

Andrew Gaul

Systems programmer interested in entrepreneurship and technical leadership opportunities in San Francisco and Tokyo. Background includes work in distributed computing, machine learning, storage, and free software.

Work experience

Determined AI Senior Systems Engineer 2018

- Developed a machine learning platform for long-lived and repeatable experiments across many GPUs and systems
- Led implementation of fault tolerance, multi-file models, and management of many experiments
- Contributed to model cost estimation, system dump, web UI, and championed upstream contributions

Consultant and Sabbatical 2016 - 2017

- Consulted for clients on Apache jclouds, goofys, s3fs, and S3Proxy and visited 17 countries across Asia and Europe

Bounce Storage Co-founder 2014 - 2015

- Created a hybrid object store for managing data spanning different clouds and geographies
- Implemented protocol translator comprised of an S3 frontend and an enhanced Apache jclouds backend
- Led outbound efforts including pitching to investors, sourcing customer prospects, and presenting at conferences

Maginatics (acquired by EMC) Staff Engineer 2013 - 2014
Senior Engineer 2011 - 2013

- Founding engineer for MagFS 1.0-2.6, a cloud-based network file system which evolved into EMC CloudBoost
- Lead developer of object data storage, SQL metadata storage, deduplication, garbage collection, backups, and snapshots
- Worked with the Apache jclouds community to improve performance and compatibility across 15 different object stores
- Addressed database concurrency, management, and space-efficiency issues and contributed to upstream MySQL projects
- Recruited and mentored engineers, supervised contractors, and interacted with customers, investors, and partners

Riverbed Technology Member of Technical Staff 2006 - 2011

Storage engineer for SteelFusion 1.0, a virtualized, WAN-optimized iSCSI storage appliance

- Contributed to core indexing structures, transactional persistence and recovery of data, and overall system performance
- Established internal dogfood setup, deployed initial customer installs, and evangelized failure-injection testing

File system engineer for Atlas 1.0, a deduplicating storage appliance which evolved into NetApp AltaVault

- Implemented algorithms for deduplication and data structures for efficient indexing and storage of compressed data
- Improved data segmentation, snapshot, and garbage collection algorithms, increasing performance and space efficiency
- Provided 2-10x speedups via algorithmic improvements, assembly-level optimizations, and implementation tuning

Tabernus (acquired by Blancco) Software Engineer 2004 - 2006

Free software

- Apache jclouds committer and PMC member (2013 - present), Vice President (2017), and GSoC mentor (2014, 2017)
 - Implemented the B2 provider, overhauled filesystem provider, enhanced multi-part upload, and expanded the portable abstraction across seven providers
 - Reviewed more than 100 pull requests by over 40 authors and individually resolved hundreds of issues
- Contributor to s3fs including over 100 commits which doubled write performance, removed head-of-line blocking, addressed data loss issues, automated CI tests, and improved POSIX compatibility, S3 interoperability, and documentation
- Author of S3Proxy, a virtual object store allowing access to various storage backends via programmable middlewares
 - Project has more than 25 contributors and over 500,000 Docker pulls
- Author of asmlint, Chaos HTTP Proxy, Java httpbin, JDBC lint, and Modernizer Maven Plugin
- Contributed to compy, goofys, H2, innodb_ruby, mbpfan, OpenSSH, Percona InnoDB tools, s3-tests, Thrift, and more

Education

University of Texas at Austin

BS Computer Sciences

1997 - 2002

- Interned at HRL Laboratories, Trilogy, and United Devices and did real-time systems research with Dr. Aloysius Mok
- Received an honorable mention at the 2001 ACM International Collegiate Programming Contest