 14, 2017. Assignee: Pure Storage.
- PA52. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao, Deduplication of volume regions. US Patent 9,589,008, issued March 7, 2017. Assignee: Pure Storage.
- PA51. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of Reactive I/O in a storage environment. US Patent 9,588,699, issued March 7, 2017. Assignee: Pure Storage.
- PA50. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of I/O in an SSD environment. US Patent 9,569,116, issued Feb. 14, 2017. Assignee: Pure Storage.
- PA49. **Ethan Miller**, Benjamin Borowiec, Steve Hodgson. Cloud alert to replica. US Patent 9,552,248, issued Jan. 24, 2017. Assignee: Pure Storage.
- PA48. **Ethan Miller**, John Colgrove, John Hayes. Multi-drive cooperation to generate an encryption key, US Patent 9,548,972, issued Jan. 17, 2017. Assignee: Pure Storage.
- PA47. John Colgrove, **Ethan Miller**, John Hayes. Storage array password management. US Patent 9,516,016, issued Dec. 6, 2016. Assignee: Pure Storage.
- PA46. John Colgrove, Lydia Do, **Ethan Miller**. Migrating data in a storage array that includes a plurality of storage devices and a plurality of write buffer devices. US Patent 9,507,532, issued Nov. 29, 2016. Assignee: Pure Storage.
- PA45. Christopher Golden, John Colgrove, **Ethan L. Miller**, Malcolm Sharpe, Steve Hodgson. Utilizing un-mapped and unknown states in a replicated storage system. US Patent 9,489,132, issued Nov. 8, 2016. Assignee: Pure Storage.
- PA44. John Colgrove, John Hayes, **Ethan Miller**, Feng Wang. Logical sector mapping in a flash storage array. US Patent 9,454,477, issued Sep. 27, 2016. Assignee: Pure Storage.
- PA43. John Colgrove, John Hayes, **Ethan Miller**, Feng Wang. Logical sector mapping in a flash storage array. US Patent 9,454,476, issued Sep. 27, 2016. Assignee: Pure Storage.
- PA42. Benjamin P. Borowiec, Jimmy T. Hu, **Ethan L. Miller**, Terence W. Noonan, Constantine P. Sapuntzakis, Neil A. Vachharajani, Daquan, Zuo. Storage array access control from cloud-based user authorization and authentication. US Patent 9,444,822, issued Sep. 13, 2016. Assignee: Pure Storage.
- PA41. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao. Safety for volume operations. US Patent 9,436,720, issued Sep. 6, 2016. Assignee: Pure Storage.
- PA40. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of reconstructive I/O read operations in a storage environment. US Patent 9,436,396, issued Sep. 6, 2016. Assignee: Pure Storage.
- PA39. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of I/O writes in a storage environment. US Patent 9,423,967, issued August 23, 2016. Assignee: Pure Storage.
- PA38. **Ethan Miller**, John Colgrove, John Hayes, Cary Sandvig. Fractal layout of data blocks across multiple devices. US Patent 9,405,486, issued August 2, 2016. Assignee: Pure Storage.
- PA37. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Grigori Inozemtsev. Performing copies in a storage system. US Patent 9,361,035, issued June 7, 2016. Assignee: Pure Storage.
- PA36. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Distributed multi-level protection in a raid array based storage system. US Patent 9,348,696, issued May 24, 2016. Assignee: Pure Storage.
- PA35. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of reactive I/O operations in an storage environment. US Patent 9,304,694, issued April 5, 2016. Assignee: Pure Storage.
- PA34. Benjamin P. Borowiec, Jimmy T. Hu, **Ethan L. Miller**, Terence W. Noonan, Constantine P. Sapuntzakis, Neil A. Vachharajani, Daquan Zuo. Providing authorization and authentication in a cloud for a user of a storage array. US Patent 9,300,600, issued March 29, 2016. Assignee: Pure Storage.
- PA33. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of I/O in an SSD environment. US Patent 9,298,376, issued March 29, 2016. Assignee: Pure Storage, Inc.

- PA32. John Colgrove, John Hayes, **Ethan Miller**, Cary Sandvig, Joseph S. Hasbani, Feng Wang. Garbage collection in a storage system. US Patent 9,251,066 issued Feb. 2, 2016. Assignee: Pure Storage.
- PA31. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Offset protection data in a RAID array. US Patent 9,244,769, issued January 26, 2016. Assignee: Pure Storage.
- PA30. John Colgrove, John Hayes, **Ethan Miller**, Cary Sandvig. Mapping in a storage system. US Patent 9,239,688, issued January 19, 2016. Assignee: Pure Storage.
- PA29. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Reconstruct reads in a RAID array with dynamic geometries. US Patent 9,229,808, issued January 5, 2016. Assignee: Pure Storage.
- PA28. **Ethan Miller**, John Colgrove, John Hayes. Efficient techniques for aligned fixed-length compression. US Patent 9,077,368, issued July 7, 2015. Assignee: Pure Storage.
- PA27. John Colgrove, John Hayes, **Ethan Miller**, Joseph S. Hasbani, Cary Sandvig. Method for maintaining multiple fingerprint tables in a deduplicating storage system. US Patent 9,069,786, issued June 30, 2015. Assignee: Pure Storage.
- PA26. John Colgrove, **Ethan Miller**, John Hayes, Cary Sandvig, Christopher Golden, Jianting Cao, Grigori Inozemtsev. Performing copies in a storage system. US Patent 9,063,967, issued June 23, 2015. Assignee: Pure Storage.
- PA25. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Intra-device data protection in a RAID array. US Patent 9,058,116, issued June 16, 2015. Assignee: Pure Storage.
- PA24. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of I/O writes in a storage environment. US Patent 9,037,827, issued May 19, 2015. Assignee: Pure Storage.
- PA23. John Colgrove, John Hayes, **Ethan Miller**, Variable length encoding in a storage system, US Patent 8,954,710, issued February 10, 2015. Assignee: Pure Storage.
- PA22. John Colgrove, John Hayes, **Ethan Miller**, Joseph S. Hasbani, Cary Sandvig. Method for removing duplicate data from a storage array. US Patent 8,930,307, issued Jan. 6, 2015. Assignee: Pure Storage.
- PA21. John Colgrove, John Hayes, **Ethan Miller**, Cary Sandvig, Joseph S. Hasbani, Feng Wang. Garbage collection in a storage system. US Patent 8,886,691 issued Nov. 11, 2014. Assignee: Pure Storage.
- PA20. John Davis, **Ethan Miller**, Brian Gold, John Colgrove, Peter Vajgel, John Hayes, Alex Ho. Data placement based on data properties in a tiered storage device system. US Patent 8,874,835, issued Oct. 28, 2014. Assignee: Pure Storage.
- PA19. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer, Scheduling of reconstructive I/O read operations in a storage environment. US Patent 8,862,820, issued October 14, 2014. Assignee: Pure Storage.
- PA18. John Colgrove, John Hayes, **Ethan Miller**, Feng Wang, Logical sector mapping in a flash storage array. US Patent 8,856,489, issued Oct. 7, 2014. Assignee: Pure Storage.
- PA17. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**, Intra-device data protection in a RAID array. US Patent 8,832,373, issued Sep. 9, 2014. Assignee: Pure Storage.
- PA16. John Colgrove, John Hayes, **Ethan Miller**, Cary Sandvig, Mapping in a storage system. US Patent 8,806,160, issued August 12, 2014. Assignee: Pure Storage.
- PA15. John Colgrove, John Hayes, **Ethan Miller**, Variable length encoding in a storage system, US Patent 8,793,467, issued July 29, 2014. Assignee: Pure Storage.
- PA14. John Colgrove, John Hayes, **Ethan Miller**, Feng Wang, Logical sector mapping in a flash storage array, US Patent 8,788,788, issued July 22, 2014. Assignee: Pure Storage.
- PA13. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**, Adaptive RAID for an SSD environment. US Patent 8,775,868, issued July 8, 2014. Assignee: Pure Storage.
- PA12. **Ethan Miller**, John Colgrove, John Hayes. Multi-drive cooperation to generate an encryption key, US Patent 8,745,415, issued June 3, 2014. Assignee: Pure Storage.
- PA11. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of reactive I/O operations in an storage environment. US Patent 8,732,426, issued May 20, 2014. Assignee: Pure Storage.
- PA10. **Ethan Miller**, John Colgrove, John Hayes, Cary Sandvig, Fractal layout of data blocks across multiple devices. US Patent 8,719,540, issued May 6, 2014. Assignee: Pure Storage.

