Company Marketing Brief

PhoenixAI

Generated 10.25.2026

Mktforge

What is this report?

This marketing brief for PhoenixAI is intended to serve as a high-level overview of PhoenixAI's ideal customers, their careabouts, how PhoenixAI can reach these customers, craft messaging for them, and position against known competitors. It is meant to serve as a starting point for focusing PhoenixAI's marketing discussions on the right frameworks and questions to quickly develop more refined and effective go-to-market strategies.

While the information here is well-sourced and guided by decades of marketing experience, it should first and foremost be seen as an early draft of the final guiding marketing materials PhoenixAI will create. Most of the material found within this document are highlights of more extensive documents created with Mktforge. If you wish to delve deeper, ask the creator of this document for the source files.

Inside This Report

Customer Overview **4**

Competitive Analysis **9**

Positioning and Messaging **12**

Customer Overview

INTRODUCTION

Who are the best people to target to drive adoption of PhoenixAI? While the diversity of buyers is as myriad as the uses for PhoenixAI's products and services, some buyers needs are better aligned with the features and value PhoenixAI offers, and so should be considered for prioritization. These include: [1]

Top Targets

Senior Engineering Manager, Data Platform

- Legacy ClickHouse/warehouse architecture can't handle unpredictable, LLM-generated SQL from AI agents
- Need isolated compute for agent query bursts so they don't contend with BI dashboards on the same data
- Query optimizer must handle novel, unplanned joins automatically without manual tuning
- Must maintain second-level SLAs for customer-facing dashboards while agents hammer the system
- Need to query petabytes of normalized data across thousands of tenants efficiently

Data Engineer

- Designing and maintaining ETL/streaming pipelines (Kafka, Flink, CDC) for near-real-time data
- Balancing cost, scalability, and complexity of warehouse/lakehouse architecture
- Ensuring data freshness at the seconds level without proliferating pipelines
- Supporting both BI dashboards and AI/ML workloads off the same underlying data
- Maintaining data governance, access control, and audit at scale

Senior Staff Software Engineer

- Complex joins across hundreds of normalized tables at massive scale (hundreds of billions of rows)
- Fraud/compliance workloads require low latency on fresh streaming data
- Need streaming updates from Kafka to be queryable within seconds
- Analysts and AI agents both need to operate on the same real-time dataset
- Legacy query engines can't provide the low-latency performance required at scale

Data Platform Engineer

- Owning end-to-end architecture and operational health of OLAP engines at scale
- Optimizing query engine performance (cost-based optimization, vectorized execution, memory)
- Maintaining sub-second latency under heavy concurrent load
- Managing streaming ingestion pipelines and cluster scalability
- Reducing infrastructure footprint and cost while sustaining performance

(Continued on next page)

[1] Source: Mktforge Find My Customer — Results for https://www.phoenixdata.ai, generated 10/25/2026

Chief Product Officer

- Delivering customer-facing analytics at terabyte scale under high concurrency
- Legacy analytics layer outgrown by product/customer growth
- Need production queries over hundreds of millions of rows to return sub-second
- Need an architecture ready for next-generation agentic workloads
- Shrinking the time from idea to production for new data-driven product features

Senior Software Engineer, Real-Time Analytics

- Building and maintaining low-latency query engines and pipelines for streaming data
- Balancing data freshness against query performance at scale
- Supporting complex multi-table joins without excessive precomputation or pre-aggregation
- Scaling systems to sustain high concurrency/QPS without latency degradation
- Debugging and optimizing distributed query execution (vectorization, parallelism, memory)

VP of Engineering / VP of Data Platforms

- Setting multi-year technology/data platform strategy and roadmap aligned with business goals
- Owning budget, headcount, and total-cost-of-ownership decisions for platform infrastructure
- Ensuring platform reliability, SLAs, and operational excellence across engineering teams
- Evaluating build-vs-buy and vendor/tool consolidation to reduce cost and architectural complexity
- Enabling new AI/agentic workloads and initiatives without destabilizing existing systems

Analytics Engineer

- Managing increasingly complex data pipelines across multiple sources and platforms
- Balancing quick fixes against long-term technical debt in transformation code
- Ensuring data quality, consistency, and testing across diverse data sources
- Translating business requirements into data models while communicating with stakeholders
- Optimizing pipeline/query performance and cost as data volume grows

Head of Data Platform

- Owning end-to-end architecture and strategic direction for the company's data platform
- Balancing cost, scalability, and reliability of data infrastructure investments
- Supporting diverse workloads (BI, ML, agentic AI) on shared infrastructure without contention
- Driving platform modernization and reducing tool/cluster sprawl
- Ensuring governance, security, and compliance across the data platform

Chief Technology Officer

- Setting overall technology vision and strategy aligned to business outcomes
- Evaluating and adopting AI/agentic technologies to maintain competitive advantage
- Managing technology risk, security, and compliance at the executive level
- Overseeing infrastructure costs and ROI across the technology stack
- Balancing innovation with legacy system stability as the company scales

(Continued on next page)

[1] Source: Mktforge Find My Customer — Results for https://www.phoenixdata.ai, generated 10/25/2026

WHERE IS BUYER X/Y/Z?

The buyers outlined above are simply the ones that best align with PhoenixAI's current messaging and feature offerings, and there are always exceptions. This list is not to say other titles aren't worth selling into or that PhoenixAI wouldn't succeed targeting other titles, only that the titles above should receive additional consideration.

TOP BUYERS - DEEPER INSIGHTS

While any of the roles outlined above could be worth targeting (and should still be considered regardless), Senior Engineering Manager, Data Platform and Data Engineer should receive further consideration for targeted outreach and messaging refinement.

