



Aidan Foley, BSc

Founder-Engineer & Tech Leader

Founder-engineer and data/platform leader building cloud-native systems, ML products and high-trust data foundations for complex, data-intensive businesses at any scale.



This CV is interactive at aidanfoley.dev — live architecture diagrams, the products themselves, and more. · [LinkedIn](#)

00 PROFILE

Founder-engineer and technical leader building data platforms, ML/statistical systems, cloud-native backends and product experiences from zero-to-one through to production scale. Equally at home as a hands-on founder or a team lead — across start-ups, scale-ups and established enterprises — holding a consistent bar for data correctness, operational resilience and product quality throughout.

Currently Head of Data & Engineering at Baller League, leading a five-person function across data, backend, mobile and computer vision. In parallel, founder of Playlistn — an operating system for music discovery — built solo across the full stack: event-driven ingestion, ML workflows and a consumer-grade product experience.

01 CORE OPERATING MODES

Technical leadership	Team building & hiring · Delivery practice · Vendor/partner management · Architectural decision-making · Cross-functional communication · Product thinking
Engineering	Full-stack development · Next.js · React · FastAPI · Go microservices · Native iOS (Swift/SwiftUI) · Analytics instrumentation · AI-leverage operating model
Data platforms	Event-driven pipelines · Ingestion & contracts · dbt modelling · Data lineage & governance · Lambda + Batch architectures · Snowflake · Athena · BigQuery
ML & statistics	Forecasting · Anomaly detection · Recommender systems · Feature engineering · Model validation & explainability · MLflow
Cloud & operations	AWS · GCP · Terraform (IaC) · ECS/Fargate · Lambda · CI/CD & trunk-based delivery · Cost-aware scaling · Monitoring & incident management

02 SYSTEMS SHIPPED

01 Live Sports Technology Platform

Baller League · Head of Data & Engineering · Jul 2025 – Present

30k+ concurrent users · 0% errors

5.9ms p95, matchday hot path

Led a five-person function and owned the rebuild from a partner-coupled Laravel monolith to a stateless FastAPI + Next.js platform on AWS. Beyond the platform, the consumer matchday app went from our initial product sketch to the App Store with a small team, a lean budget and a rapid turnaround. On matchday the platform ingests tens of thousands of Opta XML snapshots in near-real-time (~1–3s feed-to-product) and serves complex live data and assets to tens of thousands of concurrent fans at a 5.9ms p95 — validated at 30,020 concurrent users with a 0% error rate.

ARCHITECTURE — A Lambda Architecture on AWS: near-real-time matchday ingestion fans every Opta feed into a low-latency serving lane (zero-SQL hot path) and a parallel analytics lane, behind a stateless API serving native and web clients.

- Made the defining architecture call early — keep the live serving path sacred, separate from analytics, so matchday can never be compromised — then proved it at launch: 30,020 concurrent fans, 5.9ms p95, zero errors.
- De-risked the highest-stakes change in the business: replacing the partner-coupled monolith mid-season via contract-driven parallel cutover with parity tests, feature-flag phasing and a rehearsed rollback — first matchday shipped with zero customer-impact incidents.
- From concept to App Store: defined the product, designed the consumer experience, and shipped the native iOS matchday app (Swift/SwiftUI) — small team, lean budget, rapid timeline, hardened through a multi-round TestFlight programme.

- Walked the SLT back from an externally-developed computer-vision moonshot, re-scoped it properly — then beat the original promise: replaced a cloud-GPU clip pipeline with an on-device, real-time vision SDK (custom detection + pose + tracking, ~30fps on-phone) that assesses players performing drills live — event recall lifted from 83% to 98% on a 159-video blind evaluation.
- Built and led the five-person function across data, backend, mobile, CV/ML and product — instilling a rehearse-the-failure operating culture: runbooks, synthetic canaries, and emergency levers that non-engineers can pull mid-match without a deploy.
- Held partners to account: co-led a vendor audit and legal proceedings with the COO that recovered a ~\$450k settlement, and won back \$100k of AWS credits through direct appeal.
- Started a cross-functional Analytics & BI Community of Practice — making data a shared language across the business rather than an engineering silo.

