

Richard Pallangyo

Seattle, WA · rapaugustino@gmail.com · (206) 665-6484
linkedin.com/in/rapaugustino · github.com/rapaugustino · richardpallangyo.com

APPLIED AI & DATA ENGINEER

Enterprise AI systems · Data platforms · Customer-facing product delivery

Applied AI and data engineer with nearly three years of enterprise Azure experience and recent end-to-end ownership of deployed AI-enabled products. Led technical delivery across large engineering teams, translated operational workflows into software, and built governed AI systems with retrieval, tool use, structured outputs, and human approval. Brings hands-on depth across data platforms, backend systems, and customer-facing implementation.

Technical Skills

AI & Backend: Python, FastAPI, LangGraph, RAG, tool calling, structured outputs, Anthropic Claude, OpenAI, PostgreSQL, SQLAlchemy, REST APIs

Data & Cloud: SQL, PySpark, ETL/ELT, data modeling, data warehousing, Azure Data Factory, Databricks, Synapse, Stream Analytics, ADLS, AWS Glue, Athena

Product & Delivery: Next.js, React, TypeScript, technical discovery, solution design, enterprise integration, stakeholder communication, technical leadership

AI-Assisted Development: Claude Code, Codex

Experience

Independent AI & Software Engineer

Selected independent product work

Dec 2024 – Present

Seattle, WA

Shule Fanisi

shulefanisi.com

Apr 2026 – Present

- Building an AI-native education platform for Tanzania designed to unify school operations while giving teachers and students intelligent tools for teaching, learning, and school administration.

Safari King Africa — Independent Product Implementation

safarikingafrica.com

Dec 2024 – May 2026

- Partnered with the owner to evolve a static website into an integrated website, CRM, and CMS supporting the customer journey from inquiry through itinerary customization and booking management; owned product discovery, architecture, database design, frontend, backend APIs, authentication, deployment, and maintenance as the sole engineer.
- Replaced hard-coded website and offering updates that typically took **1–3+ days** with self-service controls that let the owner manage packages, pricing, destinations, itineraries, and content independently **within an hour**.
- Built **seven human-in-the-loop AI workflows** that retrieve bookings, prior communications, customer preferences, and business records through database tools and use web search when needed; reduced customized itinerary research and drafting from **3–5 hours to 30–60 minutes**, with owner approval required before outputs are sent or published.

Ubunifu Madness — College Basketball Analytics Platform

madness.ubunifu.tech.com

Feb 2026 – Present

- Conceived, built, deployed, and maintain a public analytics platform covering **700+ men's and women's Division I teams** and **165,000+ historical games**, with daily ETL, ratings, rankings, matchup predictions, bracket tools, and a database-grounded conversational assistant.
- Audited a 43-feature logistic-regression/LightGBM ensemble reporting nearly 80% historical accuracy, uncovered target leakage through controlled walk-forward evaluation, and replaced it with separately tuned Elo models that produced better Brier scores on held-out seasons.
- Built a PostgreSQL-backed Claude assistant with **seven database tools** and streaming responses for team, conference, roster, score, ranking, matchup, and bracket analysis, with code-enforced restrictions against unsupported claims.

Data Science Intern — Procurement Department
SSA Marine

Jun 2025 – Sep 2025
Seattle, WA

- Partnered with the Procurement Business Analysis Manager and IT leadership to address fragmented SharePoint policies, repetitive employee questions, manual document review, and procurement-related email back-and-forth by building a self-service AI assistant deployed in the company's Azure environment for stakeholder testing.
- Built a multi-turn LangGraph/FastAPI experience where employees could ask procurement questions, receive answers grounded in source documents with citations, retain conversational context, and generate procurement-related email drafts requiring explicit approval before dispatch.
- Designed a natural-language-configurable invoice extraction workflow using the OpenAI Batch API; users could select a prebuilt schema or create and edit one on demand, validate structured results, process tested batches of **100 invoices**, and export completed datasets to CSV or Excel.
- Integrated Azure OpenAI, Azure AI Search, App Service, APIM, Entra ID SSO, and Logic Apps into a single enterprise surface; enforced reliability through semantic and hybrid retrieval, Pydantic-gated tool calls, redacted audit logging, health checks, and LangGraph execution traces.
- Demonstrated the platform to approximately **8–15 stakeholders** across procurement, IT, data governance, and strategy, then completed a technical handoff covering architecture, documentation, training, and known limitations.

Senior Data Engineer (Team Lead)
LTIMindtree Limited

Jun 2022 – Aug 2024
Seattle, WA

- Selected by senior leadership to help establish and scale LTIMindtree's first Canada-based Data & AI delivery team while working remotely from Seattle; helped grow the organization from inception to **30+ engineers** through interviewing, onboarding, and readiness reviews.
- Led technical delivery across the 30+ engineer organization by assigning client engagements, setting delivery priorities, reviewing architectures and code, running delivery reviews, mentoring engineers, and coordinating customer escalations.
- Redesigned incident response across enterprise Azure data workloads through severity-based triage, expertise-based routing, reusable runbooks, recurring reviews, and stronger escalation practices; reduced average resolution time by **30%** while maintaining **4.8+/5 customer satisfaction** across ticket surveys and client scorecards.
- Recovered an Azure Stream Analytics pipeline that had fallen **24+ hours behind** by diagnosing sequential queries, insufficient partitioning, and compute contention; parallelized processing, increased partitions, and right-sized compute to restore data availability **within seconds**, then verified stability over several days.
- Served as a technical bridge between enterprise clients and Microsoft support and product engineering, translating production issues into reproducible bug reports and feature requests and coordinating fixes, workarounds, and implementation guidance across Azure data and AI services.
- Directly mentored four engineers and recommended them for advancement; all four later moved into leadership roles.

Junior Data Engineer
LTIMindtree Limited

Oct 2021 – Jun 2022
Seattle, WA

- Built and optimized enterprise workloads across Azure Data Factory, Databricks, Synapse, ADLS, SQL Server, SSIS, and Azure Stream Analytics, including SQL transformations, data-quality checks, cloud migrations, and production support.
- Investigated production incidents directly with clients, documented root causes and preventive actions, and worked with Microsoft support and engineering when platform-level escalation was required; promoted to Senior Data Engineer after **approximately eight months** based on performance and expanded responsibility.

Education

University of Washington
MS, Data Science

Sep 2024 – Mar 2026
GPA: 3.91

Seattle University
BA, Computer Science; Minors in Mathematics & Data Science
Cum Laude · Dean's List, 4×

Jun 2019 – Jun 2021

Seattle Central College
AS, Computer Science

Jun 2017 – Mar 2019