

Dmytro Hryniuk

Cloud & Platform Engineer | AWS, Kubernetes, Terraform | IoT & Embedded Systems

Figueira da Foz, Portugal | +351 928 325 147 | hryniuk.mitya@gmail.com | linkedin.com/in/dmytro-hryniuk | github.com/DmytroHryniuk

Open to relocation within Europe.

Ukrainian citizen, resident in Portugal. EU Blue Card eligible (MSc Computer Science).

Profile

Cloud and Platform Engineer, 5+ years on AWS, grown from delivering code into owning architecture and operations. Those years went into building and running the platform behind analytics products that a major silicon vendor and its OEM customers rely on for hardware and cloud choices: public recommendation services, customer data processing engines, automated pricing and analytics pipelines, an unattended cloud benchmarking fleet and LLM features in production, each owned from design through delivery to uptime.

Now extending that depth into Kubernetes and Terraform, with the CKA in progress and a self-hosted k3s platform on Hetzner Cloud, its own control plane and workers built from scratch, as the daily proving ground. It runs as the managed-service layer a team would otherwise rent from a cloud: private networking with the API reachable only over an SSH tunnel, GitOps delivery, automatic TLS, Prometheus and Loki with custom rules that flag stale backups, isolated Supabase tenants with nightly backups and a rehearsed disaster recovery, and zero-downtime rolling upgrades across Kubernetes minor versions with etcd snapshots and checksum-verified binaries.

The same end-to-end approach carries over to hardware: schematic, PCB and firmware on Espressif and Nordic silicon, through to the cloud service the device reports to, from LoRa sensor nodes in a garage to a standalone e-ink educational device, all backed by that same platform. Strongest where hardware, data and cloud meet.

Experience

eXalt Solutions, Inc., Boston, US. Remote, independent contractor.

Apr 2021 – Present

US enterprise software vendor whose analytics turn performance and cost data into infrastructure decisions for hardware and cloud buyers. Responsible for the AWS platform behind those products across its whole lifecycle, from architecture and delivery to production operations.

Client engagement under NDA: customers and products are described without names.

Cloud & Platform Engineer

Sep 2025 – Present

- Own and run the serverless recommendation platform behind the public advisory tool a major silicon vendor offers its customers: infrastructure as code with AWS CDK in Python (API Gateway, Lambda, DynamoDB, S3, EventBridge), one stack definition for every stage from sandbox to production, deployed stage by stage with post-deploy Synthetics canaries and alarms as the release gate.
- Designed the analyst delivery path for the recommendation models: a notebook dropped in S3 is validated, compiled into the API handler and promoted through the same stages and canary checks as any release, with rollback to the previous handler on failure, so algorithm changes reach production without hand-written application code. The pattern now runs multiple production APIs.
- Ran the performance side of a price-performance program: SPEC CPU2017, STREAM and FIO across hundreds of EC2 instance types in every current family, on a self-terminating cloud-init fleet, within a fixed budget, with pilot campaigns on GCP and Azure. Results joined with the pricing dataset for cost-per-performance comparisons and published as branded reports behind CloudFront.
- Shipped LLM features to production: Amazon Bedrock schema mapping inside a Step Functions Distributed Map for customer data intake, a Claude personalization endpoint on recommendation results, a Claude vision pipeline with tool-forced structured output, and an asynchronous LLM job pattern on SQS and DynamoDB.
- Built the product analytics pipeline on Step Functions and Lambda: daily usage telemetry validated and aggregated into partitioned Parquet on Athena and Apache Iceberg, with MERGE upserts for month-end restatement, SNS reporting and a one-command backfill.
- Set the operating conventions for every CDK stack and enforced them: TTL on transient DynamoDB tables, lifecycle rules on working buckets, retention on every log group, shared dependencies in Lambda layers rebuilt per runtime, and a canary-and-alarm pattern for each new API.
- Led an estate-wide secrets audit across every repository and deployed Lambda function and moved every service off long-lived keys onto IAM roles and Secrets Manager.
- Handled production incidents end to end and turned each into a control: an empty upstream pricing feed became an abort gate on the weekly refresh; a rollback that fired on every failure became per-query deadlines and stage-level alerts on the analytics pipeline; a dead-letter handler's first poison message became idempotent processing with an explicit retry policy.

Full Stack AWS Cloud Engineer

Sep 2022 – Sep 2025

- Re-architected the customer data processing engine, which turns spreadsheet exports from OEM and enterprise customers into normalized datasets, from a Lambda pipeline into a containerized Java 17 / Spring Boot 3 service on ECS Fargate: high-memory tasks, S3-event triggers, DynamoDB status tracking, ECR image rollouts and a declarative set of transformation actions, removing the Lambda time and memory ceilings.
- Engineered the large-dataset processing layer of that engine with Apache POI: streaming parsers, JVM memory management (ZGC), cross-file joins and a custom formula engine for inputs of several hundred megabytes, validated against a regression suite of real customer data.