- PA9. John Colgrove, John Hayes, **Ethan Miller**, Feng Wang, Logical sector mapping in a flash storage array. US Patent 8,645,664, issued Feb. 4, 2014. Assignee: Pure Storage.
- PA8. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of I/O writes in a storage environment, environment. US Patent 8,645,657, issued Feb. 4, 2014. Assignee: Pure Storage.
- PA7. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of I/O in an SSD environment. US Patent 8,589,655, issued Nov. 19, 2013. Assignee: Pure Storage, Inc.
- PA6. John Colgrove, John Hayes, **Ethan Miller**, Joseph S. Hasbani, Cary Sandvig. Method for maintaining multiple fingerprint tables in a deduplicating storage system. US Patent 8,589,640, issued Nov. 19, 2013. Assignee: Pure Storage.
- PA5. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of reconstructive I/O read operations in an storage environment. US Patent 8,589,625, issued Nov. 19, 2013.
- PA4. John Colgrove, John Hayes, **Ethan Miller**, Cary Sandvig, Joseph S. Hasbani, Feng Wang. Garbage collection in a storage system. US Patent 8,527,544, issued Sep. 3, 2013. Assignee: Pure Storage.
- PA3. **Ethan Miller**, John Colgrove, John Hayes. Efficient techniques for aligned fixed-length compression. US Patent 8,497,788, issued July 30, 2013. Assignee: Pure Storage.
- PA2. John Colgrove, John Hayes, Bo Hong, Feng Wang, **Ethan Miller**, Craig Harmer. Scheduling of I/O writes in a storage environment. US Patent 8,468,318, issued June 18, 2013. Assignee: Pure Storage.
- PA1. John Colgrove, John Hayes, Bo Hong, **Ethan Miller**. Intra-device data protection in a RAID array, US Patent 8,463,991, issued June 11, 2013. Assignee: Pure Storage.

Technical Reports

- T25. Daniel Bittman[‡], Matthew Bryson[‡], Yuanjiang Ni[‡], Arjun Govindjee*, Isaak Cherdak*, Pankaj Mehra, Darrell D. E. Long, **Ethan L. Miller**, “Twizzler: An Operating System for Next-Generation Memory Hierarchies”, Technical Report UCSC-SSRC-17-01, December 2017.
- T24. Shesha Sreenivasamurthy[‡], **Ethan L. Miller**, “SIVSHM: Secure Inter-VM Shared Memory”, Technical Report UCSC-SSRC-16-01, May 2016.
- T23. James S. Plank, **Ethan L. Miller**, Kevin M. Greenan, Benjamin A. Arnold*, John A. Burnum*, Adam W. Disney*, Allen C. McBride*, “GF-Complete: A Comprehensive Open Source Library for Galois Field Arithmetic”, Technical Report UT-CS-13-716, Computer Science Department, University of Tennessee, October 2013.
- T22. Avani Wildani[‡], **Ethan L. Miller**, Ian F. Adams, Darrell D. E. Long, “PERSES: Data Layout for Low Impact Failures”, Technical Report UCSC-SSRC-12-06, September 2012.
- T21. Avani Wildani[‡], **Ethan L. Miller**, Ohad Rodeh, “HANDS: A Heuristically Arranged Non-Backup In-line Deduplication System”, Technical Report UCSC-SSRC-12-03, March 2012.
- T20. Brian Madden*, Ian Adams[‡], Joel Frank[‡], **Ethan L. Miller**, “Analyzing User Behavior: A Trace Analysis of the NCAR Archival Storage System”, Technical Report UCSC-SSRC-12-02, March 2012.
- T19. Aleatha Parker-Wood*, Darrell D. E. Long, **Ethan L. Miller**, Margo Seltzer, Daniel Tunkelang, “Making Sense of File Systems Through Provenance and Rich Metadata”, Technical Report UCSC-SSRC-12-01, March 2012.
- T18. Ian F. Adams[‡], **Ethan L. Miller**, David S. H. Rosenthal, “Using Storage Class Memory for Archives with DAWN, a Durable Array of Wimpy Nodes”, Technical Report UCSC-SSRC-11-07, October 2011.
- T17. Brian Madden*, Ian F. Adams[‡], Mark W. Storer, **Ethan L. Miller**, Darrell D. E. Long, Thomas Kroeger, “Provenance Based Rebuild: Using Data Provenance to Improve Reliability”, Technical Report UCSC-SSRC-11-04, May 2011.
- T16. Ian F. Adams[‡], Mark W. Storer, and **Ethan L. Miller**, “Analysis of Workload Behavior in Scientific and Historical Long-Term Data Repositories”, Technical Report UCSC-SSRC-11-01, March 2011.
- T15. Andrew W. Leung[‡], Ian F. Adams[‡], **Ethan L. Miller**, “Magellan: A Searchable Metadata Architecture for Large-Scale File Systems”, Technical Report UCSC-SSRC-09-07, November 2009.
- T14. Andrew W. Leung[‡], Aleatha Parker-Wood[‡], **Ethan L. Miller**, “Copernicus: A Scalable, High-Performance Semantic File System”, Technical Report UCSC-SSRC-09-06, October 2009.

- T13. Avani Wildani[‡], Thomas J. E. Schwarz, **Ethan L. Miller**, Darrell D. E. Long, “Protecting Against Rare Event Failures in Archival Systems”, Technical Report UCSC-SSRC-09-03, April 2009.
- T12. Sasha Ames*, Carlos Maltzahn, **Ethan L. Miller**, “Quasar: A Scalable Naming Language for Very Large File Collections”, Technical Report UCSC-SSRC-08-04, October 2008.
- T11. Sasha Ames*, Carlos Maltzahn, **Ethan L. Miller**, “QUASAR: Interaction with File Systems Using a Query and Naming Language”, Technical Report UCSC-SSRC-08-03, September 2008.
- T10. Andrew W. Leung[‡], Minglong Shao, Timothy Bisson, Shankar Pasupathy, **Ethan L. Miller**, “Spyglass: Fast, Scalable Metadata Search for Large-Scale Storage Systems”, Technical Report UCSC-SSRC-08-01, Storage Systems Research Center, University of California, Santa Cruz, May 2008.
- T9. Kevin M. Greenan[‡], **Ethan L. Miller**, Thomas J. E. Schwarz, S. J., “Analysis and Construction of Galois Fields for Efficient Storage Reliability”, Technical Report UCSC-SSRC-07-09, Storage Systems Research Center, University of California, Santa Cruz, August 2007.
- T8. Mark W. Storey[‡], Kevin Greenan[‡], **Ethan L. Miller**, and Kaladhar Voruganti, “POTSHARDS: Secure Long-Term Archival Storage Without Encryption”, Technical Report UCSC-SSRC-06-03, Storage Systems Research Center, University of California, Santa Cruz, September 2006.
- T7. Sage A. Weil*, Scott A. Brandt, **Ethan L. Miller**, and Carlos Maltzahn, “CRUSH: Controlled, Scalable And Decentralized Placement Of Replicated Data”, Technical Report SSRC-06-02, Storage Systems Research Center, University of California, Santa Cruz, March 2006.
- T6. Sage A. Weil*, Feng Wang*, Qin Xin[‡], Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Carlos Maltzahn, “Ceph: a Scalable Object-Based Storage System”, Technical Report SSRC-06-01, Storage Systems Research Center, University of California, Santa Cruz, March 2006.
- T5. Geoff Kuenning and **Ethan L. Miller**, “Anonymization Techniques for URLs and Filenames”, Technical Report UCSC-CRL-03-05, Storage Systems Research Center, University of California, Santa Cruz, September 2003.
- T4. Nathan K. Edel[‡], **Ethan L. Miller**, Karl S. Brandt*, and Scott A. Brandt, “Measuring the Compressibility of Metadata and Small Files for Disk/NVRAM Hybrid Storage Systems”, Technical Report UCSC-CRL-03-04, Storage Systems Research Center, University of California, Santa Cruz, July 2003.
- T3. Timothy J. Gibson[‡] and **Ethan L. Miller**, “Long-Term File Activity in Diverse UNIX Environments”, Technical Report TR-CS-97-07, University of Maryland Baltimore County, October 1997.
- T2. R. Scott Cost*, Ian Soboroff*, Jeegar Lakhani*, Tim Finin, **Ethan Miller**, and Charles Nicholas, “TKQML: A KQML Extension to Tcl”, Technical Report TR-CS-97-04, University of Maryland Baltimore County, July 1996.
- T1. **Ethan Miller** and Jeffrey Hollingsworth, “Using Content-Derived Names for Caching and Software Distribution”, Technical Report TR-CS-96-08, University of Maryland Baltimore County, July 1996. Also available as UMIACS Technical Report TR-96-55.

