

Aaron Kuehler

Platform and infrastructure engineer with 15+ years shipping software and growing teams. I build the CI/CD, container, and data platforms other teams ship against — Kubernetes, Ruby on Rails, Go, Redis at scale. I'm accountable for outcomes and build autonomous teams where no one is a single point of failure.

Proficiency highlights

- Platform & Infrastructure Engineering (Kubernetes, containers, CI/CD, deployment automation, and workload orchestration — Go, Bash, Docker, Jenkins)
- Distributed data systems at scale (Redis, Elasticsearch, InfluxDB)
- Web application & API development (Ruby on Rails and standard web technologies — HTTP, JavaScript, CSS)
- Open source (22 merged pull requests into Shopify/shipit-engine and an upstream HAProxy feature; maintainer of a Kubernetes Redis operator; upstream bug reports and fixes across the ecosystem — RuboCop, k9s, and the Percona operator)
- Team growth and leadership (Extreme Ownership, agile and Lean principles, mentor and educator)

Recent Accomplishments

Power Home Remodeling

Lead Software Engineer

August 2025 - Present

Promoted to lead the internal platform team that powers the company's CI/CD, container platform, and shared data infrastructure — the systems every engineer in the org ships against.

Setting a 2-year direction that shrinks operational surface area: drove the platform roadmap toward a simplified recommended application stack, a versioned deployment-template library replacing copy-paste manifests, and cluster-scoped operators, collapsing per-app maintenance into a single consistent model. The result: fewer bespoke systems to babysit and a lighter on-call rotation.

Maintain the paved road itself: the company's bespoke continuous-delivery pipeline — built on Shopify's shipit-engine — and the golden-path reference application new services start from. Drive architectural direction through written RFCs, so decisions are debated in the open and outlast any single owner.

De-risking the company's most consequential platform migration: leading the cutover of 5+ applications (the flagship and its satellites) from the legacy container platform to its successor — coordinating data migrations across Redis (multiple GB), Elasticsearch (dozens of GB of indexed documents), and a widely-shared InfluxDB instance, plus DNS and change-control across multiple teams, without customer-visible downtime.

Building bus-factor resilience: distributed ownership of critical platform components (cluster-scoped operators, deployment templating, data-center decommissions) so senior teammates now lead cross-team initiatives the team once single-threaded through one person — 6 components that had a single owner now have two or more.

Shaping who joins the team and how they grow: own annual reviews, hiring loops, and team strategy for a 7-person senior platform team operating at the company's highest-stakes layer.

Senior Software Engineer

July 2023 - August 2025

Returned to Power Home Remodeling to join the platform team chartered to let application teams run reliably at scale.

Cut average Redis call latency by ~73% (7.9 ms to 2.1 ms) on the flagship Rails app by leading a Redis fleet upgrade from v3 to v6 — zero customer-facing disruption, 100% of production clusters migrated, and eliminated the commercial Redis Enterprise licensing spend by replacing it with an internally-maintained open-source operator.

Got code into upstream HAProxy: authored and merged the init-state server-configuration feature (commit 50322df, September 2024), closing a years-open issue and fixing a Redis-failover bug that had caused production incidents. The fix benefits every HAProxy user.

Shipped application-layer DDoS protection across the company's most-trafficked public properties, proactively closing a longstanding security gap.

Made deploys and CI faster for the entire engineering org: cut average production deploy time for the flagship app from ~14 to ~10 minutes and worst-case CI time by 30% — compounding daily across every engineer who ships.

Took over and modernized a critical open-source dependency: maintainer of an open-source Kubernetes Redis operator (Go), added multi-site replication, and made cross-data-center disaster-recovery exercises routine rather than a fire drill.

Multiplied the team's reach through mentorship and knowledge-sharing rituals: engineering mentor in the company's internal coding academy; led the internal Refactoring book club; founded two recurring practices still running today — "Maintenance Monday" (team-level) and "Tech Talk Thursday" (open mic across the technology department).

Shopify

Senior Software Developer

September 2021 - May 2023

Shipped the \$2.5 billion Deliverr acquisition integration ahead of schedule — defined the APIs and built the Shopify-side integration that gave Shopify Fulfillment Network (SFN) merchants a seamless experience, coordinating delivery across a geographically distributed team.

Increased Returns system throughput by 500% ahead of the SFN Returns 2022 Black Friday / Cyber Monday season. Ran capacity and resilience testing to identify choke-points. Led disaster recovery drills and defined on-call and incident response policies.

Discovered silent failure modes in warehouse management software and added observability instrumentation.

Used machine learning to build a product taxonomy that both improved how related products are surfaced to shoppers and optimized warehouse placement for faster picking and packing.

Identified bottlenecks in feature gating system. Scaled feature flag management capabilities from 10 to 2 million merchants. Added percentage-based feature rollouts, feature lifecycle automation, enhanced human-consumable feature metadata.

