

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 over 15 years of progressive experience designing and building cloud solutions that integrate AI to solve real business problems - from architecture to production deployment. Recognized expert in transforming business challenges into scalable, production-ready systems across healthcare informatics, enterprise data systems, and SaaS platform development.

Core Expertise

- **AI-Powered Cloud Solutions:** Production experience building systems that integrate LLMs and Generative AI services (Claude API, AWS Bedrock) to automate complex workflows - from infrastructure code generation to intelligent data analysis
- **Full-Stack Cloud Architecture:** Deep expertise in AWS ecosystem (CDK, Lambda, ECS Fargate, RDS, DMS, Bedrock, S3, Glue, Athena, Redshift, DynamoDB, Step Functions, SQS, EventBridge, CloudFront) and Terraform Infrastructure-as-Code
- **Data Engineering and Analytics:** Design and implementation of hybrid serverless architectures, real-time data pipelines, multi-source ETL systems, and optimized data lake architectures that turn raw data into business decisions
- **Modern Development Stack:** FastAPI, Next.js 14, Vue.js, TypeScript, Python, PostgreSQL, with full CI/CD automation via GitHub Actions
- **Leadership and Strategy:** Proven track record founding two SaaS startups (CloudPruneAI and Quira), leading cross-functional teams of 26+ professionals, and managing enterprise-scale transformation projects

Professional Differentiators

Founded and operates two SaaS products in production: CloudPruneAI (FinOps platform with AI-powered IaC code generation in 3 formats - CDK, Terraform, CloudFormation - 18 analyzers, 31+ optimization checks, B2B2B partner channel with 4 active partners across Argentina, Spain, and Uruguay) and Quira (multi-tenant industrial IoT platform for environmental monitoring, launched April 2026, validated end-to-end with real ESP32 hardware, first partner LOI signed in Colombia). Both products built solo from architecture to production, integrating Claude API for code generation, business context enrichment, and customer-facing features.

Currently leading cloud data and AI integration projects at Argeniss Software while operating CloudPruneAI and Quira. Extensive experience working with distributed teams across multiple time zones and countries using modern collaboration tools (Slack, Discord, Teams, Jira, Figma).

Results-oriented professional who translates complex business problems into scalable, production-ready solutions where AI generates real output, not just prototypes. Strategic thinker with diplomatic approach to stakeholder management and commitment to continuous learning in emerging AI technologies.

Professional Experience

CloudPruneAI | Founder and Technical Lead

November 2025 - Present | Independent SaaS Startup | Remote

Founded and built production FinOps SaaS platform from concept to live partner-channel product in 4 months, automating AWS cost optimization through AI-powered infrastructure-as-code generation across CDK, Terraform, and CloudFormation.

Key Responsibilities and Achievements

Product Development and Architecture:

- Architected and deployed hybrid serverless platform combining FastAPI on AWS Lambda + Mangum (API layer) with ECS Fargate (long-running scan workers) and async RDS PostgreSQL for data persistence
- Designed and implemented 18 specialized resource analyzers covering 31+ optimization checks across compute, storage, networking, databases, commitments, extended support, and AI/ML

services

- Built AI-powered IaC code generation using Claude API (Anthropic) with multi-format output (CDK Python, Terraform HCL, CloudFormation YAML) and business context enrichment via Claude Haiku
- Developed complete REST API with authentication (Auth0), role-based access control, and multi-tenant support
- Structured the codebase as a monorepo with Turbo for unified TypeScript/Python tooling

Business Model and Monetization:

- Designed B2B2B partner-led GTM model with white-label PDF reports for consulting partners
- Implemented gainshare pricing (10% of detected savings, 7.5% at scale) plus recurring Operate Plan subscriptions for continuous monitoring
- Integrated Stripe payment processing with automated recommendation unlock tied to payment events
- Created customer dashboard for real-time savings tracking and IaC code downloads in 3 formats

Technical Infrastructure:

- Deployed full CI/CD pipeline via GitHub Actions with automated testing and multi-environment deployment
- Implemented comprehensive monitoring and alerting using CloudWatch
- Configured multi-AZ deployment for high availability
- Established data encryption at rest and in transit using KMS
- Integrated Resend for transactional email and notifications

Compliance and Security:

- Designed GDPR-compliant architecture with Privacy Policy, Data Processing Agreement (DPA), Records of Processing Activities (ROPA), and Security Posture Document
- Implemented user-controlled data deletion via DELETE /users/me endpoint
- Documented data flows and retention policies for partner due diligence

Partner Onboarding and Validation:

- Onboarded 4 active partners across Argentina, Spain, and Uruguay
- Drove technical onboarding end-to-end: discovery calls, AWS account integration, white-label configuration, and hands-on support during first scans
- First real scan with consulting partner detected USD 5,260/year in waste across a client AWS account
- Registered USPTO trademark (Serial: 99585602)