Software

- SW5. *gferasure*: A high-performance Galois field and erasure code library written in C++ that leverages vector instructions in Intel and ARM processors. This software supersedes *GF-Complete*.
- SW4. *GF-Complete*: A high-performance Galois field library written in C that leverages specialized instruction sets in Intel processors. This software was co-authored with James Plank (University of Tennessee) and Kevin Greenan (EMC). The library is documented in a paper published in FAST 2013.
- SW3. *GaloisField*: A high-performance Galois field library written in C/C++ with an efficient Python interface. This software was co-authored with Kevin Greenan, one of my Ph.D. students.
- SW2. *DLX0S*: A CPU simulator and operating system used in undergraduate operating systems classes at the University of California, Santa Cruz, the University of Arizona, Purdue, the University of Maryland Baltimore County, and elsewhere.
- SW1. *esim*: A simple digital logic design language and simulator. Used for teaching computer architecture at the University of Maryland Baltimore County and elsewhere.

PROFESSIONAL ACTIVITIES

Work-in-Progress Posters and Presentations

- WP21. Matt Bryson[‡], Daniel Bittman[‡], Darrell Long, **Ethan Miller**, “Twizzler: The Design and Implementation of a NVM Aware OS”, 2017 Non-Volatile Memories Workshop (NVMW 2017), San Diego, CA, March 2017.
- WP20. Mathieu Ogleari*, Jishen Zhao, **Ethan Miller**, “Relaxing Persistent Memory Constraints with Hardware-Driven Undo+Redo Logging”, 2017 Non-Volatile Memories Workshop (NVMW 2017), San Diego, CA, March 2017.
- WP19. Anastasia McTaggart[‡], **Ethan L. Miller**, Sinjoni Mukhopadhyay[‡], “Clinker: Reconstructing Sharded Data Stores Efficiently”, work-in-progress presentation at the 14th Conference on File and Storage Technologies (FAST 2016), Santa Clara, CA, February 2016.
- WP18. Stephanie N. Jones*, Ahmed Amer, Rekha Pitchumani[‡], Darrell D. E. Long, **Ethan L. Miller**, “Data Layouts to Reduce Data Movement in Shingled Write Disks” work-in-progress and poster presentation at the 13th Conference on File and Storage Technologies (FAST 2015), Santa Clara, CA, February 2015.
- WP17. Joel C. Frank*, Shayna M. Frank[‡], Lincoln Thurlow[‡], Thomas M. Kroeger, **Ethan L. Miller**, Darrell D. E. Long, “Searching a Secret Split Datastore”, work-in-progress and poster presentation at the 13th Conference on File and Storage Technologies (FAST 2015), Santa Clara, CA, February 2015.
- WP16. Preeti Gupta[‡], Avani Wildani, **Ethan L. Miller**, Darrell D. E. Long, David S. H. Rosenthal “An Economic Argument For a Long Planning Horizon in Archival Storage”, poster presentation at the 13th Conference on File and Storage Technologies (FAST 2015), Santa Clara, CA, February 2015.
- WP15. Shesha Sreenivasamurthy[‡] and **Ethan L. Miller**, “VMOFS: Diskless and Efficient Object File System for Virtual Machines”, poster presentation at the 12th Conference on File and Storage Technologies (FAST 2014), Santa Clara, CA, February 2014.
- WP14. Ian F. Adams[‡], **Ethan L. Miller**, Mark W. Storer, Avani Wildani[‡], Yangwook Kang[‡], “Improved Analysis and Trace Validation Using Metadata Snapshots”, poster presentation at the 11th Conference on File and Storage Technologies (FAST 2013), San Jose, CA, February 2013.
- WP13. Daniel C. Rosenthal[‡], David S. H. Rosenthal, **Ethan L. Miller**, Ian F. Adams[‡], Mark W. Storer, and Erez Zadok, “Toward an Economic Model of Long-Term Storage”, poster and work-in-progress presentation at the 10th Conference on File and Storage Technologies (FAST 2012), San Jose, CA, February 2012.
- WP12. Rekha Pitchumani[‡], Yulai Xie*, Andy Hospodor, Ahmed Amer, and **Ethan L. Miller**, “Emulating a Shingled Write Disk”, poster and work-in-progress presentation at the 10th Conference on File and Storage Technologies (FAST 2012), San Jose, CA, February 2012.
- WP11. Andy Hospodor, Ahmed Amer, **Ethan L. Miller**, Darrell D. E. Long, Rekha Pitchumani[‡], Yangwook Kang[‡], and Yulai Xie*, “A Unified Object Oriented Storage Architecture”, poster and work-in-progress presentation at the 10th Conference on File and Storage Technologies (FAST 2012), San Jose, CA, February 2012.
- WP10. Ian F. Adams[‡] and **Ethan L. Miller**, “Challenges in Long-Term System Logging”, poster and work-in-progress presentation at the 10th Conference on File and Storage Technologies (FAST 2012), San Jose, CA, February 2012.
- WP9. Avani Wildani[‡] and **Ethan L. Miller**, “Grouping Data for Faster Rebuilds: The Art of Failing Silently”, poster and work-in-progress presentation at the 10th Conference on File and Storage Technologies (FAST 2012), San Jose, CA, February 2012.
- WP8. Ian F. Adams[‡], Mark W. Storer, and **Ethan L. Miller**, “Analysis of Workload Behavior in Scientific and Historical Long-Term Data Repositories”, poster at the 9th Conference on File and Storage Technologies (FAST 2011), San Jose, CA, February 2011.
- WP7. Yangwook Kang[‡], Jingpei Yang[‡], and **Ethan L. Miller**, “Object-based SCM: An Efficient Interface for Storage Class Memories”, poster at the 9th Conference on File and Storage Technologies (FAST 2011), San Jose, CA, February 2011.
- WP6. Aleatha Parker-Wood[‡], Christina Strong*, **Ethan L. Miller**, and Darrell D. E. Long, “Security Aware Partitioning for Efficient File Systems Search”, work-in-progress presentation and poster at the 8th Conference on File and Storage Technologies (FAST 2010), San Jose, CA, February 2010.

- WP5. Avani Wildani[‡] and **Ethan L. Miller**, “Probabilistic Reputation for Personal Trust Networks”, work-in-progress presentation and poster at the 7th Conference on File and Storage Technologies (FAST 2009), San Francisco, CA, February 2009.
- WP4. Sasha Ames[‡], Carlos Maltzahn, and **Ethan L. Miller**, “A File System Query Language”, poster at the 21st Symposium on Operating Systems Principles (SOSP 2007), Stevenson, WA, October 2007.
- WP3. Mark W. Storer[‡], Kevin Greenan[‡], and **Ethan L. Miller**, “Secure Long-Term Archival Storage with POTSHARDS”, work-in-progress presentation at the 5th Conference on File and Storage Technologies (FAST 2007), San Jose, CA, February 2007.
- WP2. Kevin Greenan[‡] and **Ethan L. Miller**, “CompulsiveFS: Making NVRAM Suitable for Extremely Reliable Storage”, work-in-progress presentation at the 5th Conference on File and Storage Technologies (FAST 2007), San Jose, CA, February 2007.
- WP1. Andrew Leung[‡] and **Ethan L. Miller**, “Scaling Security for Big, Parallel File Systems”, work-in-progress presentation at the 5th Conference on File and Storage Technologies (FAST 2007), San Jose, CA, February 2007.

