Matthew Long

Chicago, IL | mlong@magneton.io | 312-478-9956 | LinkedIn | GitHub

PROFESSIONAL SUMMARY

Principal AI & Data Engineer with 20+ years building production software, data platforms and machine-learning systems. Combines Python, SQL, Apache Spark and Databricks expertise with cloud architecture, infrastructure automation, CI/CD and production reliability across AWS, Azure and GCP. Patent holder in secure distributed messaging; acquisition contributor with experience delivering systems at scale.

EXPERIENCE

Principal AI Engineer (Founder) 2020 - Present
Magneton Labs *Chicago, IL*

- Founded and operate Magneton Labs as a solo AI engineering contractor, building autonomous, production-grade AI agent systems for client companies across multiple engagements.
- Architected AgentHero, an AI-agent marketplace for a client: built a Next.js 16/Supabase platform where Claude- and Codex-powered agents are deployed, supervised, and credit-metered, with durable execution via Inngest and automated job decomposition.
- Engineered a provider-agnostic LLM abstraction layer spanning Anthropic Claude, OpenAI Codex, Google, and local models with bring-your-own-key support, enabling per-agent model selection and graceful failover.
- Built ContextFS, a Model Context Protocol (MCP) server giving Claude Code and Codex cross-session recall, scaling past 2,300 stored memories with hybrid keyword and semantic search.
- Shipped an autonomous agent-testing framework: TypeScript MCP server (dual stdio/HTTP transport, 12 tools, 6 resources) driving browser-based agents that authenticate and exercise production workflows without human intervention.
- Designed Agent-OS, an Erlang/OTP AI operating system modeling agents as supervised OS processes with typed memory, capability-based tool access, and credit-weighted scheduling.
- Standardized a Claude and Codex multi-agent development workflow with parallel agent waves, Gemini peer review, and Codex code review, adopted across 60+ repositories.

Principal AI Engineer 2023 - Present
YonedaAI

- Built GrokRxiv, a multi-agent research pipeline orchestrating parallel Claude research agents with Gemini peer review and Codex verification to autonomously draft, peer-review, and publish formally-verified papers; live at grokrxiv.org with published reviews at github.com/GrokRxiv.
- Formalized mathematics with Codex-driven proof agents using Lean 4 (extra-high reasoning, multi-pass) on a Homotopy Type Theory programme toward the Riemann Hypothesis, producing interdependent papers and machine-checked Lean modules.
- Designed and implemented JAPL, a strict typed functional language combining Rust-style ownership, Erlang processes, and functional semantics into a compiler with 170+ tests and a 46-page documentation site.
- Delivered ferrous-bridge, an AI-driven C-to-Rust translation toolchain specified across a six-rung safety ladder, proven on libyaml and executed by autonomous agent teams across four peer-reviewed papers.
- Automated the full research-to-publish pipeline with agent workflows compiling LaTeX to PDF, generating HTML/Open Graph assets, deploying to Vercel, and notifying via Slack and email.

Lead Software Engineer & DevOps Architect July 2026 - October 26, 2026
Belden

- Migrated SQLite storage to PostgreSQL while preserving project files, version history, accounts and permissions; developed transactional import tooling and data-integrity validation.
- Implemented permission-aware PostgreSQL search and tested collaboration safeguards for idempotent retries, stale-write rejection, authorization and transactional rollback.
- Engineered backend modernization for a collaborative network-design platform using Python/aiohttp and PostgreSQL, preserving existing React/TypeScript interfaces and application behavior.
- Built DevOps tooling for reproducible Docker/Compose environments, automated Alembic migrations, application initialization, controlled updates and recovery, including Windows PowerShell setup.
- Designed an enterprise Azure architecture and authored modular Bicep templates for private Container Apps, PostgreSQL Flexible Server, Blob Storage, Key Vault and Container Registry.
- Defined enterprise security boundaries using private networking, private endpoints/DNS, least-privilege managed identities and RBAC, with isolated development, staging and production environments.
- Accelerated engineering delivery with Claude- and Codex-assisted development, parallel implementation, independent code reviews and automated regression checks, retaining human oversight and acceptance gates.
- Produced architecture diagrams and developer/IT handoff documentation covering release governance, monitoring, backup/restore and staged Azure delivery.

