

Hesham Meneisi

Head of Data and Analytics

heshammeneisi.github.io · heshammeneisi@gmail.com · +66 858 250 172 · github.com/HeshamMeneisi · linkedin.com/in/hesham-meneisi

SUMMARY

Hands-on data leader with 8+ years building data platforms for **payments, lending, and e-commerce** companies. Usually hired into chaos: near-daily failures and books that don't reconcile. Restructured an analytics org so agents handle the routine work and the pipelines still reconcile.

Focused on measuring whether AI actually works, with instrumentation and defensible numbers for executives.

TECHNICAL SKILLS

AI & Agents	Agent orchestration (3-7 personal, 10-15 team daily), subagent delegation, closed-loop agent workflows, prompt/context engineering, LLM-optimized codebase design
Languages	Python, SQL, TypeScript, C#/.NET, Go, Scala, Terraform, Shell
Data & Orchestration	Spark, Kafka, KSQL, Airflow, Dagster, dbt, Snowflake, Redshift, BigQuery
Monitoring & Visualization	Lightdash, PowerBI, Superset, Grafana (LGTM), ELK, Soda
Cloud & Infra	AWS, GCP, Azure, Kubernetes, Terraform, GitOps, GitHub Actions, Redis

EXPERIENCE

Onramp, Head of Data and Analytics

April 2025 - Present

- Restructured the data function to an AI-first operating model: 3 engineers running 10-15 parallel agents daily. Won executive buy-in by demonstrating critical dashboards built and validated in hours vs days under the old structure
- Held agent-assisted pipelines to audit-grade accuracy: **EUR 10M+ monthly volume reconciled between internal data, acquirers, banks, and wallets down to rounding errors (under EUR 5)**
- Replaced **188 static dashboards with 30** governed self-serve dashboards plus an org-wide Claude SQL skill, adopted by operations, finance, compliance, and support; data-team requests fell **from 20+/month to 0-2**
- Built a lakehouse platform on Azure (300+ governed assets) that retired the fragmented legacy stack; **reduced infrastructure cost by 43%**
- Designed the platform so agents can operate it directly: everything as code, fixed repo patterns, and an ADR for every architecture decision; agents open reviewable cross-repo PRs

Honest, Lead Data Platform Engineer (L3)

June 2023 - April 2025

- Built a self-service data mesh that pulled siloed team data into one cataloged platform
- Reduced monthly incidents from **60+ to under 10** via reliability overhaul and runbook standardization
- Halved cash-volume loss** by reworking the Probability of Default model: dropped low-signal legacy features, replaced them, and added a new data source
- Established data governance: data contracts, SLAs, ownership model, standardized feature definitions

Omise, Lead Data Engineer (team of 7)

Sep 2021 - June 2023

- Rebuilt the payment pipelines; failures fell from near-daily to a few a month, scans from **terabytes to gigabytes**
- Mentored mid-level engineers to senior
- Shipped a third-party integration **in under a month against a multi-month estimate** with a 3-engineer team

Earlier Roles

2017 - 2021

- Senior Backend Engineer, Crea (2021) - cut warehouse query time from **2 min to 1.5 s**; Devcurate (2019-21); Leastra (2018-19); Valeo Algorithms Intern (2017)

Solo Founder

- Shipped 6 products in 6 months** across mobile and web, from build through marketing and distribution

EDUCATION

Alexandria University - BS, Computer & Communication Engineering

2013-2018, Major GPA 3.92/4.0

Udacity - Machine Learning Nanodegree

2018

IELTS Band 8; Native Arabic