Conference Organization

- 2021 **Program Committee:** 19th Conference on File and Storage Technologies (FAST 2021);
13th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 2021);
2021 USENIX Annual Technical Conference.
Steering Committee: 14th Annual Israeli Experimental Systems Conference (SYSTOR 2021).
Symposium on Operating Systems Principles.
- 2020 **Program Committee:** 18th Conference on File and Storage Technologies (FAST 2020);
13th Annual Israeli Experimental Systems Conference (SYSTOR 2020).
Steering Committee: 13th Annual Israeli Experimental Systems Conference (SYSTOR 2020);
Symposium on Operating Systems Principles.
- 2019 **Program Committee:** 2019 USENIX Annual Technical Conference;
11th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 2019).
Steering Committee: 12th Annual Israeli Experimental Systems Conference (SYSTOR 2019);
Symposium on Operating Systems Principles.
- 2018 **Program Committee:** 16th Conference on File and Storage Technologies (FAST 2018);
26th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2018);
10th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 2018);
2018 Non-Volatile Memories Workshop (NVMW 2018).
Steering Committee: 11th Annual Israeli Experimental Systems Conference (SYSTOR 2018);
Symposium on Operating Systems Principles.
- 2017 **Program Committee:** 29th International Conference on Scientific and Statistical Database Management (SSDBM 2017);
9th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 2017).
Steering Committee: 10th Annual Israeli Experimental Systems Conference (SYSTOR 2017).
- 2016 **Program Committee:** 14th Conference on File and Storage Technologies (FAST 2016);
2016 Non-Volatile Memories Workshop (NVMW 2016).
36th International Conference on Distributed Computing Systems (ICDCS 2016).
Steering Committee: 9th Annual Israeli Experimental Systems Conference (SYSTOR 2016).
- 2015 **Program Committee:** 35th International Conference on Distributed Computing Systems (ICDCS 2015);
7th Workshop on Hot Topics in Storage and File Systems (HotStorage '15);
23rd IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2015);
2015 Non-Volatile Memories Workshop (NVMW 2015).
General Chair: 25th Symposium on Operating Systems Principles (SOSP 2015).

- Steering Committee:** 8th Annual Israeli Experimental Systems Conference (SYSTOR 2015);
31st Conference on Mass Storage Systems and Technologies (MSST 2015).
- 2014 **Program Committee:** 12th Conference on File and Storage Technologies (FAST 2014);
34th International Conference on Distributed Computing Systems (ICDCS 2014);
IEEE Cluster 2014;
30th Conference on Mass Storage Systems and Technologies (MSST 2014);
7th Annual Israeli Experimental Systems Conference (SYSTOR 2014);
2014 Non-Volatile Memories Workshop (NVMW 2014).
- Steering Committee:** 7th Annual Israeli Experimental Systems Conference (SYSTOR 2014);
30th Conference on Mass Storage Systems and Technologies (MSST 2014).
- 2013 **Program Committee:** 11th Conference on File and Storage Technologies (FAST 2013);
29th IEEE Conference on Mass Storage Systems and Technologies (MSST 2013);
19th IEEE International Conference on Parallel and Distributed Systems (ICPADS 2013);
2013 Non-Volatile Memories Workshop (NVMW 2013);
SNIA Storage Developers Conference (SDC), New Thinking Track.
- Steering Committee:** 6th Annual Israeli Experimental Systems Conference (SYSTOR 2013);
29th IEEE Conference on Mass Storage Systems and Technologies (MSST 2013).
- 2012 **Co-organizer:** Dagstuhl workshop: “Is the Future of Preservation Cloudy?”, November, 2012.
Program co-Chair: 28th IEEE Conference on Mass Storage Systems and Technologies (MSST 2012).
Program Committee: International Conference for High Performance Computing, Networking, Storage
and Analysis (SC12);
5th Annual Israeli Experimental Systems Conference (SYSTOR 2012);
3rd International Green Computing Conference (IGCC '12);
20th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer
and Telecommunication Systems (MASCOTS 2012);
7th IEEE International Conference on Networking, Architecture, and Storage (NAS 2012);
9th International Conference on Preservation of Digital Objects (iPRES 2012).
- Steering Committee:** 5th Annual Israeli Experimental Systems Conference (SYSTOR 2012);
28th IEEE Conference on Mass Storage Systems and Technologies (MSST 2012).
- 2011 **Program Committee:** 27th IEEE Conference on Mass Storage Systems and Technologies (MSST 2011);
The 4th Annual Israeli Experimental Systems Conference (SYSTOR 2011);
19th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer
and Telecommunication Systems (MASCOTS 2011);
IEEE Cluster 2011 Conference;
2nd Workshop on Managing Systems via Log Analysis and Machine Learning
Techniques (SLAML 2011);
8th International Conference on Preservation of Digital Objects (iPRES 2011);
7th Petascale Data Storage Workshop (PDSW 2011).
- 2010 **Program co-Chair:** The Israeli Experimental Systems Conference (SYSTOR 2010).
Program co-Chair: First USENIX Workshop on Sustainable Information Technology (SustainIT '10).
Program Committee: International Conference for High Performance Computing, Networking, Storage
and Analysis (SC10);
26th IEEE Conference on Mass Storage Systems and Technologies (MSST 2010);
18th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer
and Telecommunication Systems (MASCOTS 2010);
Workshop on Managing Systems via Log Analysis and Machine Learning Tech-
niques (SLAML 2010);
7th International Conference on Preservation of Digital Objects (iPRES 2010);
- Steering Committee:** 26th IEEE Conference on Mass Storage Systems and Technologies (MSST 2010).
- 2009 **Program Committee:** 7th Conference on File and Storage Technologies (FAST 2009);
17th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer
and Telecommunication Systems (MASCOTS 2009);
The Israeli Experimental Systems Conference (SYSTOR 2009);

- 2008 **Program co-Chair:** 1st Workshop on Hot Topics in Storage and File Systems (HotStorage '09);
5th Workshop on Hot Topics in System Dependability (HotDep 2009).
Program Committee: 16th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2008);
28th International Conference on Distributed Computing Systems (ICDCS 2008);
4th ACM Workshop on Storage Survivability and Security (StorageSS 2008);
5th IEEE International Workshop on Storage Network Architecture and Parallel I/O (SNAPI);
- 2007 **Program Chair:** 1st International Workshop on Storage and I/O Virtualization, Performance, Energy, Evaluation and Dependability (SPEED 2008);
3rd International Workshop on Software Support for Portable Storage (IWSSPS 2008).
Area Vice-Chair: 24th IEEE Conference on Mass Storage Systems and Technologies (MSST 2007).
36th International Conference on Parallel Processing (ICPP 2007), Data-Intensive Computing track.
Program Committee: 3rd ACM Workshop on Storage Survivability and Security (StorageSS 2007);
4th International IEEE Security in Storage Workshop (SISW 2007).
- 2006 **Program co-Chair:** 2nd ACM Workshop on Storage Survivability and Security (StorageSS 2006).
Program Committee: 23rd IEEE / 14th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2006);
14th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2006);
7th Workshop on Distributed Data and Structures (WDAS 2006);
2nd International Workshop on Software Support for Portable Storage (IWSSPS 2006).
- 2005 **Program Committee:** 13th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2005);
22nd IEEE / 13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2005);
International Symposium on Emergence of Globally Distributed Data [also served as publications chair];
International Workshop on Software Support for Portable Storage (IWSSPS 2005);
1st ACM Workshop on Storage Survivability and Security (StorageSS);
3rd International IEEE Security in Storage Workshop;
Invited Talks co-chair: 2005 USENIX Technical Conference.
- 2004 **Program Committee:** 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004);
2nd Intelligent Storage Workshop;
6th Workshop on Distributed Data and Systems (WDAS 2004);
12th IEEE/ACM Conference on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2004), [also served as publications chair];
Supercomputing (SC) 2004.
- 2003 **Program Committee:** 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2003) [also served as publications chair];
2nd International IEEE Security in Storage Workshop.
Steering Committee: HotOS-IX: Ninth Workshop on Hot Topics in Operating Systems.
- 2002 **Program Committee:** 1st File and Storage Technologies Conference (FAST 2002) [also served as publications chair];
19th IEEE Mass Storage System Symposium /
10th NASA Goddard Mass Storage and Technologies Conference;
1st International IEEE Security in Storage Workshop;
IEEE International Performance, Computing, and Communications Conference.