Stack: Python, AWS, FastAPI, Next.js, Aurora Postgres, Redis, Athena, dbt, Terraform, Swift/SwiftUI, C++17, Firebase, GA4

02 Music Discovery Platform

Playlistn · Founder, CEO & CTO · May 2024 - Present

8M+ artists tracked & refreshed daily
multi-source live data ingestion

Founded and built playlistn — an operating system for music discovery. A&R teams connect their Spotify and streaming accounts once; the platform ingests their world, learns their taste from every play, save and signing, and keeps tracking live as new artists catch their attention. Discovery becomes effortless: search in plain English, get a shortlist in seconds, explore a living map of 8M+ artists refreshed daily — and carry the whole journey, found to signed, in one tool.

Underneath: live multi-source ingestion, a self-growing artist graph, and a custom AI search engine tuned so precision stays high at a fraction of a cent per query. Bootstrapped solo on £30k.

ARCHITECTURE — *An event-driven enrichment mesh aggregates multi-source data into a normalised artist graph, served through a custom AI search and a Next.js product, built solo with strong engineering discipline.*

- Built the entire platform solo — by design, not necessity: a deliberate AI-leverage operating model that turns a founder budget into multi-engineer output, with the engineering discipline of a much larger team (ADR-driven decisions, contract tests, CI quality gates).
- Designed the AI search around product economics: natural language resolved to typed, deterministic execution, with retrieval techniques that hold precision high at fractions of a cent per query — so the product scales on bootstrap economics.
- Built a compounding data asset: a background similarity crawler grows the artist graph far beyond any single user's catalogue — the moat deepens around the clock, with priority-fair scheduling so growth never degrades the customer experience.
- Owned every layer of the product myself — from event-driven ingestion to an immersive 3D discovery interface and a full A&R workflow — making the product/engineering trade-offs a founder has to make daily, in both directions.
- Run it as a real company, not a side project: Playlistn Ltd incorporated, UK trademark registered, SOC 2 Type II in progress, SEIS/EIS seed round underway.

Stack: Python, FastAPI, Next.js, TypeScript, PostgreSQL, AWS, SQS, Lambda, EventBridge, DynamoDB, Three.js, OpenAI, Terraform

03 Royalty Anomaly Detection Service

Claimy · Principal Data Scientist (fractional, fixed-term) · Jul 2025 - Jan 2026

Layer 1 fast detections · inline at ingestion
Layer 2 smart detections · nightly ML batch

Fractional engagement alongside Baller League. Paired with the CTO to develop technology strategy and roadmap, then built a two-tier anomaly-detection system. Layer 1 — "fast detections" — runs deterministic SQL rules embedded directly in the royalty-processing pipeline at statement-ingestion time, catching known error patterns with zero added latency. Layer 2 — "smart detections" — runs nightly on Cloud Run, aggregating the full time-series picture and running a comprehensive ML-based classification framework to surface anomalies that only become visible at scale or over time.

ARCHITECTURE — *A GCP-native, two-tier detection system layered over the royalty-processing pipeline: deterministic rules fire inline at ingestion, while a nightly ML batch surfaces systemic anomalies that only appear at scale.*

- Brought in as a trusted fractional partner to the CTO — co-developing the long-term technology strategy, infrastructure plan and roadmap, not just delivering hands-on work.

- Designed detection as a product decision, not a model choice: deterministic rules fire inline where instant feedback matters (zero added latency); deep ML runs nightly where wider context matters — every anomaly caught at the cheapest layer that can catch it.
- Carried proven thinking across companies: extended the anomaly-framework concepts I built at Recognition into a GCP-native architecture — landing production value inside a fractional engagement run alongside Baller League.
- Took the product from sketch to production UI myself: prototyped the detection review experience in Plotly Dash, then ported it with the team into the production React app — the same full-stack product thread carried through Playlistn and Baller League.

Stack: Python, Go, BigQuery, GCP, Cloud Run, Firestore

04 Royalty Intelligence & Forecasting

~\$120M anomalous income identified · first run
~30% forecast RMSE improvement

Recognition Music Group · Lead Data Scientist · Jan 2025 – Jun 2025

Built two proprietary data products for music publishing and recording revenues — a forecasting framework deployed in three months and an anomaly-detection system that identified \$120M of outlier royalty income on first run.

