

Patrick Browne

Data Leader & Architect

719.338.4035 | Colorado
patrickbrowne@outlook.com | pbrowne.net
[LinkedIn](#)

PROFESSIONAL SUMMARY

Seasoned big data leader, engineer, and architect with over 15 years of experience. Excels at solving large-scale technical and strategic problems and thrives in fast-paced environments with limited resources. A data-driven decision maker who constantly seeks out innovative and disruptive opportunities, values candid feedback, and encourages lateral thinking among colleagues. Emphasis on continuous learning and mentoring within 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

2016 - Current

Manager, Data Engineering (2024 - Present)

- Transforming development workflows by leveraging agentic AI, MCP, and LLM across multiple teams
- Leading a multi-team Data Engineering org, encompassing product, platform, and internal teams
- Driving cross-organizational product initiatives to develop industry leading offerings
- Forecasting revenue and operational cost, including continual cost optimizations
- Serving as a technical leader for Data Governance and Compliance
- Implementing and roadmapping AI/LLM product development
- Serving as the key contact for internal and external stakeholders

Interim Manager & Staff Engineer (2022 - 2024)

- Owned of multiple observability platforms for AI/LLM and Big Data/Data Analytics workflows
- Architected and administered an AWS zero-trust multi-tenant organization
- Designed a streaming backend capable of handling 20GB/s, processing over 250 billion records daily
- Managed internal and external batch processing ETLs using Spark
- Mentored and led projects, ensuring documentation, process improvement, quality standards, and on-time delivery
- Led the support for frontend, backend, and middle-ware cloud migration from on-prem to AWS

Senior Software Engineer (2020 - 2022)

- Architect, and mentor for serverless, containerized, and big data storage, ingress, and egress
- Championed and mentored junior team members in best practices

Software Engineer II (2018 - 2020)

- Key contributor in AWS migration and development of Data Science processes
- Administered cloud and on-prem infrastructure for streaming and big data

Senior Software Developer in Test (2016 - 2018)

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

Dish Networks - Denver, CO

2015 - 2016

QA Test Engineer II

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

Drobo/Connected Data - Santa Clara, CA

2012 - 2015

QA Test Engineer

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

KEY PROJECTS

Leadership

- Organizational restructuring and alignment for backend and data departments
- Extensive experience in managing multi-team and geographically diverse groups
- Hands on influence on go-to-market strategy, technical implementation, and product roadmap
- Led a product-driven team for data wholesale offerings, managing 9+ offerings with a 55+% gross profit margin and generating \$12M in annual revenue
- Led Data Platform teams ingestion, governing, and storing over 4PB of data per week
- Led Data Science, Data Analytics, and other “data-centric” teams
- Experience in contract negotiation and strategic forecasting using metrics-driven industry research
- Evaluated ROI across products in AWS, Databricks, Snowflake, and other platforms
- Served as the point of contact for government and consumer data compliance

AI/LLM/ML

- Architect for an AI LLM powered web content categorizer for audience and deal packaging in adtech
- Developed a lookalike audience modeling crossing disparate datasets
- Structured and retrieved canonical ID data with entity resolution

Project Management

- Led large-scale migrations and refactor efforts across the company, delivering projects on time and within budget
- Evolved the full-tech stack of the adtech backend, serving both internal and external stakeholders
- Migrated of 60 node on-prem data lake and legacy data analytics architecture into AWS
- Implemented ground up tech adoption for both Snowflake and Databricks

Data Governance/Admin

- Implemented OPA and a single service layer to unify access to multiple products lines
- Created a canonical ID system creation and implemented an identity graph
- Managed a 28 node, 42+ TB Cassandra cluster in AWS for customer-facing dashboard
- Admin of Pinot Startree cluster and ETL serving frontend dashboards
- Served as an admin and architect for Snowflake, Sagemaker/EMR, Databricks

Cost Savings

- Focused on upstream and downstream optimizations to reduce infrastructure costs, including application, storage, governance, and business process workflow optimizations. Recent cost savings efforts have resulted in annual reduction of \$400K and \$570k
- Served as a key contributor in quarterly spend and revenue planning
- Tracked and forecasted cloud costs across AWS and GCP organizations and supplementary services
- Implemented 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
- Developed data science centric job monitoring and analytics platform solutions
- Conducted streaming checkpoint and lag analysis for real-time and batch process signaling

Streaming/ELT

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

Infrastructure

- Architect of AWS zero-trust multi-tenant organization and structure for Data Platform
- Developed Kubernetes cluster with applications and pipelines for big data use cases
- Created and managed on prem and cloud servers, applications, and databases with infra-as-code

TECHNICAL SKILLS

Frameworks

- AI Agentic MCP, Spark, Spark Structured Streaming, Flink, Lambda/Serverless, Glue, Kubernetes, Kafka Connect, Event-driven, APIG/ALB, REST, FastAPI, ORM, OPA, Jinja2, 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

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

EDUCATION & CERTIFICATES

Creighton University | 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