- 2001 **Program Committee:** 18th IEEE Mass Storage System Symposium / 9th NASA Goddard Mass Storage and Technologies Conference.
- 2000 **Program Committee:** 17th IEEE Mass Storage System Symposium / 8th NASA Goddard Mass Storage and Technologies Conference.
- 1999 **Program Committee:** 16th IEEE Mass Storage System Symposium / 7th NASA Goddard Mass Storage and Technologies Conference [also served as publications chair].
- 1998 **Program Committee:** 15th IEEE Mass Storage System Symposium / 6th NASA Goddard Mass Storage and Technologies Conference.
- 1996 **Program Committee:** Workshop on I/O in Parallel and Distributed Systems (IOPADS).

Reviewer of Technical Papers and Proposals

- 2019 National Science Foundation (2 panel reviews).
- 2018 National Science Foundation (panel review).
- 2017 *ACM Transactions on Storage*.
- 2016 National Science Foundation (panel review); *IEEE Transactions on Parallel and Distributed Systems*.
- 2015 National Science Foundation (panel review); 13th Conference on File and Storage Technologies (FAST 2015).
- 2014 International Symposium on Network Coding (NetCod 2014); *IEEE Transactions on Parallel and Distributed Systems*; *ACM Transactions on Information and System Security*; *ACM Transactions on Storage*.
- 2013 National Science Foundation (*ad hoc* proposal review); *IEEE Transactions on Dependable and Secure Computing*; *ACM Computing Surveys*; *ACM Transactions on Storage*.
- 2012 *ACM Transactions on Storage*; *IEEE Internet Computing*; 19th International Symposium on High Performance Computer Architecture (HPCA-19).
- 2011 National Science Foundation (SBIR panel); *ACM Transactions on Storage*; *IEEE Transactions on Computers*; *IEEE Transactions on Parallel and Distributed Systems*; *IEEE Transactions on Information Theory*; Department of Energy (ASCR proposal panel); John Wiley & Sons.
- 2010 *IEEE Transactions on Dependable and Secure Computing*; *ACM Transactions on Storage*; *ACM Transactions on Information and System Security*; *IEEE Transactions on Knowledge and Data Engineering*; 8th Conference on File and Storage Technologies (FAST); National Science Foundation (*ad hoc* review).
- 2009 *IEEE Transactions on Knowledge and Data Engineering*; *Journal of Parallel and Distributed Computing*; National Science Foundation (SBIR panel); National Science Foundation (CAREER panel); UC Discovery Grant; Addison Wesley; John Wiley & Sons.
- 2008 *IEEE Transactions on Parallel and Distributed Systems*; *ETRI (Electronics and Telecommunications Research Institute) Journal*; 6th Conference on File and Storage Technologies (FAST); British Computer Society Distinguished Dissertation Award; National Science Foundation (CISE panel).
- 2007 *IEEE Transactions on Parallel and Distributed Systems*; 5th Conference on File and Storage Technologies (FAST); *Journal of Systems and Software*; *IBM Journal of Research and Development*; IEEE Wireless Communications and Networking Conference; US-Israel Bi-national Science Foundation; National Science Foundation (CISE panel); National Science Foundation (CAREER panel); UC Discovery Grant.
- 2006 Conference on Dependable Systems and Networks (DSN); 7th Conference on Operating Systems Design and Implementation (OSDI); National Science Foundation (CAREER panel); National Science Foundation (SBIR panels).
- 2005 *IEEE Transactions on Computers*; 4th Conference on File and Storage Technologies (FAST); 19th International Symposium on Distributed Computing (DiSC 2005); National Science Foundation (CISE panel).
- 2004 3rd Conference on File and Storage Technologies (FAST); Conference on Dependable Systems and Networks (DSN); 6th Symposium on Operating Systems Design and Implementation (OSDI); *Software—Practice and Experience*; *ACM Transactions on Computer Systems*; National Science Foundation (CISE panel).

- 2003 *IEEE Transactions on Parallel and Distributed Systems*; Usenix Technical Conference; 2nd Conference on File and Storage Technologies; SC2003; National Science Foundation (SBIR panel).
- 2002 *IEEE Transactions on Computers*, *IEEE Transactions on Parallel and Distributed Systems*, Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 5th Symposium on Operating Systems Design and Implementation (OSDI), Usenix Annual Technical Conference, National Science Foundation (SBIR panel).
- 2001 Usenix Annual Technical Conference; *IEEE Transactions on Computers*; National Science Foundation (ITR panel, SBIR panel).
- 1990–2000 National Science Foundation (panel reviewer); *ACM Transactions on Computer Systems*; *IEEE Computer*; International Symposium on Computer Architecture; *Journal of Parallel and Distributed Computing*; Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS); Symposium on Operating Systems Design and Implementation (OSDI); USENIX Annual Technical Conference; Oxford University Press.

Membership in Professional Associations

- 1995–present Member, Usenix Association.
- 1990–present Association for Computing Machinery: Distinguished Member (2013).
- 1990–present IEEE Computer Society: Senior Member (2001), Fellow (2015).
- 1987–present Member, Sigma Xi.

Service to Professional Associations

- 2021 IEEE Computer Society Fellows Evaluation Committee.
- 2019 IEEE Computer Society Fellows Evaluation Committee.
- 2017 IEEE Computer Society Fellows Evaluation Committee.
- 2001–2008 Chair, *IEEE Technical Committee on Operating Systems and Applications Environments*.
- 1996–2013 Member, Executive Committee, *IEEE Technical Committee on Mass Storage Systems*.
- 1995–2000 University Liaison for the Usenix Association.

Consulting & Expert Witness Activity

- 2021 Consulting expert, WilmerHale LLP
- 2018 Expert witness, Fish & Richardson
- 2013 Consulting expert, Arnold Porter LLP
- 2009–present Consultant / Research Scientist, Pure Storage
- 2010–2011 Consulting expert, Foley Lardner LLP
- 2009–2011 Consulting and testifying expert, Sidley Austin LLP
- 2008–2009 Consulting expert, Sheppard Mullin Richter & Hampton LLP
- 2007 Hewlett Packard Laboratories
- 2006–2007 Expert witness, Fish & Richardson
- 2005–2006 Veritas (now Symantec) Corporation
- 2003 Hewlett Packard Laboratories
- 2001–2005 Expert witness, Bartlit, Beck, Herman, Palenchar & Scott
- 2000 Expert witness, Fish & Richardson
- 1998–2001 Expert witness, Hopgood, Calimafde, Judlowe & Mondolino
- 1998 Web site architect, Ambleside Logic

Editorial Duties

- 2014–2016 Associate Editor, *ACM Transactions on Storage*.
- 2009–2014 Associate Editor, *IEEE Transactions on Computers*.

Advisory Boards

2011–present	Technical Advisory Board, Chronicle of Life
2011–2015	Technical Advisory Board, 8kpc
2010–2011	Technical Advisory Board, Pancetera (purchased by Quantum in 2011)

Invited Talks

This list contains details for recent talks only.