ARCHITECTURE – *Two proprietary frameworks — one for revenue forecasting, one for anomaly detection — built over governed royalty data and productionised on Snowflake for repeatable analytics.*

- Challenged the incumbent song-decay assumptions and rebuilt the forecasting methodology — a ~30% RMSE improvement, taken from first commit to production in three months.
- The anomaly-detection system surfaced ~\$120M of anomalous royalty income on its very first run — turning a sceptical first audience into the system's sponsors, and anomaly detection into a board-level capability.
- Productionised both frameworks on governed data so the analysis is repeatable by the team — durable capability, not key-person dependency.
- Founded a Forecasting Community of Practice to align investment and data teams on methodology and evaluation — the social infrastructure that makes models trusted.

Stack: Python, Snowflake, DuckDB, PostgreSQL, AWS

05 A&R Recommendation Pipeline

~100 engineers & analysts on data platform
~0.8 FTE saved via Python automation

Kobalt Music Group · Royalties Analyst → Data Analyst → Data Scientist → Senior Data Scientist · May 2020 – Dec 2024

Progressed over four years from royalty operations into analytics and data science, leading the discovery and MVP delivery of an explainable A&R recommendation pipeline and contributing to a governed data platform supporting ~100 engineers and analysts. Kobalt's engineering standard is exceptionally high — I was fortunate to be mentored there by some incredible data and technology leaders, and everything since is built on those foundations.

ARCHITECTURE – *An explainable A&R recommendation MVP over a governed data platform: tens of thousands of temporal features feed a transparent tree-ensemble re-encoded in SQL for auditability.*

- Chose explainability over leaderboard accuracy: delivered the A&R recommender as a transparent tree-ensemble re-encoded in SQL, so analysts could audit every recommendation — trust was the feature.
- Led discovery the user-centred way (CRISP-DM): iterated with A&R stakeholders until the model surfaced long-term artist growth signals they actually believed in — adoption first, sophistication second.
- Automated the royalties grind I started in — Python tooling that saved ~0.8 FTE — and turned that initiative into a four-year runway from analyst to senior data scientist.
- Contributed validation patterns and metadata governance to a data platform serving ~100 engineers and analysts — learning what "governed" means at real scale.
- Founded the internal AI Taskforce and co-led a business-run automation team — bringing generative AI and productivity tooling to teams far outside engineering.

Stack: Python, SQL, PySpark, dbt, AWS, Tableau

03 CAREER TIMELINE

Jul 2025 - Present	Baller League	Head of Data & Engineering
May 2024 - Present	Playlistn	Founder, CEO & CTO (part-time alongside employment)
Jul 2025 - Jan 2026	Claimy	Principal Data Scientist (fractional fixed-term contract)
Jan 2025 - Jun 2025	Recognition Music Group	Lead Data Scientist
Apr 2024 - Dec 2024	Kobalt Music Group	Senior Data Scientist, Data Platform & Analytics
May 2023 - Apr 2024	Kobalt Music Group	Data Scientist
Jun 2021 - May 2023	Kobalt Music Group	Data Analyst, Tracking & Analysis
May 2020 - May 2021	Kobalt Music Group	Royalties Analyst

04 CORE STACK

Languages	Python · TypeScript · SQL · Go · Swift · C++
Backend	FastAPI · Aurora Postgres · Redis · RDS Proxy · Alembic
Data	dbt · Athena · PySpark · Snowflake · BigQuery · DuckDB
Frontend	Next.js · React · Tailwind CSS · Three.js / R3F · Framer Motion
Mobile & CV	Swift/SwiftUI · C++17 SDK · TensorFlow Lite
ML, AI & CV	MLflow · Label Studio · OpenAI API · Forecasting · Anomaly detection · Recommenders
Cloud & infra	AWS · GCP · Terraform · ECS/Fargate · Lambda · SQS · EventBridge · DynamoDB · CloudFront
Auth & analytics	Firebase Auth · Firebase Analytics · GA4 · GTM
Quality & obs.	GitHub Actions · Playwright · pytest · Sentry · X-Ray · dbt tests

05 EDUCATION

University of Leeds · BSc Theoretical Physics · 2015 – 2018

Stanford University (Online) · Machine Learning Specialisation · 2024

Supervised learning, advanced algorithms, unsupervised learning, recommenders & RL