

German Neironi

Cloud and AI Solutions Engineer | Founder at CloudPruneAI and Quira

g.neironi@gmail.com | linkedin.com/in/gneironi | gneironiar.github.io/personal_website | github.com/gneironiar

Cordoba, Argentina (Open to Remote - US/EU/LATAM)

Professional Summary

Cloud and AI Solutions Engineer with 15+ years of experience designing and building cloud solutions that integrate AI to solve real business problems - from architecture to production deployment.

Founded and operates two SaaS products in production: CloudPruneAI (FinOps platform with AI-powered IaC code generation, 18 analyzers, 31+ checks, B2B2B partner channel across 3 countries) and Quira (multi-tenant industrial IoT platform for environmental monitoring, launched April 2026). Both products built solo from architecture to production, integrating Claude API for code generation, business context enrichment, and customer-facing features.

Expert in full-stack cloud development with AWS (CDK, Lambda, ECS Fargate, Bedrock, RDS, S3, Glue, Athena), FastAPI, Next.js, PostgreSQL, and Generative AI integration (Claude API, AWS Bedrock, Prompt Engineering, RAG). Skilled in Python, TypeScript, Docker, and designing hybrid serverless architectures for production systems.

Currently leading cloud data and AI projects at Argeniss Software while operating CloudPruneAI and Quira.

Professional Experience

CloudPruneAI | Founder and Technical Lead

November 2025 - Present | Independent SaaS Startup | Remote

- Founded production FinOps SaaS platform analyzing AWS accounts and auto-generating Infrastructure-as-Code (CDK Python, Terraform HCL, CloudFormation YAML) for cost optimization using Generative AI and LLMs
- Built and deployed 18 resource analyzers covering 31+ optimization checks across compute, storage, networking, databases, commitments, extended support, and AI/ML services
- Architected hybrid serverless platform: FastAPI on AWS Lambda + Mangum for API, ECS Fargate for scan workers, async RDS PostgreSQL, Auth0, Stripe, Resend, monorepo with Turbo
- Implemented AI-powered IaC code generation using Claude API (Anthropic) with multi-format output and business context enrichment via Claude Haiku
- Designed B2B2B partner-led GTM model: white-label PDF reports, gainshare pricing (10% / 7.5% at scale) plus recurring Operate Plan subscriptions
- Onboarded 4 partners across Argentina, Spain, and Uruguay; first real scan with consulting partner detected USD 5,260/year in waste across a client AWS account
- Drove technical onboarding of partners across 3 countries: discovery calls, AWS account integration, white-label configuration, and hands-on support during first scans
- GDPR-compliant architecture with Privacy Policy, DPA, ROPA, Security Posture Document, and DELETE /users/me

Tech Stack: AWS CDK, Lambda, ECS Fargate, RDS PostgreSQL (async), API Gateway, S3, SQS, Cost Explorer, FastAPI, Next.js 14, Material UI, Auth0, Stripe, Claude API, Python, TypeScript, Docker, Turbo

Quira | Founder and Technical Lead

December 2025 - Present | Independent SaaS Startup | Remote

- Founded multi-tenant industrial IoT platform for environmental monitoring (temperature, humidity), with real-time alerts, auditable PDF/CSV reports, and native Microsoft Teams integration via Power Automate
- Designed and deployed full production stack on AWS in under 4 months: 9 CDK stacks covering ECS Fargate, ALB, WAF, Route53, ACM, RDS PostgreSQL 15, Cognito, SQS, 6 Lambdas, and SNS
- Built FastAPI backend (47 endpoints, 103 tests) with HMAC-SHA256 + HKDF device authentication, async ingestion pipeline with idempotent client-side deduplication, alert engine with severity escalation, and async report generation via WeasyPrint workers
- Designed sensor-agnostic data model and driver registry to support multi-sensor expansion without schema changes (CO2, pressure, vibration on the V2 roadmap)

- Developed Next.js 14 admin panel with 25 pages, 6-section sidebar, i18n (ES/EN), and Recharts-based reading visualizations
- Validated end-to-end with real ESP32 + SHT31 hardware: device readings reach the platform, trigger alerts, and render in production charts
- Designed B2B2B revenue model: tripartite split between platform, partner, and vendor; first partner LOI signed in Colombia (aerospace warehouse pilot, ~30 devices)
- Production launch: April 19, 2026

Tech Stack: AWS CDK, ECS Fargate, ALB, WAF, Cognito, RDS PostgreSQL 15, SQS, Lambda, SNS, Route53, ACM, KMS, FastAPI, Next.js 14, Material UI, next-intl, ESP32 (Arduino C++), SHT31, HMAC-SHA256, Python, TypeScript, PostgreSQL, Resend