2021 June	“Twizzler: Rethinking the Operating System Stack for Byte-Addressable NVM”, IBM Almaden Research Center, San Jose, CA. Talk given remotely.
2021 March	“Twizzler: Rethinking the Operating System Stack for Byte-Addressable NVM”, keynote talk at the 12th Annual Non-Volatile Memories Workshop, University of California San Diego. Talk given remotely.
2020 July	“The Future of the Past: Challenges in Archival Storage”, keynote talk at the 2020 USENIX Annual Technical Conference (ATC '20). Talk given remotely.
2019 June	“Twizzler: A Data-Centric Operating System for Persistent Memory”, IDC Herzlyia, Israel.
2019 June	“Twizzler: A Data-Centric Operating System for Persistent Memory”, Technion, Haifa, Israel.
2019 May	“Optimizing Systems for Byte-Addressable NVM by Reducing Bit Flipping”, Johannes Gutenberg University, Mainz, Germany.
2019 May	“Twizzler: A Data-Centric Operating System for Persistent Memory”, Johannes Gutenberg University, Mainz, Germany.
2019 Mar	“Archival Storage: Performance and Economics” (joint talk with Erez Zadok, Stony Brook University), 2nd Workshop on DNA for Storage, Banbury Center at Cold Spring Harbor Laboratory, NY.
2018 Oct	“Twizzler: An Operating System Designed for Non-Volatile Memory”, University of North Carolina, Chapel Hill, NC.
2018 Aug	“Adapting Systems for Next-Generation Non-Volatile Memory Technologies”, Huawei, Santa Clara, CA.
2017 Oct	“Rethinking the System Stack for Persistent Memories”, Seagate Technology Academic Summit, Fremont, CA.
2017 Sep	“CAPES: Unsupervised Storage Performance Tuning Using Neural Network-Based Deep Reinforcement Learning”, CERN, Geneva, Switzerland.
2017 Aug	“QLC and the Future of Flash Storage”, Huawei, Chengdu, China.
2016 Aug	“Performance and Durability: Storage Directions for the Next Ten Years”, Veritas, Mountain View, CA.
2016 Jun	“Leveraging Flash Storage Characteristics for High-Performance Block Storage”, IBM Haifa Research Laboratory, Haifa, Israel.
2016 Apr	“Leveraging Flash Storage Characteristics for High-Performance Block Storage”, Northeastern University, Boston, MA.
2016 Feb	“The Real Dangers in Cybersecurity”, Temple Beth El, Aptos, CA.
2016 Jan	“Rethinking Benchmarks for Non-Volatile Memory Storage Systems”, NVM Summit—Convergence of Storage and Memory—Developing the Needed Ecosystem, San Jose, CA.
2015 Nov	“Re-architecting the Systems Stack for Byte-Addressable NVRAM”, Sandisk, Milpitas, CA.
2015 Sep	“CRSS & SSRC Overview”, SK Hynix, San Jose, CA.
2014 Sep	“Muninn: a Versioning Flash Key-Value Store Using an Object-Based Storage Model”, Samsung, Milpitas, CA.
2014 Jul	“The Future is Remembering the Past: Challenges in Archival Storage”, Huazhong University of Science and Technology, Wuhan, China.
2014 Jul	“Extreme-Scale Storage for High-Performance Computing”, Huazhong University of Science and Technology, Wuhan, China.
2014 May	“The Future is Remembering the Past: Challenges in Archival Storage”, Huawei, Shenzhen, China.
2014 Mar	“The Challenges of Building Useful Archival Storage”, Huawei, Santa Clara, CA.

- 2013 Dec “Inside the Pure Storage Flash Array: Building a High Performance, Data Reducing Storage System from Commodity SSDs”, University of Toronto, Toronto, ON, Canada.
- 2013 Nov “Where’d My Photos Go? Challenges in Preserving Digital Data for the Long Term”, online seminar for Millibo, <http://www.millibo.com/>.
- 2013 Nov “Where’d My Photos Go? Challenges in Preserving Digital Data for the Long Term”, University of Houston, Houston, TX.
- 2013 Oct “Inside the Pure Storage Flash Array: Building a High Performance, Data Reducing Storage System from Commodity SSDs”, Carnegie Mellon University, Pittsburgh, PA.
- 2013 Oct “Inside the Pure Storage Flash Array: Building a High Performance, Data Reducing Storage System from Commodity SSDs”, Brown University, Providence, RI.
- 2013 Sep “Screaming Fast Galois Field Arithmetic Using Intel SIMD Instructions”, SNIA Storage Developers Conference, Santa Clara, CA.
- 2013 Sep “Percival: Securely Searching a Secret Split Archive”, 15th International Workshop on High Performance Transaction Systems (HPTS), Asilomar, CA.
- 2013 Jun “The What, Why and How of the Pure Storage Enterprise Flash Array”, Microsoft Research, Mountain View, CA.
- 2013 May “The What, Why and How of the Pure Storage Enterprise Flash Array”, Johannes Gutenberg Universität, Mainz, Germany.
- 2013 Mar “Where’d My Photos Go? Challenges in Preserving Digital Data for the Long Term”, Johannes Gutenberg Universität, Mainz, Germany.
- 2013 Feb “The What, Why and How of the Pure Storage Enterprise Flash Array”, Stony Brook University, Stony Brook, NY.
- 2012 **8 talks**, including invited talks in India and Uruguay, as well as Microsoft Research and a Library of Congress workshop.
- 2011 **4 talks**, including an invited talk at the Technion in Israel and a panel discussion on relevant research in storage at MSST 2011.
- 2010 **6 talks**, including two invited talks in Japan, one invited talk in Australia, and an invited talk for the genomics community.
- 2009 **9 talks**, including keynote talks at IEEE NAS ’09 and Symantec, invited talks at IBM Research and the Association of Moving Image Archivists, as well as international talks in the Netherlands and China.
- 2008 **17 talks**, including talks at IBM, UC Berkeley (CITRIS), an NSF workshop, and the Association of Moving Image Archivists, as well as international talks in South Korea and Germany.
- 2007 **16 talks**, including talks at UC Berkeley, UC Santa Barbara, Silicon Valley industry, and the University of New Mexico. In addition, gave five talks during a week-long visit to the Data Storage Institute at the National University of Singapore.
- 2006 **9 talks**, including talks at Brown University, the University of Rhode Island, Sandia National Laboratory, and the National Security Agency.
- 2005 **5 talks**, including a keynote address at a workshop at IBM Haifa (Israel), Stony Brook University, and the University of Maryland Baltimore County.
- 1996–2004 **16 talks**, including talks at Lawrence Livermore National Laboratory, Sandia National Laboratory, and the University of Wisconsin Milwaukee.

UNIVERSITY SERVICE

Academic Senate Service

- 2018–2019 Senate Committee on Information Technology
- 2016–2018 Senate Committee on Academic Personnel
- 2016–2017 CCA Mentor for 3 faculty
- 2015–2016 Committee on Career Advising
- 2012–2014 Committee on Academic Freedom
- 2007–2008 Committee on Faculty Research Lecture

- 2006–2007 Committee on Computing and Telecommunications (Fall & Spring)
 2004–2006 Chair, Committee on Computing and Telecommunications
 2003–2004 Committee on Computing and Telecommunications (Winter & Spring)

Service to the School of Engineering

- 2017 Baskin School of Engineering Reshaping Committee
 2001–2011 Computing Infrastructure Committee (**chair**, 2002–2005, 2007–2010)

Service to Crown College

- 2010–2014 Crown College Executive Committee

Service to the Department

NOTE: Service in 2000–2016 is to the Computer Science Department. Service beginning the 2016–17 academic year is to the Computer Engineering Department. Service beginning in the 2018–19 academic year is to the Computer Science and Engineering Department.

- 2018–2019 Faculty Search Committee
 2017–2018 Faculty Search Committee
 2016–2017 Personnel Committee
 2015–2016 Co-chair, Faculty Search Committee (2 positions)
 2014–2015 Chair, Faculty Search Committee
 2013–2015 Invited Lecturer Committee
 2011–2012 Committee on Online Education
 2010–2011 Strategic Directions Committee
 2004–2016 Personnel Committee
 2004–2005 Faculty Search Committee
 2003–2004 Faculty Search Committee
 2003–2010 Computing Committee
 2002–2003 Faculty Search Committee
 2001–2002 Faculty Search Committee

In addition to the formal service listed above, I have maintained the `git` server for the School of Engineering (`gitlab.soe.ucsc.edu`). Over the past three years, the system has hosted over 5,000 users, primarily SOE undergraduates, and approximately 10,000 student-quarters of projects. This system has provided significant benefit for students, who gain experience using industry-standard tools for developing software, which includes the ability to track changes to their software projects, view changes over time, and help with fixing software bugs they may have introduced.

Service to the University of California

- 2004–06 Committee on Information Technology and Telecommunications Policy.

Other Service to the Campus

- 2012– UCSC Extension Engineering Advisory Board
 2003 Information Technology Vision Committee.

