

Principal AI & Data Engineer

📍 **Location:** Onsite in Spokane, WA. | Remote considered for candidates in PST

Build AI-Powered Systems at a High-Growth Startup

At **Vega Cloud**, we're a **fast-moving startup** building **AI-driven FinOps solutions** that help businesses optimize their cloud costs. We're looking for a **Principal AI & Data Engineer**—a **hands-on SME** who thrives in the fast-paced world of startups and loves **building, iterating, and scaling AI-powered data systems from the ground up**.

This role is for a **true unicorn**—someone who can architect **scalable data pipelines**, deploy **ML models in production**, and **turn AI research into real-world impact**. If you're comfortable with **ambiguity, love solving hard problems, and thrive when given autonomy**, this is the role for you.

What You'll Do

- Design and optimize scalable data pipelines using Apache Kafka/Pulsar, Spark/Flink, and Iceberg/Delta Lake.
- Develop and maintain a high-performance lakehouse, ensuring governance, cost tracking (tenant tagging), and schema management.
- Build, deploy, and monitor ML pipelines using MLflow, Kubeflow, and model serving frameworks (KFServing, Seldon).
- Integrate HPC (high-performance computing) resources for large-scale AI/ML workloads.
- Own the development of AI-powered FinOps features, including cost forecasting, anomaly detection, and optimization recommendations.
- Work directly with product, engineering, and leadership to shape our AI/ML strategy.
- Thrive in an early-stage startup—where things move fast, challenges are complex, and solutions aren't always obvious.

You're a Great Fit If You Have:

- 7+ years of experience in data engineering, AI/ML, or a related field.
- Startup experience—you've worked in fast-paced environments where you had to own solutions, iterate quickly, and adapt on the fly.
- Expertise in modern data architectures (lakehouses, ETL/ELT pipelines, streaming frameworks).
- Hands-on experience with ML frameworks (TensorFlow, PyTorch, XGBoost) and MLOps tools (MLflow, Kubeflow).
- Strong programming skills in Python, Scala, or Java for data/ML workflows.
- Cloud and containerization experience (AWS/GCP/Azure, Kubernetes, Docker).
- A builder's mentality—you see a problem, design a solution, and execute.

Bonus Points If You Have:

- Worked in startups or fast-scaling companies.
- Experience with vector databases, LLMs, or Retrieval-Augmented Generation (RAG).

- Knowledge of FinOps, cloud cost optimization, or AI-powered cost analytics.

What Success Looks Like

In Your First 30 Days, You Will:

- Get familiar with our **data architecture and ML pipelines**.
- Deploy **small fixes and improvements** to existing workflows.

In 60 Days, You Will:

- Start **building and optimizing key data pipelines**.
- Improve **ML model deployment and monitoring processes**.

In 90 Days, You Will:

- **Own and drive** the development of a major AI/ML-powered feature.
- Collaborate with cross-functional teams to integrate **new predictive models** into production.

In 1 Year, You Will:

- Have built and optimized **end-to-end AI and data engineering pipelines**.
- Be recognized as a **technical leader** and AI/ML expert within the company.
- Establish **best practices** for data, AI, and MLOps workflows.

Why Join Vega Cloud?

- **Unlimited PTO** – because work-life balance matters.
- **100% medical coverage** – we take care of our people.
- **A chance to shape the future of AI-driven FinOps technology**.
- **A collaborative, innovation-driven culture** where your work makes a real impact.

 **Location Requirement:** This role is onsite, but we will consider remote candidates **only if they are located in the Pacific Time Zone (PST)**.

 **Compensation:** Salary is **negotiable** based on experience and expertise.

The Interview Process

1. **Intro Call with the VP of Engineering** – Get to know Vega Cloud and discuss your experience.
2. **Offline Technical Test** – Show off your problem-solving and coding skills.
3. **Tech Interview with the Engineering Team** – Dive into the details of the role and assess technical fit.
4. **Call with the CEO/Founder** – Get Excited
5. **Final Call with the VP of Engineering** – Wrap up and discuss next steps

 **Ready to push the boundaries of AI and data engineering?** Apply now and help us redefine the future of FinOps!