Argeniss Software | Solutions Architect / Pre-Sales Engineer

March 2026 - Present | Multiple Clients | Remote

- Technical advisor in pre-sales engagements supporting the COO during client discovery and proposal phases
- Lead technical assessments, architecture proposals, and stack recommendations for early-stage and pre-seed clients
- Created Technical Assessment Playbook (Discovery -> Assessment -> Proposal) standardizing the pre-sales workflow
- Generate additional engineering scope for delivery teams (PM + Architect + Engineer) by anchoring trust through technical depth

Tech Stack: AWS, Cloud Architecture, Solution Design, Technical Discovery, Stack Selection

Argeniss Software | AWS Data Integration Engineer

October 2025 - Present | Fintech Client | Remote

- Architected automated DMS system replicating 10-30 SQL Server databases to Redshift with CDC (30-120 sec latency)
- Built reusable Terraform modules (DMS instance, endpoints, KMS) following Infrastructure as Code best practices
- Designed multi-source architecture supporting SQL Server (DMS), MongoDB (Change Streams), REST APIs (Salesforce, NetSuite)

- Developed orchestration system with Step Functions and Lambda for zero-touch database onboarding (3 commands)
- Implemented DynamoDB-driven configuration catalog and S3 data lake for semi-structured data
- Deployed multi-AZ, Security Hub compliant infrastructure with automated KMS rotation
- Created Redshift Spectrum integration for querying S3 data alongside relational tables
- Delivered 60-70% infrastructure reusability, reducing onboarding time from weeks to days

Tech Stack: AWS DMS, Lambda, Step Functions, DynamoDB, S3, Redshift, KMS, Terraform, Python, SQL Server, MongoDB, REST APIs

Argeniss Software | AWS Cloud and AI Engineer

February 2025 - October 2025 | Multiple Clients (E-learning, Healthcare) | Remote

Concurrent client engagements covering data engineering, AI/LLM integration, and machine learning forecasting.

- Architected complete ETL system extracting student performance data from MySQL to S3 Data Lake with Parquet optimization
- Built dynamic REST API with API Gateway, Lambda, and Athena for configurable analytics endpoints
- Developed LLM-powered analysis system using AWS Bedrock (Claude 3.5 Sonnet) and Prompt Engineering for personalized educational recommendations (RAG pattern)
- Implemented 6 production Lambda functions (containerized) handling ETL, API queries, LLM analysis, and NCLEX question generation
- Created real-time update pipeline with SQS for user-specific data refreshes
- Designed Vue.js frontend with CloudFront for NCLEX question generator with iterative AI feedback
- Deployed multi-stack CDK infrastructure with automated Glue Crawlers and EventBridge schedulers
- Optimized query costs with Athena partitioning strategy, achieving sub-second API response times
- Built API data extraction and transformation pipelines with automated JSON-to-Parquet conversion and partitioned storage
- Implemented secure credential management with Parameter Store and S3 encryption
- Implemented machine learning forecast infrastructure using SageMaker Canvas with dual prediction system for short-term (180 days) and long-term (12 months) forecasting

- Designed multi-model Generative AI pipeline using AWS Bedrock with 4 concurrent models (Claude Sonnet, Haiku, Nova Lite, Nova Pro), orchestrated by Step Functions in a 3-stage LLM workflow
- Implemented automated CSV processing handling 300+ records with EventBridge triggers
- Set up GitHub Actions CI/CD pipeline with multi-environment deployment

Tech Stack: AWS CDK, Lambda (Docker), API Gateway, AWS Bedrock, Claude API, Athena, Glue, S3, DynamoDB, SQS, CloudFront, EventBridge, Secrets Manager, Parameter Store, Step Functions, SageMaker Canvas, GitHub Actions, Python, MySQL, Parquet, Vue.js, Docker

Scale AI | AI Quality Assurance Engineer

August 2023 - September 2024 | Remote

- Trained generative AI models by performing code review of LLM output
- Created code snippets with reasoning to improve AI model responses
- Identified and fixed bugs to enhance model performance
- Supported prompt engineering and quality standards for Generative AI systems

Tech Stack: Python, NumPy, Pandas, Scikit-learn, Keras, PyTorch, Docker, AWS

NTB Solutions | Senior Data Analyst

August 2018 - February 2024 | Remote

- Worked on data analysis and integration projects for healthcare clients across multiple countries
- Built ETL processes using Azure Data Factory, Notebooks, Scala, and Python
- Created BI dashboards using Power BI, Pentaho, and Tableau
- Designed and implemented health information systems (HIS, LIS, RIS)
- Set up HL7-FHIR health interoperability processes
- Achievement: Built custom report recovering billions of pesos annually in previously uncontrolled funds