PUBLIC SERVICE

- 2014– Center for Academic Engagement Faculty Fellow, Israel on Campus Coalition.
 2011–2014 Area Coordinator, Brown Alumni Interviewing Program
 2003–05 Member, Santa Cruz Hillel Board of Directors
 2002–2015 Alumni interviewer, Brown University (area chair 2012–2014)

MENTORING AND STUDENT ADVISING

Postdoctoral Scholars

Dates	Relationship	Degree Year	Name and Activities
2010 Fall–2012 Fall	Primary Supervisor		Arifa Nisar
2011 Fall–2013 Summer	Other Advisor		Yasuhiro Ohara (advisor: Darrell Long)

Advisor to Continuing Graduate Students

Allen Aboytes	Ph.D.	2019–
Daniel Bittman	Ph.D.	2016–
Matthew Bryson	Ph.D.	2016–
James Byron	Ph.D.	2016–
Yuanjiang Ni	Ph.D.	2016–
Devashish Purandare	Ph.D.	2017–

Doctoral Advisor

2015	Rekha Pitchumani	<i>Data Management for Shingled Magnetic Recording Disks</i> (won Computer Science Department Dissertation Award)
2014	Yangwook Kang	<i>High-Performance, Reliable Object-Based NVRAM Devices</i>
2013	Ian Adams	<i>Understanding Long-Term Storage Access Patterns</i>
2013	Avani Wildani	<i>The Promise of Data Grouping in Large Scale Storage Systems</i>
2009	Andrew Leung	<i>Organizing, Indexing, and Searching Large-Scale File Systems</i>
2009	Kevin Greenan	<i>Reliability and Efficiency in Erasure-Coded Storage Systems</i>
2009	Mark Storer	<i>Secure, Energy-Efficient, Evolvable, Long-Term Archival Storage</i>
2005	Qin Xin	<i>Understanding and Coping with Failures in Large-Scale Storage Systems</i>
2004	Ismail Ari	<i>Design and Management of Globally-Distributed Network Caches</i>
2002	Naomi Avigdor	<i>Building a Scalable and Reliable Parallel File System Using Commodity Computers</i>
2000	William Freeman	<i>Decentralized Security for Network-Attached Storage</i>
1998	Timothy Gibson	<i>Long-term UNIX File System Activity and the Efficacy of Automatic File Migration</i>

Doctoral Dissertation Reading Committee Member

2021	James Hughes	<i>Prof. Darrell Long</i>
2019	Matheus Ogleari	<i>Prof. Jishen Zhao (UC San Diego)</i>
2019	Yang Zhan	<i>Prof. Donald Porter (University of North Carolina)</i>
2017	Joel Frank	<i>Prof. Darrell Long</i>
2017	Veronica Estrada Galinañes	<i>Prof. Pascal Felber (Université de Neuchâtel, Switzerland)</i>
2017	Yan Li	<i>Prof. Darrell Long</i>
2016	Stephanie Jones	<i>Prof. Darrell Long</i>
2015	Alex Nelson	<i>Prof. Darrell Long</i>
2015	D. J. Capelis	<i>Prof. Darrell Long</i>
2015	Matthias Grawinkel	<i>Prof. André Brinkmann (Johannes Gutenberg University, Mainz)</i>
2014	Aviad Zuck	<i>Prof. Sivan Toledo (Tel Aviv University, Israel)</i>
2014	Raja Appusawamy	<i>Prof. Andy Tanenbaum, (Vrije Universitet, Netherlands)</i>
2013	Aleatha Parker-Wood	<i>Prof. Darrell Long</i>
2013	Dirk Meister	<i>Prof. André Brinkmann (Johannes Gutenberg University, Mainz)</i>
2010	Deepavali Bhagwat	<i>Prof. Darrell Long</i>
2009	Mohammed Khatib	<i>Prof. Pieter Hartel (University of Twente, Netherlands)</i>
2008	David Pease	<i>Profs. Richard Hughey and Darrell Long</i>
2008	Vinay Pai	<i>Prof. Erez Zadok (Stony Brook University)</i>
2008	Guozheng Ge	<i>Prof. E. James Whitehead</i>
2007	Sage Weil	<i>Prof. Scott Brandt</i>
2006	Nikolai Joukov	<i>Prof. Erez Zadok (Stony Brook University)</i>
2006	Lawrence You	<i>Prof. Darrell Long</i>
2006	Feng Wang	<i>Prof. Scott Brandt</i>
2005	Bo Hong	<i>Profs. Darrell Long and Scott Brandt</i>
2005	Scott Banachowski	<i>Prof. Scott Brandt</i>
2002	Ahmed Amer	<i>Prof. Darrell Long</i>
2002	Tsozen (Frank) Yeh	<i>Prof. Darrell Long</i>
2000	Ian Soboroff	<i>Prof. Charles Nicholas (UMBC)</i>

Doctoral Qualifying Exam Committee Member

2021	Austen Barker*	<i>Prof. Darrell Long</i>
2020	Devashish Purandare	<i>Prof. Ethan Miller</i>
2020	James Hughes*	<i>Prof. Darrell Long</i>
2019	Chandranil Chakrabortii*	<i>Prof. Heiner Litz</i>
2019	Kamala Ramasubramanian*	<i>Prof. Peter Alvaro</i>
2018	Daniel Bittman	<i>Prof. Ethan Miller</i>
2018	James Byron	<i>Prof. Ethan Miller</i>
2018	Yuanjiang Ni	<i>Prof. Ethan Miller</i>
2018	Oceane Bel*	<i>Prof. Darrell Long</i>
2018	Sinjoni Mukhupadhyay*	<i>Prof. Darrell Long</i>
2018	Yang Zhan	<i>Prof. Donald Porter, University of North Carolina</i>
2017	Xin Li	<i>Prof. Chen Qian</i>
2014	Preeti Gupta	<i>Prof. Ethan Miller</i>
2014	Yan Li*	<i>Prof. Darrell Long</i>
2014	Joel Frank*	<i>Prof. Darrell Long</i>
2014	Alex Nelson*	<i>Prof. Darrell Long</i>
2012	Rekha Pitchumani	<i>Prof. Ethan Miller</i>
2012	Yangwook Kang	<i>Prof. Ethan Miller</i>
2012	Christina Strong*	<i>Prof. Darrell Long</i>
2012	Stephanie Jones*	<i>Prof. Darrell Long</i>
2012	Aleatha Parker-Wood*	<i>Prof. Darrell Long</i>
2012	Ian Adams	<i>Prof. Ethan Miller</i>
2011	Avani Wildani	<i>Prof. Ethan Miller</i>
2011	D. J. Capelis*	<i>Prof. Darrell Long</i>
2010	Lanbo Zhang	<i>Prof. Yi Zhang</i>
2008	Andrew Leung	<i>Prof. Ethan Miller</i>
2008	Dhananjay Sampath	<i>Prof. J. J. Garcia-Luna</i>
2008	Kevin Greenan	<i>Prof. Ethan Miller</i>
2008	Neerja Bhatnagar	<i>Prof. Ethan Miller</i>
2007	Mark W. Storer	<i>Prof. Ethan Miller</i>
2007	Deepavali Bhagwat*	<i>Prof. Darrell Long</i>
2006	Sage Weil*	<i>Prof. Scott Brandt</i>
2006	Elias Sinderson	<i>Prof. E. James Whitehead</i>
2006	Nikolai Joukov	<i>Prof. Erez Zadok (Stony Brook University)</i>
2005	David Pease	<i>Prof. Darrell Long</i>
2004	Guozheng Ge*	<i>Prof. E. James Whitehead</i>
2004	Damian Cieslicki	<i>Prof. Thomas Schwarz (Santa Clara University)</i>
2003	Scott Banachowski	<i>Prof. Scott Brandt</i>
2003	Qin Xin	<i>Prof. Ethan Miller</i>
2003	Bo Hong	<i>Prof. Darrell Long</i>
2003	Feng Wang	<i>Prof. Scott Brandt</i>
2002	Ismail Ari	<i>Prof. Ethan Miller</i>
2001	Tsozen Yeh	<i>Prof. Darrell Long</i>
1999	Naomi Avigdor	<i>Prof. Ethan Miller (UMBC)</i>
1999	William Freeman	<i>Prof. Ethan Miller (UMBC)</i>
1998	Ian Soboroff	<i>Charles Nicholas (UMBC)</i>
1997	Timothy Gibson	<i>Prof. Ethan Miller (UMBC)</i>
1995	Vincent Marier	<i>Prof. Deepinder Sidhu (UMBC)</i>

Note: * indicates that I served as qualifying exam committee chair.