Grew 2 interns into autonomous Individual Contributors in weeks through pairing, synchronous 1:1 support, and defining right-sized work items.

Power Home Remodeling Group

Senior Software Developer / Substitute Instructor / Mentor

May 2018 - September 2021

Slashed deployment times by 82% by transforming the deployment topology to self-hosted, Kubernetes-based container orchestration.

Became a sustained upstream contributor to Shopify's open-source deploy tool, shipit-engine — 22 merged pull requests from 2019 to 2021, as an outside contributor, spanning domain modeling (elevating the repository to a first-class concept, renaming PullRequest to MergeRequest, extracting webhook handlers) and maintenance-grade internals (task timing on a monotonic clock, git-cache-corruption recovery, GitHub push-event handling). Primary author of its Review Stacks feature (PR #1102): ephemeral, full-stack preview environments spun up per pull request — ~8,400 lines across 106 files, delivered with a small team and reviewed by Shopify maintainers.

Took Talkbox — our in-house telephony system — from ~5 daily outages to one: it was taking the call center offline multiple times a day until I identified the constrained resources and decoupled/distributed the system by call type, keeping it online through the workday.

Expanded development organization capacity by 20% by growing "Power Code Academy" student developers into Individual Contributors. Functioned as a substitute instructor and post-graduation mentor.

Bauer Xcel Media

Technical Team Lead / Senior Software Developer

February 2016 - May 2018

Boosted ad revenue by ~200% by optimizing impression value and advertising rank. Deferred ad load requests until ad slots were just about to scroll into reader's view.

Cut non-production infrastructure spend by ~50% by moving stakeholder preview and review environments from always-on to on-demand — spun up per change and torn down after use, instead of running continuously.

Reduced work-item cycle-time by ~34% by encouraging adoption of pairing and building a healthy code review culture.

Increased site traffic by ~300% by optimizing metadata for search engine and social media ranking.

exercism.io

Volunteer Mentor

August 2019 - September 2021

Mentored over 60 students along the Ruby and Emacs-Lisp language tracks.

Example Open Source Contributions

Shopify/shipit-engine

<https://github.com/Shopify/shipit-engine>

Shipit is a deployment tool that makes shipping code better for everyone. It's especially great for large teams of developers and designers who work together to build and deploy GitHub repos. I have 22 merged pull requests here, including primary authorship of the Review Stacks feature.

percona-xtradb-cluster-operator

<https://github.com/percona/percona-xtradb-cluster-operator>

Percona's Kubernetes operator for XtraDB Cluster. I contributed a fix merged upstream (PR #521) that resolved a crash when the operator pod started without a HAProxy PodDisruptionBudget.

rubyforgood/mutual-aid

<https://github.com/rubyforgood/mutual-aid>

Mutual Aid is when people get together to build community by voluntarily sharing resources with each other. Mutual Aid groups are more concerned about local resiliency than global campaigns, and prefer solidarity before charity. I contributed 10 merged pull requests, largely hardening the Docker-based development environment for new volunteers.

Sample Side Projects

asdf-guile

<https://github.com/indiebrain/asdf-guile>

An asdf version manager plugin to manage GNU Guile installations

resque-alive

<https://github.com/indiebrain/resque-alive>

A Kubernetes Liveness probe for Resque

Talks

Git Internals

https://github.com/indiebrain/talks/blob/master/git_internals/git_internals.org

Explains the inner workings of git and sheds a bit of light on how the internal structure of git as a data store influences its user interface.

Git Workflow

https://github.com/indiebrain/talks/blob/master/git_workflow/git_workflow.org

Treats version-control history as project metadata: a workflow built for reason-ability (what, exactly, is in production?), communication, and a semantic commit history that records the intent code alone can't.

Release It!

https://github.com/indiebrain/talks/blob/master/release_it/release_it.org

A brief, and messy, history of releasing software — from ENIAC and the 1968 "software crisis" to how modern practices arrived at continuous, lower-risk delivery.

Elsewhere

- <https://aaronkuehler.com>
- <https://www.github.com/indiebrain>
- <https://www.linkedin.com/in/aaronkuehler/>

Education

West Chester University of Pennsylvania

Bachelor of Science, Computer Science with Information Assurance Minor

January 2006

- Magna Cum Laude
- Award for Academic Excellence (2006)
- Dean's list (2005 and 2006)

Research

- Small File Affects on Hadoop Distributed File System
 - Whitepaper PDF

The Hadoop Distributed File System is a high throughput distributed File system designed to accommodate large data sets; average file sizes in the gigabyte-terabyte range. However when a data set is composed of large amounts of small files, say in the kilobyte range, the storage system's semantics introduce high amounts of overhead in terms of file system block storage and read latency. This paper explains the architectural attributes which cause these problems and examines techniques to mitigate their impact when working with data sets comprised of large numbers of small files.