

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, MCP tooling, 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	Claude & Claude Code, ChatGPT Codex, Hex AI (Context Studio), MCP Tooling & Integration (FastMCP), 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 through peer code review and closed 84 tickets across data modeling, architecture, infrastructure, product analytics, AI enablement, and incident response in four months.

- **Owned source-to-warehouse ingestion across the company’s full first- and third-party data estate**, spanning production Aurora PostgreSQL app data (users, workouts, billing), Stripe finance data, paid media and marketing platforms (Meta, Google Ads, TikTok), and product and support tools (Amplitude, Intercom, Iterable, Shopify, Google Search Console), handling ingestion, identity resolution, and downstream modeling from raw source to production-ready marts.
- **Modeled the warehouse through a layered staging-to-mart dbt architecture (Medallion-style bronze/silver/gold)**, setting the engineering standard behind it: enforced dbt 1.9+ conventions and generic-test-argument syntax project-wide to close a class of silent CI failures, then led the staged deprecation of a retired product line, archiving 22 models across two coordinated phases with full documentation preserved for reactivation.
- **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, generated a LookML-derived semantic context layer covering 539 measures across 58 views, and helped stand up MCP tooling across Snowflake, dbt, Looker, Hex, Hightouch, and Linear, including a MetricFlow generator producing 53 semantic models, so AI tools and workflows could query, validate, and operate real 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 non-technical teammates to query governed production data without writing SQL.
- **Led the Amplitude product-analytics revamp across both production projects (mobile app and web)**, rebuilding the event and user-property taxonomy and partnering with engineering to ship a canonical-ID fix that ended cross-device double-counting.
- **Led incident response as primary vendor contact** through a 13-day, 2.4-billion-row Aurora PostgreSQL replication failure, restoring pipelines with verified parity, and automated a stale affiliate-spend reporting process.

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**, unlocking cross-platform engagement reporting that hadn't previously existed, and laid the groundwork for enterprise dbt adoption 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