- Established company LinkedIn presence and go-to-market materials
- Deployed production environment: app.cloudpruneai.com

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

Quantifiable Impact:

- 4-month delivery from concept to live partner channel
- 18 production-grade analyzers covering 31+ optimization checks
- 3 IaC output formats (CDK, Terraform, CloudFormation) auto-generated by AI
- 4 active partners across 3 countries (Argentina, Spain, Uruguay)
- USD 5,260/year detected in first real partner scan

Quira | Founder and Technical Lead

December 2025 - Present | Independent SaaS Startup | Remote

Founded and deployed multi-tenant industrial IoT platform for environmental monitoring (temperature, humidity), built solo from architecture to production launch in under 4 months. Validated end-to-end with real ESP32 hardware and first partner LOI signed in Colombia.

Key Responsibilities and Achievements

Product Development and Architecture:

- Designed and deployed full production stack on AWS via 9 CDK stacks: ECS Fargate behind ALB with WAF, Route53 + ACM, RDS PostgreSQL 15, Cognito, SQS queues, 6 Lambdas, and SNS for alert delivery
- Built FastAPI backend with 47 endpoints and 103 tests covering authentication, ingestion, alerts, devices, organizations, sites, reports, and webhooks
- Implemented HMAC-SHA256 + HKDF device authentication with per-device key derivation, replay protection, and idempotent client-side deduplication on the ingestion pipeline
- Designed alert engine with severity escalation logic and async report generation via WeasyPrint workers (PDF/CSV)

Sensor-Agnostic Data Model:

- Designed sensor-agnostic data model and driver registry to support multi-sensor expansion without schema changes
- Built CO2, pressure, and vibration support into the V2 roadmap without requiring database migrations
- Implemented native Microsoft Teams integration via Power Automate for enterprise alerting

Frontend and User Experience:

- Developed Next.js 14 admin panel with 25 pages organized in a 6-section sidebar
- Implemented full i18n (Spanish and English) using next-intl
- Built reading visualizations with Recharts for time-series temperature and humidity data

Hardware Integration and Validation:

- Validated end-to-end with real ESP32 + SHT31 hardware: device readings reach the platform, trigger alerts, and render in production charts
- Built device firmware in Arduino C++ with HMAC signing, OTA-ready architecture, and configurable sampling intervals

Business Model and Validation:

- Designed B2B2B revenue model with tripartite split between platform, partner, and vendor
- Signed first partner LOI in Colombia (aerospace warehouse pilot, ~30 devices)
- Production launch: April 19, 2026

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

Quantifiable Impact:

- Concept to production in under 4 months, solo
- 9 CDK stacks deployed on AWS
- 47 backend endpoints, 103 tests
- 25 frontend pages, full ES/EN i18n
- End-to-end validated with real ESP32 + SHT31 hardware
- First partner LOI signed (Colombia, aerospace pilot ~30 devices)

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. Anchors trust through technical depth and converts discovery conversations into delivery scope.

Key Responsibilities and Achievements

Pre-Sales Technical Leadership:

- Lead technical assessments, architecture proposals, and stack recommendations for early-stage and pre-seed clients
- Run discovery calls alongside the COO to surface technical requirements and constraints
- Translate ambiguous business goals into concrete technical roadmaps and stack choices

Process and Standardization:

- Created Technical Assessment Playbook (Discovery -> Assessment -> Proposal) standardizing the pre-sales workflow across the company
- Documented decision frameworks for stack selection (cloud provider, language, framework, data layer)
- Established architecture proposal templates reusable across engagements

Business Impact:

- Generate additional engineering scope for delivery teams (PM + Architect + Engineer) by anchoring trust through technical depth
- Bridge the gap between commercial conversation and engineering execution
- Enable the COO to engage technical buyers credibly during early-stage sales cycles

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

Argeniss Software | AWS Data Integration Engineer

October 2025 - Present | Fintech Company | Remote

Lead engineer for enterprise-scale automated database replication system supporting 10-30 SQL Server databases with plans to expand to MongoDB and REST API sources.

Key Responsibilities and Achievements

System Architecture and Design:

- Architected complete automated DMS (Database Migration Service) system replicating SQL Server databases to Amazon Redshift with Change Data Capture (CDC) achieving 30-120 second latency
- Designed multi-source extensible architecture supporting SQL Server (via DMS), MongoDB (Change Streams), and REST APIs (Salesforce, NetSuite) through unified orchestration layer
- Created reusable Terraform module library (DMS instance, endpoints, KMS keys) following Infrastructure-as-Code best practices and semantic versioning

Orchestration and Automation:

- Built sophisticated Step Functions workflow with Lambda (Python) orchestration enabling zero-touch database onboarding (3 commands: configure, authenticate, execute)
- Implemented DynamoDB-driven configuration catalog for centralized source management and metadata tracking
- Developed S3 data lake architecture with Redshift Spectrum integration for querying semi-structured data (JSON/Parquet) alongside relational tables

