

Jay Fallon

Founding Engineer · Platform Architect · Full-Stack Engineer

Boston, MA · 617-406-8027 · jayfallon@gmail.com · jayfallon.net · github.com/jayfallon

SUMMARY

Principal-level engineer who designs and operates multi-service production platforms end-to-end. Currently building Knokr, a music-industry platform managing **60,000+ artists, 1,500+ festivals, 1,000+ venues, and 25,000+ events** across a federated architecture: three Next.js apps, a Python/FastAPI prediction service, and four BullMQ background workers sharing a single Postgres/Redis/S3 backbone.

Twenty-plus years of engineering experience across enterprise (Dassault Systèmes, Bose, State Street, Cengage) and consumer health (Form Health). Deep on backend architecture, data modeling, and applied ML — pgvector and HNSW indexes, graph algorithms (Node2Vec, Louvain), AI-powered recommendation systems — with a long front-end track record that informs how I build APIs and design systems.

Seeking founding engineer, principal IC, or staff backend roles where one engineer can own architecture and ship across the stack.

SELECTED OUTCOMES

- **Built a multi-service platform from zero to production solo.** Multiple Next.js apps and Python/FastAPI services, and BullMQ workers, multiple Postgres/Redis/S3 backbone instances. 3,000+ tests across the system to run both the user-facing Knokr.com and its backend.
- **Production-scale data ingestion.** Designed and built Knokr Fetch, a Python-based pipeline tool designed to aggregate and consolidate data from multiple sources to include festival and venue websites along with popular musical event services and custom parsers for bespoke websites, populating 60K artists, 1.5K festivals, 1K venues, and 25K events for 2026.
- **Graph-based discovery on 3.3M edges.** Knokr Predictor — a FastAPI service running Node2Vec embeddings and Louvain community detection over festival co-occurrence data, with weighted scoring for emerging-artist surfacing.
- **Sub-100ms semantic search at scale.** pgvector with HNSW indexes powering CLIP image search, text embedding search, and "more like this" recommendations across the platform.
- **Six-tier RBAC with Redis-cached permissions.** 90–95% latency reduction on permission checks; three-layer route protection (middleware → layout → component).
- **Senior IC track at Dassault Systèmes.** Four years as a hands-on specialist working alongside director-level engineers on the SOLIDWORKS product suite — front-end platform consistency, a new documentation CMS, and consumer-facing UI work across multiple products, including the design, implementation and maintenance of multiple internal and consumer-facing products.

EXPERIENCE

Founder & Principal Engineer · Knokr

2024 – Present

Sole architect and engineer of a federated music-industry platform spanning a consumer-facing application and an operator/CMS backend. Chose multi-service over monolith from day one: separate Next.js apps, a Python/FastAPI prediction service, BullMQ workers, all sharing multiple Postgres/Redis/S3 layers.

- **Consumer application:** Brings together artists, venues, festivals, and fans in geographic context — helping festivals and venues find their audiences, and enabling fan-to-fan interaction. Includes

geographic and semantic discovery, animated festival map, dream lineup builder, personal calendar with RSVPs, an extensive media library (photo/video albums tied to entities, CLIP-based visual search), and user-to-user direct messaging with location-tagged threads.

- **Operator platform & CMS:** Handles all data sourcing, hygiene, and management for the music-industry domain. Manages organizations and individual accounts, plus the projects within them — venues, festivals, events, and websites. Includes a dynamic content management system that can host venue and festival sites either as shared tenants on the main domain or as standalone branded sites with custom domains.
- **Backend & data:** Designed the Postgres schema for a complex music-industry domain — entity resolution across artists, venues, festivals, and events; slug stability; ownership transfer; multi-tenant isolation; six-tier RBAC with three-layer route protection. pgvector + HNSW for semantic search. Denormalized rating rollups, composite indexes, and Redis caching delivering 90–95% latency reductions on hot paths.
- **Background processing:** BullMQ workers for graph computation, CLIP embeddings (with NSFW + near-duplicate classification), insights aggregation, and social-card generation. Standalone Railway deploys communicating over HTTP with service-token auth.
- **Applied ML:** Node2Vec + Louvain on 3.3M co-occurrence edges for artist discovery and "scene" detection. CLIP image embeddings via @xenova/transformers for visual search and near-duplicate detection. OpenAI text embeddings for festival/artist semantic search. Anthropic Claude API for the AI festival-recommendation engine and lineup extraction from festival posters.
- **Data ingestion:** KnokrFetch — Python sourcing system pulling from traditional event services and bespoke websites. Populated 60K+ artists, 1.5K+ festivals, 1K+ venues, and 25K+ events for 2026.
- **Infrastructure & ops:** Railway deploys, AWS S3 + CloudFront, Sentry (errors, performance, session replay), Winston structured logging, GitHub Actions CI/CD with pre-commit, pre-push, and CI gates.
- **Test coverage:** Designed and own ~2,800 tests across the admin CMS and public site (Vitest, Testing Library, Playwright, Testcontainers, Supertest, Memlab). Integration suite hits a real Postgres via Testcontainers — no DB mocking — covering artist normalization, duplicate detection, geocoding, and event-request provisioning end-to-end. Unit suite covers Zod validation schemas, RBAC permission gates, and HeroUI form flows; component suite uses Testing Library against rendered React Server Component output under Happy-DOM.
- **Front-end:** Next.js 16 / React 19 / TypeScript / Tailwind v4 / HeroUI. Database-driven routing, custom subdomain handling, golden-ratio image cropping, animated Leaflet festival map, drag-and-drop lineup builder.