- Developed the multi-cloud pricing dataset behind the cost models: AWS, Azure and GCP price lists collected, normalized to one schema and filtered to comparable SKUs, refreshed weekly with change detection, and served through Athena for price-performance and TCO comparisons.
- Maintained the product's front-end runtime (Webpack, Chart.js 4) and built its embeddable LLM chatbot widget with async submit-and-poll and file upload.
- Built the visual-regression QA service that runs before every product release: Puppeteer screenshots on ECS Fargate, pixel diffs and an annotated Excel report; and the headless-Chrome print service that replaced a legacy ESB export path for customer-facing PDFs.
- Inventoried the Lambda estate and ran the Python runtime cleanup: triaged every function, retired the dead ones with archive-and-redeploy tooling, and upgraded the rest to Python 3.12.

Backend & Cloud Engineer

Apr 2021 – Sep 2022

- Built the first serverless generation of the customer data processing engine on AWS Lambda (Node 18, ExcelJS, HyperFormula): validation, action execution, S3-event architecture and DynamoDB status tracking.
- Built the customer benchmark intake pipeline, portal upload to S3 to JSON to Athena tables and views with SES notifications, still running in production five years later.
- Owned deployments end to end: Lambda packaging, IAM roles, S3 event wiring and DynamoDB data modelling across Node.js and Python services.

Earlier roles

Flutter & IoT Developer, independent contractor for a German pet-tech startup. Prague, Czechia, Feb 2021 – Jan 2022. Raspberry Pi based remote pet-monitoring and live-streaming system: cellular connectivity, RTMP/HLS video, a Firebase and AWS backend and a Flutter app.

Flutter Developer, Dynamic Division. Remote, May 2021 – Oct 2021. Second engagement with the same client: rebuilt the industrial cleaning-robot control app in Flutter for the startup's seed-stage product, replacing the 2019 native Android proof of concept; robot control over MQTT.

Flutter Developer, MySpool LLC. Remote, Aug 2020 – Apr 2021. Companion app for a family of Wi-Fi sensor devices: onboarding over the device's own access point, alerts and push notifications, per-sensor thresholds and charts, on an AWS Amplify backend (Cognito, Lambda).

Web Application Developer, SHERP / Quadro International. Kyiv, Sep 2020 – Feb 2021. Spare-parts catalogue for a Canadian all-terrain-vehicle dealer: exploded-view drawings with clickable image maps, revision-aware part tables and e-mail ordering, for three vehicle models.

Mobile Application Developer, Sensorium. Kyiv, Sep 2019 – Nov 2020. Android application for Fourier's einstein Tablet+ classroom science tablets: physics and chemistry lessons as live experiments, driven by built-in sensors, external lab probes and Bluetooth sensors through the vendor SDK, with real-time charts and a USB camera feed; three languages, three tablet generations.

Mobile Application Developer, Dynamic Division. Kyiv, Oct 2019 – Dec 2019. Proof-of-concept control app for a fleet of industrial cleaning robots over MQTT.

Skills

Cloud & AWS: AWS, AWS CDK (Infrastructure as Code, IaC), Lambda, API Gateway, Step Functions, EventBridge, SQS, SNS, SES, DynamoDB (Streams), S3, ECS Fargate, ECR, EC2 (AMI builds, cloud-init), ALB, CloudFront, Route 53, ACM, IAM, VPC (peering, endpoints), Secrets Manager, CloudWatch (Synthetics, alarms), Athena, Cognito, Amplify, Hetzner Cloud (VMs, Object Storage), Cloudflare

GenAI & LLM: Amazon Bedrock, Anthropic Claude API, LLM integration, tool use and structured output, vision, prompt engineering, async LLM orchestration (SQS, DynamoDB), agent-ready API schemas

Platform & DevOps: Kubernetes (k3s), Terraform, ArgoCD (GitOps), Helm, Docker (buildx multi-arch), GitHub Actions, observability (Prometheus, Grafana, Loki), cert-manager, Traefik, nginx, Linux administration (Ubuntu, systemd, hardening), Lambda layers, secrets scanning and security remediation, backup and disaster recovery (pg_dump, restore drills), runbooks and ADRs

Programming Languages: Python (boto3, pandas, Jupyter), TypeScript / Node.js (Express), Java 17 (Spring Boot 3, Apache POI, JVM tuning), Bash

Data & Databases: Apache Iceberg, Parquet, Athena pipelines, Glue Data Catalog, DynamoDB, PostgreSQL, MongoDB Atlas, MySQL

Web & App: Next.js, React, Tailwind CSS, Supabase, Firebase (Firestore, Cloud Messaging, Cloud Functions), Webpack, Chart.js, Puppeteer, Flutter

IoT, Embedded & Hardware: Espressif, Nordic Semiconductor, RAKwireless, Heltec, Zephyr / nRF Connect SDK, LoRa, BLE, MQTT, KiCad, EasyEDA, Onshape

Multi-cloud & Performance: GCP Compute Engine, Azure Virtual Machines, cloud pricing APIs, performance benchmarking (SPEC CPU2017, STREAM, FIO), cloud cost optimization (FinOps), vendor-optimized compiler toolchains

Education

M.Sc., Computer Science, National Pedagogical Dragomanov University, Kyiv 2018 – 2020

B.Sc., Software Engineering, National Pedagogical Dragomanov University, Kyiv 2016 – 2018

Associate, Computer Software Engineering, Kyiv National Economic University, Kyiv 2012 – 2016

Languages

Ukrainian, native. English, B2 professional. Portuguese, B1 intermediate. Spanish, currently learning.