Security and Compliance:

- Deployed multi-AZ infrastructure with automatic failover capabilities
- Achieved AWS Security Hub compliance with automated security scanning
- Implemented KMS encryption with automatic key rotation policies
- Configured CloudWatch monitoring with custom metrics and automated alerting

Operational Excellence:

- Delivered 60-70% infrastructure reusability across source types through modular design
- Reduced database onboarding time from weeks to days via automation
- Created comprehensive documentation and runbooks for operations team
- Implemented automated testing for infrastructure changes

Technical Stack: AWS (DMS, Lambda, Step Functions, DynamoDB, S3, Amazon Redshift, Glue, Athena, KMS, Secrets Manager, CloudWatch), Terraform, Python, SQL Server, MongoDB, REST APIs, Security Hub

Quantifiable Impact:

- 10-30 databases replicated with 30-120 sec latency

- 60-70% infrastructure code reusability
- 3-command onboarding process (previously multi-week manual setup)
- Multi-AZ deployment with 99.9% uptime target
- Zero data loss with CDC implementation

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. Sole engineer responsible for design and implementation across multiple production systems.

Engagement 1: Educational Analytics Platform with LLM Integration (E-learning and Healthcare Client)

Sole engineer responsible for designing and implementing complete educational analytics platform serving thousands of nursing students with real-time performance insights and AI-powered recommendations.

Data Pipeline and ETL:

- Architected end-to-end ETL system extracting student performance data from MySQL to S3 Data Lake with Parquet optimization and user-based partitioning strategy
- Implemented EventBridge-scheduled daily extraction (24-hour cadence) with automated Glue Crawler management for schema evolution
- Developed real-time update pipeline using SQS for user-specific data refreshes without full ETL re-runs, reducing update latency from hours to minutes

API Development:

- Built production dynamic REST API (API Gateway + Lambda + Athena) supporting configurable analytics endpoints with zero-code deployment (DynamoDB-driven configuration)
- Implemented auto-generated API documentation endpoint listing all available analytics with calculation methods and examples
- Achieved sub-second query response times through optimized Athena partitioning and query patterns
- Designed multiple response formats: heatmaps, time series, performance summaries, detailed drill-downs

AI and LLM Integration:

- Developed LLM-powered analysis system using AWS Bedrock (Claude 3.5 Sonnet) and RAG (Retrieval-Augmented Generation) pattern for personalized educational recommendations
- Implemented NCLEX-RN exam question generator with iterative AI feedback loop for quality improvement
- Created methodology configuration system via DynamoDB enabling dynamic Prompt Engineering without code deployment
- Built batch processing system for flashcard generation with OAuth2 integration to client systems

Frontend and User Experience:

- Designed and deployed Vue.js frontend for NCLEX question generator with CloudFront CDN distribution
- Implemented basic authentication and session management
- Created intuitive UI for selective field regeneration and feedback submission

Infrastructure and DevOps:

- Deployed multi-stack CDK infrastructure (3 independent stacks: ETL, API, LLM) with cross-stack references
- Implemented 6 production Lambda functions (containerized) with proper error handling and logging
- Configured EventBridge schedulers for automated maintenance tasks
- Established Secrets Manager integration for secure credential management

Engagement Impact:

- 300+ student records processed per batch
- Sub-second API query response times
- 6 production microservices deployed
- 10+ analytics endpoints configured
- 24/7 system availability with automated monitoring
- Zero data breaches (FERPA compliant)

Engagement 2: API-to-Data Lake Pipeline (Confidential Client)

- Built dual-Lambda architecture for API data extraction and transformation using AWS CDK
- Implemented incremental processing system with timestamp-based change detection to prevent data duplication

- Developed automated JSON-to-Parquet conversion with partitioned storage strategy (year/month/day hierarchy)
- Configured AWS Glue Crawler for automated schema cataloging and Athena query optimization
- Designed pagination system for large-volume API data extraction with exponential backoff retry mechanisms
- Integrated Parameter Store for secure credential management and S3 server-side encryption (KMS)
- Created control file system for processing state management and data lineage tracking

Engagement Impact:

- 70% reduction in storage costs (JSON to Parquet)
- 10x faster Athena queries through partitioning
- Automated daily processing with zero manual intervention

Engagement 3: SageMaker Forecast System (Confidential Client)

- Implemented machine learning forecast infrastructure using SageMaker Canvas models for predictive analytics
- Created dual forecasting system supporting both short-term (180 days) and long-term (12 months) predictions
- Configured S3 data flows with input/output separation for model training and inference
- Set up GitHub Actions CI/CD pipeline with automated testing
- Implemented multi-environment deployment (dev/staging/prod) with CDK parameterization

Engagement Impact:

- Automated forecast generation reducing manual effort by 95%
- Multi-environment deployment capability
- Continuous delivery pipeline