Tech: PostgreSQL/Prisma, pgvector, Redis, BullMQ, Node.js, Python/FastAPI, asyncpg, scikit-learn, Pandas, NumPy, Selenium, BeautifulSoup, Next.js 16, React 19, TypeScript, Tailwind, Clerk, Stripe, AWS S3/CloudFront, Sentry, Railway, Anthropic Claude API, OpenAI API.

Consulting Engineer · Form Health

2022 – 2024

- **Front-end rewrite:** Rewrote the front-end from scratch to incorporate the site redesign. Developed a componentized design system library for implementation across the platform.
- **Marketing & product partnership:** Worked with marketing and product teams on the customer-facing platform for a medical weight-loss company; built and maintained the marketing site infrastructure and shared component library.
- **Design system tooling:** Implemented Storybook documentation and supporting tooling to standardize component usage and reduce front-end fragmentation across teams.

Tech: Next.js, TypeScript, Tailwind CSS, Prismic, Storybook.

Senior Software Engineering Manager · Dassault Systèmes SE · Waltham, MA

2021 – 2023

Senior IC role despite the manager title — hands-on engineer working alongside director-level peers on the SOLIDWORKS Cloud Platform.

- **Front-end consistency:** Implemented and maintained UI features across products to enforce design-system consistency and reduce accumulated technical debt.
- **Documentation CMS:** Designed and built a new content management system serving the entire product suite, replacing fragmented per-product docs.
- **New product partnership:** Worked with new product teams on consumer-facing UIs while keeping product and design language consistent across the portfolio.

Software Engineer (Contract) · Dassault Systèmes SE · Waltham, MA 2019 – 2021
Contractor working across two separate products in the SOLIDWORKS Apps for Kids suite.

SOLIDWORKS Apps for Kids 2019 – 2021

- **Front-end rewrite:** Rewrote the entire front-end in Angular 1.x, JavaScript, and CSS. Developed a componentized design system to unify the separate applications in the product suite.
- **API consolidation:** Consolidated disparate API streams into one unified structure for easier maintenance across the suite.
- **End-to-end testing:** Rewrote the end-to-end test suite to expand unified test coverage across the product suite using Mocha, Chai, and Puppeteer.

SOLIDWORKS Apps for Kids — Classroom 2019 – 2021

- **Sole engineer on Classroom:** Teacher-facing application for 30,000 students aged 8–14 — classroom and assignment scheduling, grading tools, custom newsletter app, COPPA-compliant data and account flows.
- **Local-dev orchestration:** Containerized an Express suite of 12 interdependent services (auth, calendar, mail, etc.) using Docker Compose, replacing manual per-service startup.
- **Front-end rebuild:** Introduced a design system the application had never had. Designed and implemented every component in HTML/JS/plain CSS against static designs, on a Dassault-licensed customized Angular 1.x build.

Tech: Angular 1.x (Dassault-licensed build), JavaScript, HTML, CSS, Node.js, Express, Docker, Docker Compose, Mocha, Chai, Puppeteer, Cassandra.

Software Engineering & Consulting · Earlier Roles 2000 – 2019

Engineering and consulting roles spanning financial services, education, consumer audio, and pharmaceuticals. Work included full-stack web development, design system architecture, and front-end engineering at scale. Stack depth from this period includes **six years of Angular.js**, **four years of Ruby on Rails**, **multi-year Express.js applications**, and the full pre-React CSS lineage (CSS, SCSS, PostCSS, Stylus). Clients and employers via direct hire and consulting: **Bose Corporation**, **State Street Global Advisors**, **Cengage Learning**, **Boston Technologies**, **netNumina Solutions** (Fidelity, State Street, Johnson & Johnson, Pfizer, ADQSR), and **Style Me Pretty**.

TECHNICAL SKILLS

Backend & Data: PostgreSQL, Prisma, pgvector + HNSW, Redis, BullMQ, Node.js, Express, Python, FastAPI, asyncpg, GraphQL, REST, WebSockets, multi-tenant systems, database-driven routing, microservices, Cassandra, MongoDB, MySQL, Supabase.

Applied ML & Search: Anthropic Claude API, OpenAI API, vector embeddings (CLIP, all-MiniLM, OpenAI text-embedding-3), graph algorithms (Node2Vec, Louvain community detection), semantic + full-text search, scikit-learn, Pandas, NumPy.

Infrastructure & DevOps: Railway, AWS (S3, Lambda, CloudFront), Vercel, Docker, GitHub Actions CI/CD, Sentry, Winston, Husky pre-commit/pre-push gates, Testcontainers.

Data Pipelines & Automation: Selenium, BeautifulSoup, Puppeteer, asyncio, custom site parsers, data validation pipelines, scheduled BullMQ jobs.

Testing: Vitest, Playwright, Testing Library, Testcontainers, Supertest, Memlab, Happy-DOM.

Frameworks & Front-End: Next.js 16, React 19, TypeScript, Tailwind, HeroUI, Framer Motion, PayloadCMS, Prismic, Storybook.

Prior Stack Depth: Angular.js (6 years), Ruby on Rails (4 years), Express.js applications across multiple engagements, plus the full pre-React CSS lineage (CSS, SCSS, PostCSS, Stylus).

Auth, Payments & Misc: Clerk, Zod validation, Stripe (checkout, billing portal, webhooks), Firebase, Agile.

Languages: Fluent in English and Spanish.

EDUCATION

Master of Liberal Arts (ALM), Digital Media Design · Harvard University 2017 - 2019

Harvard Extension School. Coursework in user experience design, interaction design, and digital product development.

B.S. in Business Administration (International Business, Finance) · Northeastern University 1993 - 1998

Northeastern's cooperative education program — placements at Boston's Better Snacks, FMRC, and PictureTel.