Masters of Science Advisor

2019	Dev Purandare	non-thesis
2018	Kenneth Chang	non-thesis
2018	James Byron	title
2018	Matthew Bryson	non-thesis
2018	Kenneth Chang	non-thesis
2017	Sinjoni Mukhopadhyay	non-thesis
2017	Vedang Joshi	non-thesis
2017	Ethan Vadai	non-thesis
2017	Aneesh Neelam	non-thesis
2017	Preeti Gupta	non-thesis
2015	Runchen Liu	non-thesis
2015	Xiaoyuan Lu	<i>Storage Workload Characterization and Performance Prediction for Better I/O Traffic Management</i>
2014	Daniel Rosenthal	non-thesis
2014	Thomas Marlette	non-thesis
2012	Rekha Pitchumani	non-thesis
2011	Nathan Edel	<i>MRAMFS: A Compressing File System for Byte-Addressable Non-Volatile RAM</i>
2010	Keren Jin	<i>Deduplication on Virtual Machine Disk Images</i>
2010	Aleatha Parker-Wood	<i>Security Aware Partitioning for Efficient File System Search</i>
2010	Ian Adams	non-thesis
2009	Mrunal Gawade	non-thesis
2008	Casey Marshall	non-thesis
2008	Danni Fu	non-thesis
2007	Jeff Hagen	<i>Ladon: A Framework for Peer-to-Peer Backup</i>
2007	Max Mehech	non-thesis
2007	Andrew Leung	<i>Security in Scalable Storage Systems</i>
2006	Kevin Greenan	non-thesis
2006	Mark Storer	non-thesis
2005	Christopher Olson	non-thesis
2005	Chengyu Sung	<i>Integrating Pictorial Identity into Secure Email</i>
2005	Sasha Ames	non-thesis
2004	R. J. Honicky	<i>Object Placement Algorithms for OBSD Systems</i>
2001	Kennedy Akala	non-thesis
2000	Vivekand Krishnamoorthi	<i>A Comparison of Long Term File Migration Algorithms</i>
1999	Ting Chen	non-thesis
1998	Junli Liu	non-thesis
1998	Michael Shapiro	non-thesis
1998	Dan Shen	<i>Experiments with Large-Scale N-gram Based Information Retrieval</i>
1998	Changgong Zhang	non-thesis
1997	Eric Robertson	non-thesis
1996	Arun C. Mahendran	non-thesis
1996	Amen Zwa	non-thesis

Masters of Science Reading Committee Member

2021	Steffen Eiden	<i>Prof. Andre Brinkmann (Johannes Gutenberg University, Mainz)</i>
2020	Barbara Moretto Dama	<i>Prof. Katia Obraczka</i>
2020	Krystine Carrington	<i>Prof. Mircea Teodorescu</i>
2019	Staunton Sample	<i>Prof. Darrell Long</i>
2018	Pinglei Guo	<i>Prof. Peter Alvaro</i>
2018	Nakul Dhotre	<i>Prof. Darrell Long</i>
2016	Borui Wang	<i>Prof. Jishen Zhao</i>
2016	Jun Yuan	<i>Prof. Jishen Zhao</i>
2015	Ignacio Corderi	<i>Prof. Darrell Long</i>
2014	Michael McThrow	<i>Prof. Darrell Long</i>
2014	Erik Steggall	<i>Prof. Darrell Long</i>
2013	Alex Nelson	<i>Prof. Darrell Long</i>
2012	Akhilesh Malivalli	<i>Prof. Darrell Long</i>
2010	Stephanie Jones	<i>Prof. Darrell Long</i>
2008	Rosie Wacha	<i>Prof. Darrell Long</i>
2007	Corrie Scalisi	<i>Prof. Manfred Warmuth</i>
2006	Suma Potluri	<i>Prof. E. James Whitehead</i>
2006	Nikhil Bobb	<i>Prof. Scott Brandt</i>
2005	Travis Odegaard	<i>Prof. Scott Brandt</i>
2004	Svetlana Kagan	<i>Prof. Darrell Long</i>
2004	Deepa Tuteja (thesis)	<i>Prof. Scott Brandt</i>
2003	Suruchi Malapture (thesis)	<i>Prof. Scott Brandt</i>
2003	Karen Glocer	<i>Prof. Darrell Long</i>
2003	Ravindra Vaishampayan	<i>Prof. J. J. Garcia-Luna</i>
2003	Caixue Lin	<i>Prof. Scott Brandt</i>
1999	Amy Germida	<i>Prof. James Plusquellic (UMBC)</i>
1996	Greg Sylvain	<i>Prof. Tim Finin (UMBC)</i>
1996	Chetan Shah	<i>Prof. Charles Nicholas (UMBC)</i>
1996	Scott Stewart	<i>Prof. Tim Finin (UMBC)</i>

Bachelors of Science Thesis Advisor

2017	Varun Arora	<i>Capsule: Clientside Test System to Encompass and Execute ORAM Algorithms for Data Access Pattern Obfuscation</i>
------	-------------	---

COURSES TAUGHT

Note that classes from Fall 1994 through Spring 2000 are semester classes, and classes from Fall 2000 onward are quarter classes. Courses in bold italics are those for which I either developed or significantly revised the curriculum.

Undergraduate

Spring 2021 *Embedded Operating Systems*
 Fall 2020 Principles of Computer Systems Design
 Fall 2019 Principles of Computer Systems Design
 Winter 2019 *Principles of Computer Systems Design*
 Fall 2018 Computer Architecture
 Winter 2018 Computer Architecture
 Fall 2016 Introduction to Programming in Python
 Spring 2016 Operating Systems
 Fall 2015 Computer Security
 Spring 2015 Operating Systems
 Fall 2014 Computer Security
 Spring 2014 Introduction to Programming in Python
 Winter 2014 Computer Security
 Spring 2012 Operating Systems
 Spring 2011 Introduction to Programming in Python
 Winter 2011 Computer Security
 Spring 2010 Computer Security
 Spring 2009 Operating Systems
 Fall 2008 *Intro to Programming in Python*
 Spring 2008 Distributed Systems
 Fall 2007 Operating Systems
 Spring 2007 Computer Security
 Spring 2006 Operating Systems
 Winter 2006 *Distributed Systems*
 Winter 2005 Computer Security
 Fall 2004 Operating Systems
 Spring 2004 Computer Security
 Fall 2003 Operating Systems
 Spring 2003 *Computer Security*
 Winter 2003 Operating Systems
 Winter 2002 Introduction to Data Structures
 Fall 2001 Operating Systems
 Fall 2000 Operating Systems
 Fall 1999 Operating Systems
 Spring 1999 Operating Systems
 Fall 1997 Computer Architecture
 Spring 1997 Computer Architecture
 Spring 1996 Computer Architecture
 Spring 1995 *Computer Architecture*
 Fall 1994 Computer Architecture

Graduate

Winter 2020 Advanced Operating Systems
 Fall 2017 *Advanced Topics In Computer Engineering: Non-Volatile Memory Systems*
 Winter 2017 Advanced Operating Systems
 Winter 2016 Advanced Operating Systems
 Winter 2015 Storage Systems
 Fall 2013 Distributed Systems
 Fall 2011 *Special Topics in Computer Systems: Archival Storage*
 Winter 2011 Advanced Computer Security
 Winter 2010 Distributed Systems
 Winter 2009 Advanced Operating Systems
 Winter 2008 Storage Systems
 Spring 2007 Distributed Systems
 Fall 2005 Advanced Operating Systems
 Spring 2005 Distributed Systems
 Winter 2004 Storage Systems
 Fall 2002 Advanced Operating Systems
 Spring 2002 *Distributed Systems*
 Spring 2001 Computer Security
 Spring 2000 Computer Architecture
 Spring 2000 Storage Systems
 Fall 1998 Computer Architecture
 Spring 1998 Storage Systems
 Fall 1997 Computer Architecture
 Spring 1997 Operating Systems
 Fall 1996 Computer Architecture
 Fall 1995 Computer Architecture