Brent Burgess

Data leader who ships AI automation.

Benicia, California · brent@brentburgess.com · LinkedIn (https://www.linkedin.com/in/brent-burgess-126984a/) · GitHub (https://github.com/bburgess42) · brentburgess.com (https://brentburgess.com/)

Summary

Data and AI leader with 21 years in data, 18 of them in enterprise data architecture and 7 leading data engineering, machine learning and analytics teams.

At Wine.com I grew from Lead Data Architect to Director of Data and Analytics and built the data foundation that planning, finance and operations ran on, then the forecasting, anomaly detection and production AI agents on top of it.

Selected results

- A governed lakehouse on AWS and Databricks at 99.9% uptime across five years, replacing an on-premise SQL Server estate domain by domain.
- One certified metric layer, with an owner, grain and lineage on every definition, that finance, operations, product and marketing all reported from.
- Production AI agents that answered from certified metric definitions through MCP servers behind SSO and OAuth, and a Claude Code suite that engineers and analysts used daily.
- A team of 7 across data engineering, ML and analytics, and data standards adopted across more than 8 engineering teams.

Selected systems

- Sleep Anchor: an AI and Python publishing pipeline (https://brentburgess.com/work/sleep-anchor): An eight-hour render fell from about 4.8 hours to about 25 minutes.
- An MCP server for this site (https://brentburgess.com/work/mcp-server): AI agents can search and read this site’s writing, projects and published numbers, read-only.
- This site’s own data (https://brentburgess.com/work/nightly-data-pipeline): Numbers go live only after their tests pass, and 157 of 157 passed on data through 3 November 2026.

Experience

Wine.com

Director of Data and Analytics (from 2023) · Senior Data Engineering Manager (from 2022) · Data Engineering Manager (from 2018) · Lead Data Architect (from 2017)
San Francisco, California, October 2017 to March 2026

- Set the technical direction and roadmap as product owner of the data platform: priorities, risk, the platform budget, build-versus-buy, vendor contracts and the investment cases I defended with the executive team.
- Replaced an on-premise SQL Server, SSAS and SSIS estate, domain by domain, with a medallion lakehouse on AWS and Databricks (Delta Lake, Unity Catalog, Spark, Redshift, S3): change data capture from more than 20 source systems, Airflow, Terraform, CI/CD, parity checks before every cutover, and 99.9% uptime across five years.
- Ran a data products program with named domains, owners, published interfaces and contracts for schema, freshness and quality, with self-healing checks that quarantined bad records and backfilled them, and designed one certified metric layer, with an owner, grain and lineage on every definition, that finance, operations, product and marketing reported from.
- Led forecasting and predictive modeling to production: demand and volume forecasting, propensity and churn models, a feature store, model serving, drift monitoring and retraining, and anomaly detection with KPI alerting on statistical process control limits.
- Built the supply chain and operations data behind planning (delivery, fulfillment, inventory, purchasing, spend, an ERP migration), the leadership scorecards and the financial reporting layer for FP&A, and retired Power BI onto governed Tableau and Sigma.
- Shipped production AI agents for internal tooling, a business workflow and a customer-facing feature, with tool calling, multi-agent orchestration and MCP servers that give agents governed access to data, and set the standards they run

under: SSO and OAuth, least-privilege access, PII rules, monitoring, cost discipline and human-in-the-loop review, including in a Claude Code development suite.
- Hired, coached and developed a distributed team of 7, growing engineers into lead and manager roles, while writing production SQL, PySpark and Terraform and setting data standards adopted across more than 8 engineering teams.

UC Berkeley Graduate School of Education

Lead Data Architect
Berkeley, California, March 2017 to September 2017

- Designed and delivered the data warehouse for the DRDP statewide pre-kindergarten to kindergarten assessment platform used across California public schools, alongside psychometric researchers and under public-sector privacy and security constraints.

Bluenet Technologies

Lead Data Architect
Tampa, Florida, February 2013 to March 2017

- Customer-facing lead architect in a consultancy across 5 client organizations and more than 100 databases, advising client CTOs, CIOs and CDOs on data strategy, architecture and investment trade-offs, and delivering executive reporting and Cognos analytics for healthcare, insurance and industrial clients.

Quality Distribution

Data Architect · Lead Data Engineer
Tampa, Florida, October 2008 to February 2013

- Designed analytical databases and BI for logistics and fleet operations at a NASDAQ-listed bulk logistics company: dispatch, fleet utilization, terminal and safety reporting, truck telematics at volume, financial reporting under SOX and ITGC controls, and benefits claims reporting with automated de-identification of PHI.

Lifestyle Family Fitness

Senior DBA · Data Engineer
St. Petersburg, Florida, July 2005 to October 2008

- Administered and optimized SQL Server environments and built the first data integration and reporting infrastructure for a multi-site membership business: member activity, retention and billing reporting.

Skills

Data architecture
Data architecture, Data engineering, Medallion lakehouse, Databricks on AWS (Delta Lake, Unity Catalog, Spark, Feature Store, MLflow), Redshift, S3, EMR, Lambda, Batch and near-real-time CDC ingestion, Apache Airflow, Terraform, CI/CD

AI automation
AI automation, Agent development with tool calling, Model Context Protocol, Multi-agent orchestration, Claude Code, LLM features in production, SSO and OAuth for AI tools, LLM observability and cost management, Human-in-the-loop governance

Decision intelligence and ML
Demand and volume forecasting, Time series, Propensity and churn models, Anomaly detection, Statistical process control, Feature pipelines, model serving, drift monitoring and retraining

Governance and metadata
Unity Catalog (lineage, classification, RBAC), Certified semantic and metric layers, Data contracts, Self-healing quality checks, CCPA and GDPR, PCI DSS, SOX and ITGC, HIPAA and PHI

Analytics and BI
Analytics and BI, Tableau, Sigma, Power BI, Cognos, Executive scorecards

Languages and tools
SQL (Spark SQL, T-SQL, PostgreSQL, MySQL), Python, PySpark, dbt, DuckDB

Leadership
Data and AI strategy, Roadmap and backlog ownership, Platform budget and investment cases, Vendor evaluation and contract negotiation, Hiring, coaching and performance management

Education

- B.A., History, University of South Florida, Tampa, Florida
- Certifications in progress: Databricks Data Engineer Professional; Databricks Generative AI Engineer Associate