Sean Dugan

Senior Software Engineer | Backend-Heavy Full-Stack | Scala · Akka · React · Python
Austin, TX | lunacity360@gmail.com | linkedin.com/in/sean-dugan-604a545a

SUMMARY

Senior backend and full-stack engineer with 11+ years building distributed JVM systems, most of them in functional Scala, including Dun & Bradstreet's Hoovers / DataVision B2B data platform. Deep in Scalaz and Cats, Akka (Typed, Cluster, Persistence), Play, and sharded PostgreSQL, with React / TypeScript front ends. Designed a KMS-backed hybrid JWT session platform and a Redis Streams usage-metering system, and shipped monolith features as production agent-callable tools in a shared Python MCP codebase. Carried technical-lead duties for the DataVision project, making backend architecture decisions, running scrums and onboarding new hires; experienced with AI-assisted development (Cursor, Codex, Claude, Gemini).

SKILLS

Languages: Scala (2 & 3), Java, TypeScript / JavaScript, Python, SQL; familiar with Node.js / Next.js
AI & Agentic: Cursor, OpenAI Codex, Claude, Gemini; Python MCP (Model Context Protocol) services
Backend & FP: Scalaz, Cats, Akka (Typed, Cluster, Persistence), Play, Lift, Slick, pac4j, HikariCP, REST, GraphQL
Data & Messaging: PostgreSQL (sharded, SKIP LOCKED queues), Redis Streams, Elasticsearch aggregations, AWS S3 (SDK v2), ETL pipelines
Security: JWT / JWKS, RS256, AWS KMS, OAuth 2.0, SSO, token revocation, HttpOnly session cookies, CVE remediation
Frontend: React (TypeScript, RTK Query, React Testing Library), Ember, Vue, Chart.js
Infra & Tooling: AWS, GCP, systemd / RPM, Chef, ELK, sbt, scalafmt / scalafix, Git
Integrations: Salesforce (REST / OAuth), HubSpot (CRM v3), Marketo, Segment CDP

EXPERIENCE

Dun & Bradstreet · Austin, TX **Mar 2018 – Oct 2026**

Senior Software Engineer, Hoovers / DataVision

Senior Software Engineer (2021 – 2026) · promoted from Software Engineer (2018 – 2021)

Core engineer on Hoovers / DataVision, a B2B data and analytics SaaS platform that was one of D&B's most profitable lines, uniquely profitable for several years. Owned work across the Scala backend, data pipelines, CRM integrations, and Ember / React front ends.

- **Carried DataVision's technical-lead duties:** ran scrums and standups, made the backend architecture decisions, led critical bug fixes, and became the owner of the codebase and its customers; onboarded and mentored new hires, including a graduate-school recommendation for a mentee.
- **Visitor Intelligence ingestion:** first big DataVision project: designed and owned the application-side S3 loaders, the stable workhorse ingesting canonicalized data that another team's upstream AWS Lambdas aggregated; extended them into a loader framework (Prospect, Lattice, Segment CDP; delta / zip / gzip, rediscovery, priority queue, retention janitor) with daily monitoring reports, ending as the only engineer left on the product.
- **Connector framework:** designed and built the framework behind every DataVision export type, from CRM connectors (Salesforce, HubSpot, Marketo) with per-connector, per-product scheduling config to scheduled exports, including a solo weekend rewrite before AI coding tools; iterated on it for years, with deep knowledge of the complex monad-based export pipeline.

Auth, Metering & Platform Modernization (2023 – 2026)

- **Agentic MCP services:** migrated monolith features into a shared agentic Python MCP (Model Context Protocol) codebase, exposing platform capabilities as production agent-callable tools.
- **JWT authentication platform:** designed, built and supported a hybrid JWT session architecture spanning the legacy Lift webapp and the new Play service, which became the auth layer every production MCP service and AI agent relies on. RS256 signing goes through AWS KMS with hybrid local signing, a /.well-known/jwks.json endpoint resolves the latest enabled key version, and JWKS keys are cached for an hour. Migrated Lift tokens from HS256 to RS256 and unified issuance across SSO, CRM redirect, account switching, and restored sessions behind a secure /token endpoint with HttpOnly cookies.
- **Auth hardening:** hardened auth and closed security gaps with persistent token revocation via an Akka Cluster PersistentActor blacklist tied to logout; a dedicated JWT dispatcher, timeouts, and an admin safety toggle; remediated CVE-2026-29000 (pac4j auth bypass) by upgrading to pac4j 6.3.3 with tests that reject none-algorithm tokens; blocked cross-site list subscriptions and reserved-subdomain UUID collisions.
- **Usage metering:** built the Records Usage billing-metering system for contact reveal, CRM export, and financial-health and report downloads: an Akka Cluster router with per-site workers (no cross-site head-of-line blocking), loss-free passivation, synchronous capacity reservation, chunked ON CONFLICT bulk writes, and audit loss markers.