Engagement 4: Healthcare Data Lake (Confidential Client)

- Designed and implemented serverless data processing architecture for medical records using AWS CDK
- Developed Python Lambda functions for JSON to Parquet data transformation across multiple medical record types
- Implemented robust error handling and logging with CloudWatch integration
- Designed and implemented data cataloging system using AWS Glue with automatic schema discovery
- Configured IAM roles and policies following least privilege principle

- Implemented data encryption using KMS and granular access control with Lake Formation

Engagement Impact:

- Significant reduction in storage costs through format optimization
- Improved query performance via Parquet and efficient partitioning
- HIPAA-compliant security implementation

Combined Technical Stack: AWS (CDK, Lambda Docker, API Gateway, Bedrock Claude 3.5 Sonnet, Athena, Glue, S3, DynamoDB, SQS, CloudFront, EventBridge, Secrets Manager, Parameter Store, KMS, Lake Formation, IAM, SageMaker Canvas, CloudWatch), GitHub Actions, Python, MySQL, Parquet, Vue.js, Docker, OAuth2

Argeniss Software | Bedrock AI Analysis Engineer

September 2024 - February 2025 | Innovative Client | Remote

Designed and deployed production multi-model AI consensus system for automated analysis workflows.

Key Responsibilities and Achievements

AI Architecture:

- Designed complete multi-model Generative AI pipeline using AWS Bedrock with 4 concurrent models: Claude Sonnet, Claude Haiku, Nova Lite, Nova Pro
- Built serverless architecture with Step Functions orchestrating 3-stage workflow: data preprocessing, parallel AI analysis, consensus output generation
- Implemented EventBridge triggers for automated CSV processing handling 300+ records per batch

System Development:

- Developed Lambda functions for multi-model consensus analysis with conflict resolution logic
- Created structured output generation supporting both JSON and CSV formats with schema validation
- Implemented DynamoDB job tracking system with complete lifecycle management (queued, processing, completed, failed)
- Built retry mechanisms with exponential backoff for API rate limiting

Performance and Reliability:

- Achieved 30-60 second processing time per account with cross-model validation
- Implemented parallel execution of AI models for maximum throughput
- Configured CloudWatch alarms for job failures and performance degradation

Technical Stack: AWS (Bedrock Claude/Nova models, Step Functions, Lambda, DynamoDB, S3, EventBridge, CloudWatch), Python

Quantifiable Impact:

- 300+ accounts processed per automated run
- 4-model consensus validation improving accuracy
- 30-60 second per-account processing time
- 95%+ job success rate

Scale AI | AI Quality Assurance Engineer

August 2023 - September 2024 | Remote

Contributed to training and improvement of generative AI models through systematic code review and quality assurance processes.

Key Responsibilities and Achievements

Model Training and Evaluation:

- Performed comprehensive code review of AI model outputs across multiple programming languages
- Created high-quality code snippets with detailed reasoning to improve model response accuracy
- Developed test cases and edge cases for model evaluation
- Identified and documented bugs in model behavior with reproduction steps

Quality Assurance:

- Established quality standards for AI-generated code outputs
- Reviewed model responses for correctness, efficiency, and best practices adherence
- Provided feedback on Prompt Engineering for improved model responses
- Escalated complex technical issues to ML engineering teams

Cross-Functional Collaboration:

- Worked with ML engineers to implement model improvements
- Contributed to documentation of model capabilities and limitations
- Participated in calibration sessions to ensure consistency across QA team

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

Impact:

- Reviewed 1000+ model outputs
- Contributed to improvement of model accuracy
- Identified critical bugs leading to model retraining

NTB Solutions | Senior Data Analyst

August 2018 - February 2024 | Multiple Healthcare Clients | Remote/Hybrid

Led data analysis and integration projects for healthcare institutions across multiple countries (Argentina, Paraguay, Uruguay), focusing on improving patient care, optimizing resources, and enabling data-driven decision-making.

Key Responsibilities and Achievements

Healthcare Analytics and BI:

- Analyzed complex healthcare data from multiple sources (EHR, hospital management systems, government registries, insurance claims)
- Built comprehensive ETL processes using Azure Data Factory, Databricks Notebooks, Scala, and Python
- Created executive dashboards using Power BI, Pentaho, and Tableau for C-level healthcare executives
- Designed and implemented health information systems (HIS, LIS, RIS) integrations

Data Integration and Interoperability:

- Integrated disparate healthcare systems using MIRTH Connect middleware
- Implemented HL7-FHIR health interoperability standards for cross-system communication
- Managed SQL and NoSQL databases (MySQL, Oracle, PostgreSQL, MongoDB) for multi-facility healthcare networks

- Developed Pentaho Data Integration ETL processes for automated data synchronization

Stakeholder Management:

- Consulted directly with healthcare executives, clinical staff, and IT teams to define project goals and success criteria
- Translated complex technical concepts into actionable business insights
- Managed stakeholder expectations across distributed teams and time zones
- Conducted training sessions for end-users on BI tools and dashboards

Major Achievement - Revenue Recovery System:

Built custom reporting solution using Pentaho Report Designer that identified previously uncontrolled revenue streams, recovering approximately \$6 billion pesos annually (2021 baseline) for healthcare institution. System automated detection of billing discrepancies and insurance claim gaps.