Tech Stack: Azure Data Factory, Python, Spark, PySpark, SQL, Pentaho, Power BI, MIRTH Connect, MySQL, Oracle, PostgreSQL, MongoDB

Anyone AI | Machine Learning Engineer (Part Time)

July 2022 - December 2022 | Remote

- Developed sentiment analysis, image classification, credit risk analysis, and salary prediction models
- Built and trained ML models using Python and various libraries
- Integrated models into API-based web services with Docker deployment

Tech Stack: Python, Scikit-learn, Keras, PyTorch, TensorFlow, Docker

Sanatorio la Entrerriana | Data Analyst and IT Manager

July 2008 - August 2018 | Parana, Argentina

- Led multi-year implementation of integrated HIS/LIS/RIS health information system across the institution, achieving full staff adoption and replacing legacy paper-based workflows
- Created executive dashboards and operational reports using Power BI and Google Data Studio
- Managed Oracle databases and led data-driven decision-making processes for clinical and administrative operations

Tech Stack: Oracle, SQL, Power BI, Google Data Studio, ETL processes

Core Technical Skills

Cloud and Infrastructure:

AWS (CDK, Lambda, Step Functions, ECS Fargate, ALB, WAF, Cognito, ACM, Route53, SNS, DMS, Bedrock, S3, Glue, Athena, Redshift, DynamoDB, SQS, KMS, EventBridge, Secrets Manager, CloudFront, SageMaker), Terraform, Azure

Generative AI and LLMs:

AWS Bedrock, Claude API (Anthropic), Prompt Engineering, RAG (Retrieval-Augmented Generation), LLM Integration, AI Agents, Generative AI, SageMaker Canvas

Programming and Frameworks:

Python, TypeScript, SQL (T-SQL, PL/SQL), JavaScript, FastAPI, Next.js 14, Vue.js, R, Arduino C++ (ESP32), Embedded Systems

Data Engineering:

ETL/ELT, Parquet, JSON, Pandas, PySpark, Apache Spark, Data Pipeline Design

Databases:

PostgreSQL, MySQL, SQL Server, MongoDB, Oracle, Amazon RDS, Amazon Redshift, PostGIS

DevOps and Tools:

Docker, Git, GitHub Actions, CI/CD, Infrastructure as Code (CDK, Terraform), Jira

Business Intelligence:

Power BI, Tableau, Pentaho, QuickSight, SSRS

Key Projects and Achievements

CloudPruneAI Platform:

Production B2B2B FinOps SaaS with 18 analyzers and 31+ optimization checks. AI-powered IaC code generation in 3 formats (CDK, Terraform, CloudFormation). 4 active partners across Argentina, Spain, and Uruguay. First real scan detected USD 5,260/year in client savings.

Quira IoT Platform:

Multi-tenant industrial IoT platform for environmental monitoring, launched April 2026. 9 CDK stacks on AWS, 47 backend endpoints, end-to-end validated with real ESP32 hardware. First partner LOI signed in Colombia (aerospace pilot).

Multi-Source DMS Architecture:

Automated replication system supporting SQL Server, MongoDB, and REST APIs with 60-70% infrastructure reusability.

Educational Analytics Platform with LLM Integration:

Complete ETL system with RAG-based LLM analysis, dynamic REST API, and NCLEX question generator using AWS Bedrock.

Healthcare Data Recovery:

Custom reporting solution recovering billions of pesos annually in previously uncontrolled healthcare funds.

Education and Certifications

Bachelor's Degree in Business Management

Universidad Autonoma de Entre Rios (UADER), Parana, Argentina

Certifications:

- AWS Certified Cloud Practitioner
- AWS Cloud Services (Kinesis, EMR, Glue, Athena, QuickSight, Redshift)
- Azure Fundamentals and Data Engineering (Microsoft)
- Machine Learning Engineer - Anyone AI
- C1 English Language - EFSET
- Scrum Foundation Professional
- ITIL Foundation Certificate in IT Service Management

In Progress:

- AWS Certified Solutions Architect - Associate (SAA-C03) - exam scheduled April 2026

Soft Skills

- Strategic Thinking and Problem Solving
- Team Leadership and Collaboration
- Continuous Learning and Growth Mindset
- Strong Written and Verbal Communication in English
- Detail-Oriented and Goal-Driven
- Remote Work Excellence
- Cross-Functional Team Coordination
- Stakeholder Management