- **Usage hand-off:** re-architected the hand-off twice, first a lease-based Postgres SKIP LOCKED queue (batched acks, idempotent redelivery, depth / liveness metrics), then replaced it with a Redis Streams producer and poller behind a kill switch.
- **Metering admin UI:** built the Records Usage ops admin UI in React / TypeScript with RTK Query, shared form and table components, and React Testing Library tests, backed by a Scala admin service with parallelized summary reads.
- **UUID migration:** migrated the monolith's sharded PostgreSQL tables from legacy founder-era varchar IDs to native UUIDs, spanning billions of rows, using a resumable migration program (ctid-to-pkid pagination, intra-shard concurrency, partial temp indexes, Slick UUID mapping, feature-toggled cutover).
- **Lift-to-Play rewrite:** laid the foundation for the Scala 3 rewrite with legacy compatibility: pac4j redirect flow and JWT backend, session renewal, Slick codegen with a custom generator, and a thread-safe rate-limit filter; built a multi-database shard manager and sharded user / site sync with deterministic UUIDs to reach 1:1 auth parity with Lift.
- **Resiliency hardening:** in the AI-assisted era, drove adoption of circuit breakers, Guava caches, and backoffs, hardening sensitive synchronous Lift web threads while rewriting them as asynchronous Play endpoints with pac4j security; added export backoff and replica-lag retries and a prepared-statement leak fix.
- **HubSpot CRM v3:** migrated lists to the CRM v3 API behind a real-time feature toggle (dual v1 / v3 list properties, dynamic client selector), with partial-success batch handling, blank-email filtering, and sampled list telemetry; cut SSO-login database load by caching allowed framing URLs with scheduled refresh.
- **Salesforce connector:** built the backend end to end, including OAuth authorize / token / refresh / revoke, companies fields API, listview selection with typeahead, compound-field handling, and China-instance support; started the Salesforce export connector; partnered with a senior engineer on scheduled account imports and fixed Marketo list creation.
- **Runtime and deployment:** rewrote a core subsystem on Akka Typed; moved the service from Tomcat to a systemd-managed RPM with Chef-managed config, partnering with a senior engineer on the per-environment cluster rollout; built a lead-node cluster project with a React cluster-status page; upgraded to AWS SDK v2 with hardened S3 timeouts and retries.

Data Platform, Integrations & Analytics (2018 – 2022)

- **DataVision admin app:** rebuilt the admin application (Ember + Scala API) with loader management, consolidated datamart pages, queues, scheduled jobs, reindexing, and permission-checked admin routes; introduced bulk export with job-tracker stats and zip / CSV exports; hardened industry-upload CSV validation, partnering with a teammate on parse-error reporting.
- **Elasticsearch aggregations:** extended the aggregation layer (2018–2019), moving every DataVision visualization and CSV download onto aggregations: keyed range, reverse nested, date range, and corporate-family aggregations, a nested-aggregation request builder, crosstab / custom views, and Charts CSV export; later added an Elasticsearch-backed distinct-column cache and VI charts in Chart.js.
- **Data and admin tooling:** shipped a Salesforce bulk downloader; the record-usage (RUM) credits admin and site-usage reporting; the analytics-models admin; NAICS 2017 and 2022 industry migrations; a GraphQL product-tag filter separating Hoovers and DataVision users; and sharded DUNS batch lookups that replaced per-record export-credit queries.

FATHOM Water Management · Austin, TX **Sep 2016 – Mar 2018**

Software Engineer II

- **Integrations:** a growth period moving from Java / Spring Boot to Scala, Akka, and functional programming, unlearning Java habits; delivered actor-based REST and SOAP integrations with meter-vendor head-end servers, replacing unstable partner interfaces that were blocking major implementations.
- **Meter data:** scaled high-throughput meter data collection on Akka Cluster with a reactive, fault-tolerant pattern; contributed to multi-tenant alarm and event-notification pipeline reworks on AWS (S3, SNS, SQS, Kinesis, DynamoDB).
- **Cloud platform:** groundwork for a Scala / Nashorn platform running JavaScript services on AWS (S3, SQS, DynamoDB).
- **Partners and deployments:** worked out technical approaches for new partner integrations with partner technical contacts, and handled customer deployments and pilot demos.

DaVincian Healthcare · Austin, TX **Jun 2015 – Sep 2016**

Software Engineering Intern → Software Engineer

- **Messaging and security:** designed pub / sub messaging between microservices on the core Revenue Cycle Management SaaS product (Java, Spring Boot, ActiveMQ / JMS); replaced the legacy security service with OAuth2 / JWT.
- **PNC Bank / TransUnion financial screening:** audits and routes real-time and batch requests (REST, SFTP, Spring, JMS, JPA).
- **Tooling:** ran JMeter stress tests, fixed Hibernate concurrency bugs, moved CVS to Git / Gradle, built HL7 test tooling.

Earlier: Primus Real Estate, Austin, TX — Web Developer / Web Administrator (2012 – 2013): maintained the company's commercial-property website (PHP, HTML, jQuery), did Photoshop and other digital design work, and kept property documents and unit vacancy status current.

EDUCATION

B.S. Computer Science, St. Edward's University, Austin, TX 2015