Technical Stack: Azure (Data Factory, Databricks), Python, Spark, PySpark, SQL, Pentaho (Data Integration, Report Designer), Power BI, Tableau, MIRTH Connect, MySQL, Oracle, PostgreSQL, MongoDB, HL7-FHIR

Quantifiable Impact:

- \$6B+ pesos annual revenue recovery through automated reporting
- 15+ healthcare facilities integrated across 3 countries
- 100+ dashboards and reports deployed
- 50+ ETL processes automated

Anyone AI | Machine Learning Engineer (Part-Time)

July 2022 - December 2022 | Educational Projects | Remote

Completed intensive machine learning bootcamp with multiple real-world projects demonstrating end-to-end ML development capabilities.

Projects Completed

Sentiment Analysis System:

- Built sentiment analysis model for movie streaming service product reviews
- Implemented word embeddings for text vectorization (Word2Vec, GloVe)
- Achieved 85%+ accuracy on test dataset

- Deployed via Flask API with Docker containerization

Image Classification Model:

- Developed CNN-based image classifier for vehicle make and model prediction
- Implemented data augmentation to improve model generalization
- Used transfer learning with pre-trained models (ResNet, VGG)
- Achieved 90%+ accuracy on validation set

Credit Risk Analysis:

- Built binary classification model for loan default prediction
- Performed feature engineering and hyperparameter tuning
- Implemented SMOTE for handling class imbalance
- Deployed predictive model with scikit-learn pipeline

Salary Prediction:

- Developed regression model for salary prediction based on job features
- Performed exploratory data analysis and feature selection
- Implemented ensemble methods (Random Forest, XGBoost)
- Created API endpoint for real-time predictions

Technical Stack: Python, Scikit-learn, Keras, PyTorch, TensorFlow, Pandas, NumPy, Docker, Flask, Jupyter Notebooks

Skills Developed:

- End-to-end ML model development lifecycle
- Model deployment and API integration
- Docker containerization for ML models
- MLOps best practices

Sanatorio la Entrerriana | Data Analyst and IT Manager

July 2008 - August 2018 | Parana, Argentina | On-site

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.

Key Responsibilities and Achievements

Major Project - HIS/LIS/RIS Implementation:

- 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
- Managed cross-functional team of clinical and technical staff through phased rollout across departments
- Coordinated change management across all departments (clinical, administrative, financial)

Data Analysis and Business Intelligence:

- Established data analytics function from ground zero
- Created executive dashboards and operational reports using Power BI and Google Data Studio
- Performed ad-hoc analysis for clinical and administrative decision-making
- Developed KPIs and metrics for hospital performance monitoring

Database Management:

- Managed Oracle databases for mission-critical healthcare applications
- Optimized SQL queries for report generation and data extraction
- Implemented backup and recovery procedures
- Ensured data security and HIPAA-equivalent compliance

Leadership and Team Management:

- Built and led IT team from 2 to 8 professionals
- Managed vendor relationships for software and hardware procurement
- Coordinated with clinical leadership on technology adoption strategies
- Conducted user training across the institution

Major Achievement:

Led successful adoption of the integrated HIS/LIS/RIS system, which remains in production use 12+ years later, supporting all clinical and administrative processes.

Technical Stack: Oracle Database, SQL, Power BI, Google Data Studio, Health Information Systems, Windows Server, Network Administration

Quantifiable Impact:

- Full institutional adoption of integrated HIS/LIS/RIS
- 12+ years system longevity in production

- Zero critical system failures post-implementation
- Replaced legacy paper-based clinical and administrative workflows

Technical Skills

Cloud Platforms and Infrastructure

AWS (Advanced): CDK, Lambda, Step Functions, ECS Fargate, ALB, WAF, Cognito, ACM, Route53, AWS DMS, AWS Bedrock (Claude API), API Gateway, S3, Glue, Athena, Amazon Redshift, DynamoDB, SQS, SNS, KMS, EventBridge, Secrets Manager, CloudFront, Cost Explorer, SageMaker Canvas, Lake Formation, IAM, CloudWatch

Infrastructure as Code: Terraform, AWS CDK

Azure: Data Factory, Databricks, Synapse Analytics, Fundamentals certified

Containerization and Orchestration: Docker, Container-based Lambda deployments

Generative AI and LLMs

AWS Bedrock (Claude 3.5 Sonnet, Haiku, Nova Lite, Nova Pro), Claude API (Anthropic), Prompt Engineering, RAG (Retrieval-Augmented Generation), LLM Integration, AI Agents, Generative AI, SageMaker Canvas

