Jake Sterns

Software Engineer | AI Automation & Data Engineering | Platform Integrations

(210) 999-0555 | San Antonio, TX | jakesterns21@gmail.com | GitHub | LinkedIn

Software and data engineer with 5+ years of production experience building AI-powered agentic workflows, backend services, and data pipelines across enterprise platforms. Hands-on with LLM APIs (Claude, OpenAI), MCP (Model Context Protocol), LangChain-style orchestration frameworks, RAG/LightRAG pipelines, and Python-first backend systems. Proven track record shipping AI tools that real users depend on — including automated query optimization workflows generating multi-million-dollar monthly savings opportunities. Experienced integrating systems across Salesforce-style CRMs, Snowflake/BigQuery-class data warehouses, Slack, REST APIs, and observability platforms (Grafana, OpenTelemetry, Datadog). Comfortable with autonomy, ambiguity, and cross-functional collaboration with non-technical stakeholders.

AI & PLATFORM PROJECTS

AI-Assisted Query Optimization Workflow (Production | Python, MCP, LLM APIs)

- Designed a near-real-time agentic workflow combining query telemetry, platform resource metrics, and multi-database context across UDB, Oracle, and SQL Server — aligned with LLM-driven recommendation patterns used in RAG pipelines.
- Built a tool-integrated backend using MCP-style server patterns to expose structured performance context to LLM-driven automation workflows, enabling context-aware optimization guidance with a human-in-the-loop review model.
- System targets multi-million-dollar monthly savings by surfacing inefficient query patterns at scale — demonstrating end-to-end ownership of an AI tool in a production environment.

Query Performance ML Classification Engine (Production | Python, XGBoost, scikit-learn)

- Developed an ML pipeline to classify millions of rows of query performance data and identify inefficient query patterns across enterprise database platforms using XGBoost, DBSCAN, HDBSCAN, and scikit-learn inference workflows.
- Delivered measurable cost reduction of $10K+/month by prioritizing optimization opportunities — shipped as a production tool with real operational impact.

AgentLink (Public Release | Python, WebSocket, Agentic Orchestration)

- Building a secure, local-first AI agent platform for orchestrating distributed nodes, agents, tools, and services via typed protocols, policy controls, audit events, WebSocket communication, and a full-stack control plane.
- Implements governance controls, audit trails, and human-in-the-loop escalation patterns — directly aligned with responsible AI deployment standards for sensitive business data.

OpenClaw / Mission Control (Private | Multi-Agent, Observability, Tool Routing)

- Extended a self-hosted multi-agent AI environment with custom agent governance, service integrations, tool routing, gateway communication, runtime debugging, and observability-focused operational workflows.

Hermes Agent (Private | Python, Agentic Automation, Tool Orchestration)

- Developing an internal AI agent for structured automation, tool orchestration, context-aware task execution, and reliable backend workflow support across integrated services — reusable agentic skill pattern aligned with GTM automation use cases.

Centralized Tool Registry Hub (In-Flight | LightRAG, RAG-Anything, Python)

- Designed an internal AI-powered knowledge platform using retrieval-augmented generation (RAG) via LightRAG to centralize tool governance, review status, use cases, cost visibility, and adoption guidance across enterprise teams.

CI/CD and Deployment Automation Platform (Production | GitLab CI/CD, Docker, Python)

- Modernized enterprise SDLC workflows by integrating automated testing, deployment pipelines, and ticket-generation processes — improving deployment consistency, release throughput, and operational traceability.

Observability & AI/ML Platform Monitoring (Production | Grafana, OpenTelemetry, Datadog, ELK)

- Built and enhanced monitoring and logging workflows using OpenTelemetry, Datadog, Grafana, and ELK/Logstash to improve reliability and visibility for automation and analytics services in production.

n8n Open Source Contribution (Public | Workflow Automation, Integration Orchestration)

- Contributed open-source changes to n8n with hands-on exposure to workflow automation, integration orchestration, and low-code/agent-adjacent tooling patterns — directly relevant to agentic workflow automation stacks.

WORK EXPERIENCE

Software Engineer — Data Insights & Analytics | USAA (United Services Automobile Association) May 2020 – Present

- Built production ingestion and automation applications across 10+ enterprise data platforms, enabling scalable ETL workflows and operational reporting across SQL Server, Oracle, Snowflake, UDB, Cassandra, Couchbase, Aurora/Postgres, Neptune, and Netezza.
- Designed and shipped an ML-based query performance classification workflow (XGBoost, DBSCAN, HDBSCAN, scikit-learn) identifying high-cost query patterns across multi-platform environments — generating $10K+/month in sustained savings as a production tool.
- Expanding query optimization capabilities via a near-real-time AI-assisted agentic recommendation workflow combining query-level telemetry, platform resource metrics, and MCP-style tool integrations — generating context-aware SQL optimization guidance at scale.
- Modernized CI/CD and deployment workflows by integrating Talon Batch automation into the SDLC, reducing manual release effort and improving testing, deployment, and ticketing consistency across data and analytics applications.
- Collaborated with cross-functional teams as an active Product Owner to scope, prioritize, and deliver software solutions aligned with business objectives — including communicating technical trade-offs to non-technical stakeholders.
- Built and maintained observability and monitoring infrastructure using Grafana, OpenTelemetry, Datadog, and ELK/Logstash for production automation and analytics services.

TECHNICAL SKILLS

Languages: Python (primary), Go, C#, Bash, SQL, YAML, JSON

AI / ML & Agentic Frameworks: LLM APIs (Anthropic Claude, OpenAI), MCP (Model Context Protocol), LangChain-style orchestration, RAG / LightRAG / RAG-Anything, XGBoost, DBSCAN, HDBSCAN, scikit-learn, ML Inference, AI-assisted Automation, Claude Code, Copilot, NLP (Natural Language Processors)

Backend / APIs: Python (FastAPI), REST APIs, Microservices, Systems Integration, Serverless / Containerized Services

Data / Warehouses: Snowflake, BigQuery-class SQL, Oracle, SQL Server, UDB/DB2, Cassandra, Couchbase, Neptune (graph), Aurora/Postgres, Netezza, IMS, DocDB

DevOps / Cloud Platform: Docker, Kubernetes, Helm, OpenShift, GitLab CI/CD, GitHub, GCP-compatible Cloud Run / Functions patterns, JFrog, Linux, UrbanCode, Puppet, Control-M

Observability: Grafana, OpenTelemetry, Datadog, ELK/Logstash, LangSmith-style logging, ServiceNow

Workflow Automation: n8n (contributor), dbt, Talon Batch, Agile/Scrum

EDUCATION

B.S. Computer Science (Cybersecurity) | Texas A&M University – Corpus Christi | GPA: 3.85