Lead Software Engineer & DevOps Architect 2020 - May 2026

TSI

- Developed the Rescue Coordination System, a browser-based search-and-rescue platform built with React/TypeScript, PHP/Laravel, PostgreSQL and Redis.
- Built incident-management and distress-beacon workflows, including Cospas-Sarsat alerts, registry search, operational logging and role-based access.
- Implemented geospatial search planning with editable search/tasking areas, route and waypoint tools, coordinate handling, and incident/sighting map overlays.
- Integrated real-time updates with Laravel Echo/Pusher and aircraft-tracking data through the ADS-B module to support shared operational awareness.
- Expanded automated delivery and regression coverage with GitHub Actions, Cypress, Jest and PHPUnit, alongside Docker/Helm workflows for cloud and on-premises environments.
- Achieved 99.95% uptime by migrating critical applications to AWS with auto-scaling Kubernetes clusters, eliminating downtime during peak traffic periods.
- Improved application performance 40% by optimizing PostgreSQL queries and implementing Redis caching, reducing average response time from 2.5s to 1.5s.
- Reduced infrastructure costs 25% (FinOps) through resource rightsizing, reserved-instance planning, and IaC best practices; deployed Terraform configurations for AWS EKS, RDS, ElastiCache, and VPC networking.

Lead Software Engineer & DevOps Architect 2024 - 2025

Zoro/Grainger | Data Platforms & DevOps

- Built and maintained Palantir Foundry data pipelines, transformations and operational datasets supporting order, shipment and returns workflows.
- Developed Foundry data quality checks and monitoring views to improve fulfillment-data traceability and production troubleshooting.
- Reduced shipment query latency 60% by architecting a real-time tracking service processing 2M+ daily package updates with sub-second response times.
- Increased returns processing efficiency 40% by building an automated return service handling complex partial shipment logic, reducing manual intervention from 30% to 5%.
- Eliminated 95% of webhook failures by implementing FedEx integration with Datadog monitoring and automated retry logic.
- Scaled order APIs to 10K requests/second with high-performance order management microservices supporting peak holiday traffic.
- Automated infrastructure provisioning with Terraform modules for GCP resources (Cloud Run, Pub/Sub, BigQuery) and streamlined CI/CD with GitHub Actions blue-green deployments, reducing release cycles from weekly to daily.

Senior Full-Stack/DevOps Engineer 2023 - 2024

Cogability | Data Engineering & AI *Florida*

- Processed 1.2M+ documents via a scalable Lambda-based pipeline with automated quality validation, and improved data extraction reliability 92% with Python ETL and multi-stage validation, reducing manual corrections from 25% to 2%.
- Reduced document processing time 85% by building an AI-powered automation platform that decreased ticket processing from 40 hours to 6 hours per month for municipal governments handling 100K+ parking tickets monthly.
- Achieved 99.5% classification accuracy by integrating AWS Textract OCR with Vision LLMs (Claude, GPT-4V) for intelligent document classification, eliminating manual sorting.
- Built React admin dashboards with live processing metrics, role-based access controls, and audit logging for SOC2 readiness.

Senior Software/DevOps Engineer 2022 - 2023

Shell | Data Engineering & MLOps *London*

- Scaled to 10M+ daily events by architecting an event-driven microservices platform on AWS EKS/GCP GKE with Kafka streaming and under 5s latency.
- Accelerated model iteration 50% by implementing MLflow and Airflow orchestration, reducing model retraining cycle from 4 days to 2 days, and improved deployment velocity 3x with GitOps via ArgoCD across 15+ microservices.
- Reduced forecasting errors 30% by developing ensemble ML models using TensorFlow and scikit-learn, improving prediction accuracy from 60% to 90% and saving $600K+ annually.

Senior Full-Stack Engineer 2021 - 2023

PredictiveHR | Data & ML Platforms *Braintree, MA*

- Reduced forecasting time 95% by architecting automated data pipelines with AWS Glue and Redshift, decreasing workforce planning cycle from 30 days to 1.5 days.
- Improved attrition prediction 250% by building an ML-powered analytics platform increasing accuracy from 40% to 85%, enabling proactive retention strategies that reduced turnover 15%.
- Delivered enterprise SSO and SCIM provisioning with SAML/OIDC and SCIM user/group provisioning for multi-tenant B2B customers, enabling automated onboarding/offboarding and centralized identity management.
- Created interactive React dashboards visualizing predictive models, headcount trends, and cost forecasts used by 200+ HR professionals.

