

Patrick Browne

Data Platform Leader & Engineer

719.338.4035 | 5515 W 63rd Ave, Arvada, CO 80003

patrickbrowne@outlook.com

pbrowne.net | linkedin.com/in/patrickmbrowne

PROFESSIONAL SUMMARY

Big data business leader, engineer, and architect with 15 years of experience. Interested in solving large-scale technical and strategic problems. Thrive in fast-paced environments with limited resources. Data-driven decision maker that constantly seeks out innovative and disruptive opportunities, candid feedback, and lateral thinking from colleagues. Emphasis on learning and mentoring in a highly collaborative environment.

SKILLS

- Enterprise Data Platform Architecture
- Product Roadmapping and Development
- Revenue Forecasting & Cost Optimization
- Data Org Management and Structuring
- AI/LLM/ML Development & Implementation
- Project Management & Agile Leadership
- Data Governance Systems
- Hybrid & Native Cloud Best Practices

WORK HISTORY

Sovrn Holdings Inc. - Boulder, CO

Engineering Manager (2024 - Present)

- Leading and developing multi team Data Engineering organization, spanning product, platform, and internal/external business end customers
- Driving cross-organization product initiatives for industry leading offerings
- Forecasting revenue and operational cost, including continual cost optimizations
- Technical leader of Data Governance and Compliance
- Implementing and roadmapping AI/LLM product development
- Key contact for internal and external stakeholders

Interim Manager & Staff Engineer (2022 - 2024)

- Owner of multiple observability platforms for AI/LLM and Big Data/Data Analytics workflows
- Architect and admin of AWS zero-trust multi-tenant organization
- Designer of streaming backend handling: 10GB/s, 150+ billion records daily
- Owner of Internal and external batch processing ETLs via Spark
- Mentor and project lead responsible for documentation, process improvement, quality standards, and on time delivery
- Lead support for front, backend, and middle-ware cloud migration from on-prem to AWS

Senior Software Engineer (2020 - 2022)

- Design lead, architect, and mentor for serverless, containerized, and big data storage ingress/egress projects
- Champion and Mentor for best practices to junior team members

Software Engineer II (2018 - 2020)

- Key contributor in AWS migration and creation of Data Science processes
- Cloud and On-prem admin for streaming and big data infrastructure

Senior Software Developer in Test (2016 - 2018)

- Test automation lead for Data Platform team of 9 developers
- Designed and maintained ephemeral system and integration test suites

Dish Networks - Denver, CO

QA Test Engineer II (2015 - 2016)

- Lead Quality Assurance for Data Architecture team
- Designed and maintained front-end Java-based test automation framework

Drobo/Connected Data - Santa Clara, CA

QA Test Engineer (2012 - 2015)

- Lead tester for the launch of the Transporter private cloud storage device
- Lead tester for Drobo storage devices, both enterprise and commercial

TECHNICAL SKILLS

Frameworks

- Spark, Spark Structured Streaming, Flink, Lambda/Serverless, Glue, Kubernetes, Kafka Connect, Step Functions, API Gateway/ALB, REST, FastAPI, ORM, OPA, Ninja2, Structured and Unstructured data

Databases/Storage

- S3, Databricks, GCS, Snowflake, EMR, Aerospike, Cassandra, Pinot/Startree, MySQL, Postgres, DynamoDB, RDS, Redis, Hadoop, MapR, Hive, Looker, Hue, Zeppelin, EBS, on-premise Linux

Languages

- Python, Go, Bash, Java, Scala, Rust

Infrastructure

- Docker, AWS Networking (VPC, IG, TGW, R53, ALB, etc.), EKS/Kubernetes, Glue Data Catalog, EC2, Airflow, Kafka, Kinesis, SNS, SQS, EventBridge, DataDog, Graphite/Grafana, VictorOps, Jenkins, GoCD, Vagrant, VirtualDock, Terraform, Terramate, Ansible, Helm, Cloud Formation

AI/LLM

- RAG analysis, LLM Data Categorization, Lookalike Modeling, Traffic Shaping, Agentic workflows, Ollama, Langchain, PyTorch, TensorFlow

RELEVANT PROJECTS

Leadership

- Organizational restructuring and alignment for backend and data departments
- Multi-team and geo-diverse management experience
- Hands on influence on go-to-market strategy, technical implementation, and product roadmap
- Leader of a product-driven team for data wholesale offerings | 9+ offerings, 55+% gross profit margin, \$12M of revenue per year
- Leader of a Data Platform team(s) responsible for the ingestion, governance, and storage of 4PB+ of data per week
- Leader of Data Science, Data Analytics, and other “data-centric” teams
- Experience in contract negotiation and strategic forecasting based on metrics-driven industry research
- Evaluations and negotiations across products in AWS, Databricks, Snowflake, and others
- Point-of-contact for government and consumer data compliance

AI/LLM/ML

- Architect for AI LLM powered web content categorizer used to drive insights and revenue for audience and deal packages on a wholesale and ad exchange level
- Lookalike audience modeling crossing disparate datasets
- Improved the checkout process to ensure that customers purchase products easily

Project Management

- Large-scale migrations and refactor efforts that spanned the company with timely on-budget delivery
- Full-tech stack evolution of adtech backend, outputting to both internal and external stakeholders
- Migration of 60 node on-prem data lake and legacy data analytics architecture into AWS
- Ground up tech adoption implementation for both Snowflake and Databricks

Data Governance/Admin

- OPA and a single service layer to unify access to multiple products lines
- Canonical ID system creation and identity graph implementation
- 28 node, 42+ TB Cassandra cluster in AWS for customer facing dashboard
- Snowflake, Sagemaker/EMR, Databricks admin and architect

Cost Savings

- Upstream and downstream optimizations focused on infrastructure bill reduction. Including but not limited to application, storage, governance, and business process workflow optimizations - Recent separate savings efforts include \$400K and \$570k in yearly cost reductions
- Key contributor in quarterly spend and revenue planning
- Tracking and forecasting cloud costs across AWS/GCP organizations and supplementary services
- Granular automated metrics to expose cost-to-revenue ratios across product lines and platforms

Observability

- ETL and data quality and efficiency monitoring for peta-byte scale ingest and egress workflows
- Data Science centric job monitoring and analytics platform solutions
- Streaming checkpoint and lag analysis for real-time and batch process signaling

Streaming/ELT

- Modular Spark Streaming framework ingesting ~250 billion records per day across teams and products
- Modular Flink streaming application to Delta Lake: 9 GB/s, 200+ billion records daily
- Adtech audience products with Databricks, AWS, and GDPR/CCPA data compliance requirements
- Spark in Snowflake/Databricks/EMR/AWS Glue for critical financial data and external data products
- Serverless event-based ingest monitoring and auditing system for unifying cross-product data sets and workflows
- Migration data processing applications into k8s wired for auto-scaling in EKS
- Kafka Connect ingestion from bare metal geo-dispersed servers to S3

Infrastructure

- Architect and admin of AWS zero-trust multi-tenant organization and structure for Data Platform
- Kubernetes application and pipeline development for big data use cases
- On-prem and cloud server and services management

EDUCATION & CERTIFICATES

Creighton University - (2008 - 2012)

B.S. in Business Intelligence Analysis

AWS Solutions Architect - Professional Level

Databricks Lakehouse Architecture

HOBBIES

AI, Classic Car Restoration, Welding, Wine-making, Snowboarding, Outdoors, Piano, Electronics, Exercise

REFERENCES AVAILABLE UPON REQUEST