Programming and Scripting Languages

Expert: Python (Pandas, NumPy, PyArrow, Boto3), SQL (T-SQL, PL/SQL, PostgreSQL)

Proficient: TypeScript, JavaScript, R, Scala, Arduino C++ (ESP32, Embedded Systems)

Working Knowledge: [VB.NET](#)

Web Frameworks and APIs

FastAPI, Next.js 14, Vue.js, REST API Design, OAuth2, Stripe API Integration

Databases and Data Storage

Relational: PostgreSQL, MySQL, SQL Server, Oracle, Amazon RDS, Amazon Redshift

NoSQL: MongoDB, DynamoDB

Data Formats: Parquet, JSON, CSV, Avro

Data Warehousing: Star/Snowflake schemas, dimensional modeling, partitioning strategies

Data Engineering and ETL

ETL/ELT Tools: AWS Glue, Azure Data Factory, Pentaho Data Integration, SSIS, Custom Python pipelines

Big Data: Apache Spark, PySpark, Databricks

Data Integration: MIRTH Connect, Mule, HL7/FHIR standards

Stream Processing: Amazon Kinesis, SQS-based event-driven architectures

AI/ML and Advanced Analytics

Generative AI: AWS Bedrock (Claude 3.5 Sonnet, Haiku, Nova models), Prompt Engineering, RAG, AI Agents

ML Frameworks: Scikit-learn, PyTorch, Keras, TensorFlow, XGBoost, LightGBM

ML Deployment: SageMaker Canvas, Docker-based model serving, API integration

Statistical Analysis: Hypothesis testing, regression analysis, time series forecasting

Business Intelligence and Visualization

Power BI, Tableau, Amazon QuickSight, Pentaho, Google Data Studio, SSRS, SSAS

DevOps and CI/CD

Version Control: Git, GitHub, Bitbucket

CI/CD: GitHub Actions, automated testing pipelines

Monitoring: CloudWatch, Application Insights

Security: KMS encryption, IAM policy design, Security Hub compliance, least privilege principles

Methodologies and Frameworks

Agile/Scrum (Certified), DevOps, ITIL, Infrastructure as Code (IaC), Test-Driven Development, Documentation-as-Code

Featured Projects and Achievements

CloudPruneAI - Production FinOps SaaS Platform

Role: Founder and Technical Lead | Timeline: November 2025 - Present

Project Overview:

Production B2B2B FinOps SaaS platform that automatically analyzes AWS infrastructure and auto-generates Infrastructure-as-Code in 3 formats (CDK Python, Terraform HCL, CloudFormation YAML) for cost optimization, distributed via a partner-led GTM model.

Technical Architecture:

- Hybrid serverless: FastAPI on Lambda + Mangum (API) + ECS Fargate (scan workers) + async RDS PostgreSQL
- 18 specialized analyzers covering 31+ optimization checks across compute, storage, networking, databases, commitments, extended support, and AI/ML services
- AI-powered IaC code generation using Claude API with multi-format output and business context enrichment via Claude Haiku
- Multi-tenant architecture with Auth0 authentication
- Monorepo with Turbo unifying TypeScript and Python tooling
- Stripe payment integration with gainshare pricing plus recurring Operate Plan subscriptions

Key Achievements:

- 4 active partners across Argentina, Spain, and Uruguay
- First real partner scan detected USD 5,260/year in waste across a client AWS account
- 4-month delivery from concept to live partner channel
- USPTO trademark registered (Serial: 99585602)
- GDPR-compliant architecture with Privacy Policy, DPA, ROPA, and Security Posture Document

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

Production URL: app.cloudpruneai.com

Quira - Multi-Tenant Industrial IoT Platform

Role: Founder and Technical Lead | Timeline: December 2025 - Present

Project Overview:

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.

Validated end-to-end with real ESP32 hardware. Production launch April 19, 2026.

Technical Architecture:

- Full AWS stack via 9 CDK stacks: ECS Fargate behind ALB with WAF, Route53 + ACM, RDS PostgreSQL 15, Cognito, SQS, 6 Lambdas, SNS
- FastAPI backend with 47 endpoints and 103 tests covering ingestion, alerts, devices, organizations, sites, reports, webhooks
- HMAC-SHA256 + HKDF device authentication with idempotent client-side deduplication
- Async report generation via WeasyPrint workers (PDF/CSV)
- Sensor-agnostic data model and driver registry to support multi-sensor expansion (CO2, pressure, vibration on V2 roadmap) without schema changes
- Next.js 14 admin panel with 25 pages, full ES/EN i18n via next-intl, Recharts visualizations
- ESP32 + SHT31 firmware in Arduino C++ with HMAC signing

Key Achievements:

- Concept to production in under 4 months, solo
- End-to-end validated with real ESP32 + SHT31 hardware: readings reach the platform, trigger alerts, and render in production charts
- First partner LOI signed in Colombia (aerospace warehouse pilot, ~30 devices)
- B2B2B revenue model with tripartite split between platform, partner, and vendor

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, Recharts, ESP32 (Arduino C++), SHT31, WeasyPrint, Python, TypeScript, Resend

Multi-Source DMS Replication Architecture

Role: Lead AWS Data Integration Engineer | Timeline: October 2025 - Present

Project Overview:

Enterprise-scale automated database replication system supporting multiple source types (SQL Server, MongoDB, REST APIs) with unified orchestration.