Senior Software Engineer 2020 - 2023
Fidelity Labs *Boston, MA*

- Achieved 99.995% uptime by architecting multi-cloud infrastructure (AWS/GCP) with auto-failover and disaster recovery, exceeding 99.99% SLA with zero critical incidents.
- Improved compliance accuracy 45% by implementing BERT-based NLP models for regulatory document classification, reducing false positives from 30% to 15%.
- Scaled to 10K+ requests/second with high-performance Java Spring microservices and ReactJS frontend handling peak regulatory transaction volumes.

Senior Full-Stack/DevOps Developer 2018 - 2022
Regis Corporation *Minneapolis, MN*

- Increased client retention 25% by building an ML-powered document management system freeing 60 hours/month per advisor for client engagement.
- Achieved 94% classification accuracy with AWS SageMaker ML models automatically categorizing financial documents, reducing manual review from 100% to 6%.
- Built a ReactJS portfolio management platform serving 500+ franchisees with real-time analytics, improving decision-making speed 40%.

Lead Full-Stack Engineer 2020 - 2021
BetaCom *Bellevue, WA*

- Reduced incident detection time 95% by building a real-time monitoring dashboard with sub-second WebSocket updates, decreasing alert latency from 8 minutes to 20 seconds and reducing SLA penalties 75%.
- Scaled to 50M+ daily events via a Kafka streaming platform with Python FastAPI async backend processing high-volume 5G network telemetry with 99.95% delivery guarantee.
- Optimized query performance 80% by designing a PostgreSQL time-series schema with partitioning and indexing, reducing dashboard load times from 15s to 3s.

Senior Software Engineer 2016 - 2018
Evident.io *Pleasanton, CA*

- Core engineer contributing critical features that positioned Evident.io for its $300M acquisition by Palo Alto Networks.
- Reduced AWS costs $70K annually (FinOps) by leading cloud cost analysis through rightsizing, reserved instances, and architectural improvements, decreasing monthly spend 35%.
- Scaled to 1M+ events/day by engineering a cloud-native microservices architecture processing security events with 99.9% detection accuracy and under 30s latency.
- Awarded Bug Fix of the Year 2017 for resolving a critical performance bottleneck that improved detection speed 5x and prevented customer-facing incidents.

SKILLS

Data Engineering: Databricks, Apache Spark, Python, SQL, ETL/ELT, data pipelines, data modeling, data quality and validation
Data Platforms & Streaming: Palantir Foundry, Kafka, Airflow, AWS Glue, Redshift, BigQuery, Snowflake, PostgreSQL, Redis, DynamoDB
AI/ML & MLOps: MLflow, TensorFlow, PyTorch, scikit-learn, SageMaker, BERT/Transformers, model training and deployment
AI Agents & LLMs: LLM orchestration (Claude, OpenAI Codex, Gemini), Model Context Protocol (MCP), multi-agent systems
DevOps & Delivery: GitHub Actions, GitLab CI, Jenkins, CircleCI, AWS CodePipeline, GitOps, ArgoCD, CI/CD, automated testing
Cloud & Infrastructure: AWS (EKS, Lambda, CDK, S3, RDS), Azure (Container Apps, PostgreSQL Flexible Server, Blob Storage, Key Vault, managed identities), GCP, Terraform, Bicep, Kubernetes, Docker, Helm
Languages: Python, SQL, PowerShell, PHP, TypeScript/JavaScript, Go, Rust, Haskell, Java, Scala, Ruby
Frameworks: Next.js, React, Vue 3, Node.js, Laravel, aiohttp, FastAPI, Django/DRF, Spring Boot
Formal Methods & Verification: Lean 4, theorem proving, formal verification, type systems
Enterprise & Security: SSO (SAML/OIDC), SCIM, multi-tenancy, RBAC, IAM, SOC2, HIPAA, PCI, NIST

CERTIFICATIONS

- AWS Certified Solutions Architect
- Microsoft Certified: Azure Solutions Architect Expert
- Additional credentials: Red Hat Certified; Microsoft Certified Professional

AWARDS & HONORS

- Bug Fix of the Year 2017, Evident.io
- DARPA Distinguished Technology Innovator and Leadership Alliance Member
- Patent Holder: Secure Distributed Messaging Architecture

EDUCATION

University of Colorado | M.S. and B.S., Computer Science.
Aion University | PhD-Level Research, Foundations of Mathematics and Physics.