Matty Serapio

Analytics Engineering | Data Products | Infrastructure Capacity and Utilization Analytics
Denver, CO | linkedin.com/in/matthewserapio | serap.io

PROFESSIONAL SUMMARY

Data and analytics professional with 11 years of experience turning infrastructure telemetry into metrics, data products, and self-service analytics for cloud capacity, utilization, and cost at hyperscale. Expert in SQL (Presto, Trino, Spark SQL, Databricks SQL), Python, Airflow, ETL, and data warehousing on Databricks and Amazon Web Services (AWS), and in designing semantic layers and ontologies that give engineers, executives, and AI agents consistent business definitions and trustworthy metrics. Co-built a self-service analytics platform adopted by 1,500 users and delivered over $70MM in documented annual infrastructure savings.

EXPERIENCE

Databricks, Denver, CO March 2025 to Present

Delivery Solutions Architect, Strategic Accounts, Travel and Hospitality

- Build performance benchmarks and cost visibility tooling where none existed, defining a two-sided KPI framework that measured 9x faster job execution and 97% lower compute cost for a platform migration decision
- Lead migration of Trino SQL workloads to Databricks SQL (DBSQL) for a strategic account, converting queries and data pipelines to the lakehouse
- Model cost per query attribution from Redshift query telemetry, producing a chargeback dataset that ties compute consumption to consuming teams and quantifies modernization savings
- Monitor CPU and memory utilization for serverless Machine Learning (ML) inference workloads, right-sizing capacity to balance cost, performance, and production reliability
- Optimize Apache Spark and Spark SQL pipelines across data engineering and lakehouse architecture for two of Databricks' largest Travel industry strategic accounts
- Lead migration of self-hosted notebooks on Amazon Elastic Kubernetes Service (EKS) to managed Databricks, delivering $2MM in cost savings and eliminating infrastructure support overhead
- Unify the ML model lifecycle in a single environment for ML scientists and engineers, removing development to production handoffs and increasing A/B testing velocity

Amazon Web Services, Denver, CO July 2020 to March 2025

Senior Optimization Solutions Architect, Strategic Accounts and Financial Services

- Advised 70 enterprise accounts, including 5 of the top 10 AWS customers by spend (about $1B in annual recurring revenue), on FinOps strategy and the KPIs they used to measure the value of their cloud investments
- Co-developed a zero to one self-service data product in Python and Streamlit that scaled the expertise of 40 cloud financial operations (FinOps) specialists to 1,500 global users
- Built its LLM chat interface to fan each question out to multiple models, an early multi-agent pattern, and instrumented adoption metrics by compute configuration
- Built SQL cost allocation models on the AWS Cost and Usage Report (CUR) that mapped spend to multi-level business unit hierarchies and reallocated shared costs to consuming teams for chargeback
- Drove Graviton and AMD adoption campaigns using CUR-based treemaps of spend and adoption by business unit, delivering $40MM in annual infrastructure savings, including $5MM for a single customer
- Defined capacity and utilization metrics for these accounts, tracking identified vs. actioned savings month over month and driving $30MM in annual savings from optimization recommendations
- Designed a Kubernetes cost and utilization data model with AWS Split Cost Allocation Data, attributing EKS workload cost by business unit to enable chargeback and route unallocated capacity to the platform team
- Partnered with ML engineering teams on GPU capacity and cost optimization, matching instance types to training (P4d) and serving (G4dn) workloads based on utilization
- Presented 5 AWS videos, including The Keys to AWS Optimization series, teaching cost reallocation in SQL, CUR analysis, and rightsizing
- Recognized with the 2024 AWS OPTICS All Star Award and 15 Culture Cloud Awards

Liberty Energy, Denver, CO February 2018 to July 2020

Engineering Data Team Lead, Data Engineering and Analytics Engineering

- Led a team of 7 data specialists across data engineering and analytics engineering, building ETL pipelines, PostgreSQL data models, and dashboards for engineering and operations
- Contributed to the data and semantic layer behind FracTrends, a self-service well analytics tool covering 60,000+ horizontal wells with daily updates, used by engineers and customers
- Built Tableau dashboards for internal teams and customers, and replaced manual data entry with workflows for comparing completion designs and proppant, saving the engineering team 400 hours annually
- Automated Excel reporting workflows with R scripts, saving 100 hours of analyst time annually

ConocoPhillips, Houston, TX May 2014 to October 2017

Infrastructure Performance Analyst (FinOps) and Business Analyst

- Built TCO models for migrating 10% of on-premises compute to AWS, with Elastic Compute Cloud (EC2) sizing recommendations that saved $340K annually
- Managed budget and operations for a $200K high performance computing (HPC) project
- Served as systems analyst for an optimization software deployment that produced $600K in savings

EDUCATION

University of Pennsylvania August 2022

Master of Computer and Information Technology (Computer Science), GPA 3.84

University of Colorado Boulder May 2014

Bachelor of Science in Business Administration, GPA 3.74

TECHNICAL SKILLS

Data and Analytics Engineering: SQL, Presto, Trino, Spark SQL, Databricks SQL, Python, ETL, Apache Airflow, data modeling, semantic layers, ontology design, data warehousing, metrics development, data pipelines, Apache Spark, Delta Lake, Apache Iceberg, Apache Kafka

Visualization and Self-Service: Databricks AI/BI dashboards, Tableau, Looker, Power BI, Amazon QuickSight, Streamlit, Spotfire

Infrastructure Domain: Capacity and utilization analysis, cloud cost optimization (FinOps), Kubernetes, GPU infrastructure, Graviton, Terraform, Databricks, AWS

AI and ML: LLM applications, AI-assisted development (Claude Code, Cursor), MLflow, scikit-learn

Languages: Python, SQL, R, Java; Mandarin Chinese (professional working proficiency)

Tools: Git, Docker

Certifications: Databricks Certified Data Engineer Professional, Data Engineer Associate, Data Analyst Associate, Machine Learning Associate, and Generative AI Engineer Associate; AWS Certified Cloud Practitioner; Professional Scrum Master I (PSM I); Tableau Desktop Specialist