Technical Architecture:

- AWS DMS for SQL Server CDC replication (30-120 sec latency)
- Step Functions orchestration with Lambda automation
- DynamoDB configuration catalog for centralized management
- Terraform modules for infrastructure reusability

- S3 data lake with Redshift Spectrum integration

Key Achievements:

- 60-70% infrastructure reusability across source types
- 3-command onboarding process (previously weeks of manual work)
- Multi-AZ deployment with automatic failover
- Security Hub compliant with automated scanning

Tech Stack: AWS (DMS, Step Functions, Lambda, DynamoDB, Redshift, S3, Glue), Terraform, Python, SQL Server

Impact: 10-30 databases replicated, weeks-to-days onboarding time reduction

Educational Analytics Platform with LLM Integration

Role: Solo AWS Data Engineer | Timeline: March 2025 - October 2025

Project Overview:

Complete educational data platform serving thousands of nursing students with real-time performance analytics and AI-powered insights using RAG pattern and LLM integration.

Technical Architecture:

- MySQL-to-S3 ETL pipeline with Parquet optimization
- Dynamic REST API (API Gateway + Lambda + Athena) with DynamoDB configuration
- AWS Bedrock integration (Claude 3.5 Sonnet) for personalized recommendations using RAG
- NCLEX question generator with Vue.js frontend
- Real-time SQS update pipeline for user-specific refreshes

Key Achievements:

- Sub-second API response times through optimized Athena queries
- 300+ student records processed per batch
- 6 production microservices deployed via CDK
- 10+ analytics endpoints with zero-code configuration
- FERPA compliant security implementation

Tech Stack: AWS (CDK, Lambda, Bedrock, API Gateway, Athena, Glue, S3, DynamoDB, SQS, CloudFront), Python, MySQL, Vue.js

Impact: 24/7 availability, automated insights for thousands of students

Healthcare Revenue Recovery System

Role: Senior Data Analyst | Timeline: 2018-2024 | Client: NTB Solutions

Project Overview:

Custom reporting solution that identified and automated recovery of previously uncontrolled healthcare revenue streams.

Technical Implementation:

- Built custom reports using Pentaho Report Designer
- Integrated data from multiple healthcare systems (EHR, billing, insurance)
- Implemented automated discrepancy detection algorithms
- Created executive dashboards for revenue tracking

Key Achievement:

\$6 billion pesos annual revenue recovery (2021 baseline) through automated detection of billing discrepancies and insurance claim gaps.

Tech Stack: Pentaho (Report Designer, Data Integration), SQL, Oracle, MySQL

Impact: Billions in annual revenue recovered, ongoing automated monitoring

Education and Certifications

Bachelor's Degree in Business Management

Universidad Autonoma de Entre Rios (UADER), Parana, Argentina | 2017

Professional Certifications and Advanced Training

In Progress

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

AWS Cloud Services (2024)

- Introduction to Amazon Kinesis Analytics (LATAM Spanish)
- Amazon EMR Getting Started
- AWS Glue Getting Started
- Amazon Athena
- Amazon QuickSight Getting Started and Visualization
- AWS IoT Analytics Getting Started
- Amazon Redshift Getting Started

Microsoft Azure Ecosystem (2023)

- Get started with data engineering on Azure
- Microsoft Azure Fundamentals (Cloud concepts, Architecture, Management)
- Data engineering with Apache Spark on Azure Synapse
- Data analysis solutions with serverless SQL groups on Azure Synapse
- Data modeling, querying and exploring on Azure Synapse
- Data storage and transformation with Data pipelines on Azure Synapse Analytics

Machine Learning and AI (2022)

- Machine Learning Engineer - Anyone AI
- Introduction to AI (Elements of AI) - University of Helsinki

Project Management and Methodologies (2020-2021)

- Project Manager Officer - National Technological University (FRBA)
- Scrum Foundation Professional Certificate - CertiProf
- Scrum Grand Master - National Technological University (FRBA)

Language Proficiency

- C1 English Language - EFSET (2023)
- B2.2 English - Freedom Private Institute (2023)

IT Service Management

- ITIL Foundation Certificate in IT Service Management - PeopleCert (2018)
- ITIL V3 - National Technological University (FRBA)

Database and Programming Fundamentals (2016-2020)

- Database programming (PL-SQL Language) - National Technological University
- Database Fundamentals - National Technological University

