

Chase Burkhalter

Senior Analytics Engineer | Data Engineer | AI-Ready Data Infrastructure

chase@chaseburkhalter.com | (334) 333-4308 | Dothan, AL (Open to Remote) | linkedin.com/in/chaseburkhalter | chaseburkhalter.com | github.com/ctburkhalter

PROFESSIONAL SUMMARY

Senior analytics engineer with 6+ years designing, building, and operating the analytics layer of modern data stacks, with deep expertise in dbt, Snowflake, and Fivetran from raw ingestion through governed, production-grade marts. Recent work concentrates where demand is highest: cutting Snowflake compute costs ~80% on targeted pipelines through incremental re-architecture, making warehouses AI-ready with persisted documentation, custom MCP servers, and semantic context layers, and automating recurring data workflows with Python. Pairs software-engineering rigor (version control, CI/CD, testing, documentation) with close stakeholder partnership and a consistent record of measurable outcomes.

CORE SKILLS & TOOLS

Analytics Engineering	dbt (Core & Cloud), Snowflake, Fivetran, Hightouch, Airflow, MetricFlow, Semantic Layer Design, Incremental Modeling
AI & LLM Tooling	MCP Server Development (FastMCP), Claude & Claude Code, Hex AI (Context Studio), LookML-to-Semantic-Context Generation, Prompt Engineering, AI-Ready Data Design
Programming	SQL, Python, JavaScript / TypeScript, Jinja, LookML, React, Next.js
Cloud & Data Infrastructure	AWS (Lambda, S3, Redshift, Glue), BigQuery, Postgres, Docker
BI & Analytics	Looker / LookML, Amplitude, Hex, Tableau, Power BI, Apache Superset, Chartbeat, Mixpanel
Instrumentation & Product Analytics	Google Tag Manager, GA4, Segment, AppsFlyer, Airship, Avo, DataHub, OpenMetadata, Pendo
Marketing & Data Activation	Iterable, Braze, Stripe, Salesforce, HubSpot, LaunchDarkly, Optimizely
Practices	Data Quality & Governance, Data Lineage, CI/CD, Warehouse Cost Optimization (FinOps), A/B Testing, Cohort Analysis, Dimensional Modeling

PROFESSIONAL EXPERIENCE

Future Research, Inc. | Remote

February 2026 - May 2026

Senior Analytics Engineer (Contract)

- Sole analytics-engineering contractor on the core Analytics team of a consumer fitness platform; shipped ~30 dbt PRs and closed 84 tickets across modeling, infrastructure, AI enablement, and incident response in four months.
- **Cut Snowflake compute costs ~80%** across the 11 highest-cost dbt models (682 to 131 credits/month, roughly \$12K annualized) by converting full-refresh models to incremental builds and archiving dead pipelines; savings were verified in ACCOUNT_USAGE and sustained while schema query volume grew 58%.
 - **Built the company's AI-ready data foundation:** enabled persist_docs to write governed model and column documentation into Snowflake as queryable metadata, authored dedicated documentation for 18 core mart models, and generated a LookML-derived semantic context layer covering 539 measures across 58 views for AI tools and BI consumers.
 - **Developed custom MCP servers (FastMCP)** for Snowflake, Hightouch, and Linear, plus a MetricFlow generator producing 53 semantic models, tooling that let AI-assisted workflows query, validate, and operate production data systems directly.
 - **Launched natural-language data access** by designing the Snowflake-to-Claude integration: a scoped, least-privilege role architecture (read-only, mart-layer only) that enabled 13 non-technical teammates to query governed production data without writing SQL.
 - **Owned end-to-end source-to-warehouse integrations** for Intercom, Amplitude, Iterable, and Google Search Console, handling Fivetran ingestion, identity resolution, and downstream modeling from raw source to production-ready marts.
 - **Led incident response** for a 13-day, 2.4-billion-row Aurora PostgreSQL replication failure as primary vendor contact, restoring business-critical pipelines with verified source-to-warehouse parity; separately recovered a \$1,000 vendor credit by escalating a Fivetran billing dispute.

- **Automated recurring manual data workflows**, replacing a stale, agency-maintained affiliate-spend spreadsheet with a scheduled API-driven pipeline and shipping source-freshness monitoring automation, which eliminated a chronic source of stale data and false alerts.

The Atlanta Journal-Constitution | Atlanta, GA (Hybrid)

August 2024 - June 2025

Senior Data & Analytics Engineer

- **Promoted from Senior BI Analyst to Senior Data & Analytics Engineer** in February 2025, in recognition of expanding ownership of the company's data platform and engineering practices.
- **Owned the end-to-end rollout of Amplitude analytics across 7 digital properties**, independently handling JavaScript instrumentation, taxonomy design, QA, and ongoing governance to establish the company's first unified behavioral analytics layer.
- **Migrated core event tracking and reporting from Redshift to Snowflake**, improving query performance for editorial and executive dashboards and building a more scalable, cost-efficient foundation for downstream analytics.
- **Rebuilt the GTM/GA4 tracking infrastructure from the ground up** after a team restructuring, assuming sole ownership of tagging, debugging, and data QA across all web properties with no handoff documentation to work from.
- **Integrated Airship, AppsFlyer, and Iterable into the analytics stack**, partnering with product and platform teams to unlock cross-platform engagement reporting that hadn't previously existed.
- **Laid the groundwork for enterprise dbt adoption and data governance** through documentation, taxonomy standardization, and data-lineage mapping.
- **Built Power BI and Tableau dashboards for election and campaign coverage**, giving editorial and executive leadership faster, more reliable performance reporting during high-stakes news cycles.

Shortcut | New York, NY (Remote)

October 2022 - May 2024

Senior Data Analyst

- **Served as the sole analyst supporting 6 cross-functional departments**, including sales, finance, marketing, product, engineering, and customer support, owning the full analytics workflow from requirements through delivery.
- **Drove \$100K+ in annual savings** by independently evaluating and executing a full migration off proprietary visualization tools to Apache Superset, with zero loss of reporting coverage.
- **Administered the company's entire modern data stack solo**, including Segment, Amplitude, Superset, Google Analytics, Snowflake, dbt, Fivetran, Hightouch, and Pendo.
- **Built and maintained reporting models in SQL and dbt** from first- and third-party data sources, becoming the primary source of trusted insight for finance, sales, marketing, and product stakeholders.
- **Standardized event-tracking processes across teams**, establishing shared definitions that enabled non-technical stakeholders to self-serve customer behavior analysis.

SteadyApp / SteadyIQ | Atlanta, GA (Remote)

March 2020 - September 2022

Product & Data Analyst

- **Saved \$1.4M annually** by analyzing roughly 3 billion transactions and associated fraud indicators to redesign the company's data-aggregator infrastructure, minimizing data loss while cutting processing cost.
- **Promoted from Product Analyst to Data Analyst** in January 2022 in recognition of expanding technical ownership of the analytics stack.
- **Managed the company's core analytics stack**, including Segment, Amplitude, Tableau, Braze, and Extol, serving as the primary technical resource driving data-informed decisions across product, growth, and executive teams.
- **Built and maintained source-of-truth reporting models in Snowflake via dbt**, replacing inconsistent ad hoc reporting with scalable, reliable models that executive stakeholders trusted without re-verification.
- **Designed data governance plans and tracking strategies** for new feature launches, UX optimizations, and A/B testing initiatives, establishing measurement frameworks before features shipped rather than after.

EDUCATION

The University of South Alabama | Mobile, AL | May 2017

B.S. Biology / Biochemistry