- Programming Fundamentals - e-Learning Center UTN FRBA
- R Language Programmer - Edutin
- AWS Cloud Practitioner Essentials, CloudFormation, ECS, EC2, Serverless development

Core Competencies

Technical Leadership

- Solution Architecture: Design scalable, production-ready cloud architectures balancing performance, cost, and security
- Technical Strategy: Translate business requirements into technical roadmaps and implementation plans
- Team Building: Recruit, mentor, and develop high-performing technical teams
- Vendor Management: Negotiate and manage relationships with technology vendors and service providers

Data Engineering Excellence

- Pipeline Design: Create robust, fault-tolerant data pipelines with comprehensive error handling
- Performance Optimization: Optimize queries, data structures, and processing workflows for maximum efficiency
- Data Quality: Implement data validation, cleansing, and quality monitoring systems
- Cost Optimization: Design cost-effective data architectures through smart partitioning and storage strategies

AI and ML Integration

- LLM Integration: Production experience with AWS Bedrock (Claude models) for business applications
- Prompt Engineering: Design effective prompts for consistent, high-quality AI outputs
- RAG Implementation: Build Retrieval-Augmented Generation systems for context-aware AI responses
- Multi-Model Systems: Build consensus systems leveraging multiple AI models for improved accuracy

Cloud Architecture

- Hybrid Serverless: Design architectures combining Lambda, containers, and managed services optimally
- Infrastructure as Code: Comprehensive CDK and Terraform expertise for repeatable deployments
- Security and Compliance: Implement secure architectures meeting industry standards (HIPAA, FERPA, Security Hub)
- Cost Management: Architect cost-effective solutions through right-sizing and resource optimization

Business and Leadership

- Stakeholder Management: Effectively communicate technical concepts to non-technical audiences
- Project Management: Deliver complex projects on time and within budget using Agile methodologies
- Change Management: Lead organizational transformation initiatives with cultural change components
- Strategic Planning: Align technology initiatives with business objectives and ROI requirements

Soft Skills and Professional Attributes

Collaboration and Communication

- Remote Work Excellence: 5+ years experience with distributed teams across multiple time zones
- Cross-Functional Leadership: Proven ability to work with clinical, business, and technical stakeholders
- Written Communication: Strong documentation skills with comprehensive technical writing experience
- Verbal Presentation: Confident presenter to executive leadership and technical audiences
- Multilingual: C1 English proficiency, native Spanish speaker

Problem-Solving and Innovation

- Analytical Thinking: Break down complex problems into actionable solutions
- Creative Problem-Solving: Find innovative approaches to technical and business challenges
- Critical Thinking: Evaluate multiple solutions objectively with data-driven decision-making

- Continuous Learning: Proactively learn emerging technologies and industry best practices
- Growth Mindset: Embrace challenges as opportunities for professional development

Professional Excellence

- Detail-Oriented: Meticulous attention to code quality, documentation, and system reliability
- Goal-Driven: Results-focused with track record of delivering measurable business outcomes
- Diplomatic Approach: Navigate complex organizational dynamics with professionalism
- Strategic Thinking: Balance short-term deliverables with long-term architectural vision
- Team Player: Collaborative approach fostering positive team dynamics and knowledge sharing
- Accountability: Take ownership of deliverables and proactively manage risks

Leadership and Management

- Team Development: Mentor junior engineers and facilitate knowledge transfer
- Conflict Resolution: Mediate technical disagreements and build consensus
- Adaptability: Thrive in fast-paced environments with changing priorities
- Emotional Intelligence: Navigate stakeholder concerns with empathy and professionalism
- Decision Making: Make informed technical decisions under uncertainty with calculated risk assessment

Tools and Collaboration Platforms

Development and Version Control: Git, GitHub, Bitbucket, Visual Studio Code, PyCharm, Jupyter Notebooks

Project Management: Jira, Trello, ClickUp, MS Project, Asana

Collaboration and Communication: Slack, Discord, Microsoft Teams, Zoom, Google Meet

Design and Documentation: Figma, [Draw.io](#), Lucidchart, Confluence, Notion, Markdown

CI/CD and DevOps: GitHub Actions, Docker, Docker Compose, AWS CloudFormation, Terraform

Additional Information

Languages

- Spanish: Native proficiency
- English: C1 Advanced proficiency (EFSET certified)

Work Authorization

- Argentina: Citizen
- Remote Work: Available for remote positions globally

Availability

- Current Status: Open to full-time opportunities, consulting engagements, and contract work
- Notice Period: Negotiable based on project commitment
- Work Schedule: Flexible, experienced with multiple time zones (EST, PST, GMT)

Professional Networks

- Active member of AWS community forums
- Contributor to open-source projects (GitHub: gneironiar)
- Regular attendee of cloud computing and data engineering webinars

Professional References

Available upon request.

Last Updated: April 2026