These buyers are not only worth targeting because their needs align with PhoenixAI's general value proposition, but because PhoenixAI's specific features and services map best to their needs outlined below: [2][3]

Where Senior Engineering Manager, Data Platform's Needs Align with PhoenixAI's Offerings

Needs	Where PhoenixAI Wins
Reduce platform complexity by consolidating clusters, eliminating redundant ETL, and enabling direct queries on normalized data	• Cluster consolidation & cost reduction • Distributed MPP joins on normalized data • Cost-based optimizer for AI-generated…
Enable customer-facing analytics and embedded dashboards with sub-second latency under heavy concurrent load	• Sub-second freshness with streaming… • Isolated warehouses for workload…

Proof Points to Consider

Cluster consolidation & cost reduction

Demandbase consolidated 49 ClickHouse clusters to 1 PhoenixAI cluster with 60% hardware reduction and 90% storage cost savings, eliminating sprawl and operational overhead.

Distributed MPP joins on normalized data

Runs complex multi-table joins on normalized schemas with sub-second latency. Coinbase queries 573B+ rows across 300+ tables at <1s—no pre-flattening or denormalization required.

Sub-second freshness with streaming…

Ingests Kafka and Flink CDC at 100K+ events/sec with updates queryable in under 5 seconds. Intuit achieves 2s total latency for real-time ML predictions without the FINAL tax.

Cost-based optimizer for AI-generated…

Automatically handles novel joins and LLM-generated SQL without manual tuning. Demandbase replaced ClickHouse architecture that couldn't handle unpredictable AI agent queries.

Isolated warehouses for workload…

Agent queries, BI dashboards, and ingestion each run in their own compute warehouse on the same data. Agent bursts don't contend with production dashboards.

[2] Source: Mktforge Battle Card Generator — Competitive Battle Card: ClickHouse, Target Job Title: Senior Engineering Manager, Data Platform, generated 10/25/2026
[3] Source: Mktforge Battle Card Generator — Competitive Battle Card: ClickHouse, Target Job Title: Data Engineer, generated 10/25/2026

Where Data Engineer's Needs Align with PhoenixAI's Offerings	
Needs	**Where PhoenixAI Wins**
Build and maintain scalable real-time data pipelines from streaming sources (Kafka, Flink CDC) with sub-second data freshness	• Streaming upserts queryable in seconds
Design lakehouse architectures that unify streaming and historical data (Iceberg, Delta) for analytics and ML workloads	• Unified SQL on streaming and lakehouse
Support AI agent and ML workloads with low-latency query infrastructure that handles unpredictable, complex SQL patterns	• Cost-based optimizer for AI-generated…
Optimize query performance and reduce infrastructure costs through right-sized compute, partitioning, and workload isolation	• Isolated warehouses for workload…
Ensure high concurrency and sub-second latency for customer-facing analytics and embedded dashboards	• Distributed MPP joins without denorm

Proof Points to Consider

Streaming upserts queryable in seconds

Primary key tables handle 100K+ events/sec with upserts and deletes queryable in under 5 seconds—no FINAL tax or latency degradation. Intuit achieves 2s total latency for real-time ML predictions.

Unified SQL on streaming and lakehouse

Query Iceberg, Delta, and Hudi in place alongside streaming tables in one SQL statement. 8.2× faster than Trino on Iceberg in TPC-DS 1TB benchmark. No data copy, no duplicate pipeline to maintain.

Cost-based optimizer for AI-generated…

Automatically handles novel joins and LLM-generated SQL without manual tuning. Demandbase replaced ClickHouse architecture that couldn't handle unpredictable AI agent queries; optimizer adapts at runtime.

Isolated warehouses for workload…

Agent, BI, and ingestion workloads each run in their own compute on the same data. Demandbase achieved 60% hardware reduction and 49'1 cluster consolidation after replacing ClickHouse sprawl.

Distributed MPP joins without denorm

Complex multi-table joins on normalized schemas with sub-second latency. Coinbase queries 573B rows across 300+ tables at <1s latency. No pre-flattening or pre-aggregation required.

(Continued on next page)

[3] Source: Mktforge Battle Card Generator — Competitive Battle Card: ClickHouse, Target Job Title: Data Engineer, generated 10/25/2026
[4] Source: Mktforge Marketing Opportunities — Results, Job Titles Provided: Data Engineer, generated 10/25/2026

WHERE TO FIND THESE BUYERS

Many various channels and activities exist to reach these buyers, and some are likely already known. Below is a collection of some of the best matches based on the initiatives of Senior Engineering Manager, Data Platform's and Data Engineer's and where they'd be most receptive to hearing about PhoenixAI's offerings: [4]

Name	Channel	Link
Current 2026 (Confluent)	Events and Tradeshows	https://current.confluent.io/
Iceberg Summit 2026	Events and Tradeshows	https://www.icebergsummit.org/
Apache Iceberg Meetup ATL	Smaller Group Events	https://www.meetup.com/apache-iceberg-meetup-atl/
NYC Apache Iceberg Community Meetup	Smaller Group Events	https://luma.com/apache-iceberg
Alex Merced	Influencers	https://iceberglakehouse.com/
Data Engineer Things	Newsletters	https://dataengineerthings.substack.com/

[4] Source: Mktforge Marketing Opportunities — Results, Job Titles Provided: Data Engineer, generated 10/25/2026

Competitive Analysis

INTRODUCTION

Determining PhoenixAI's competitors requires a combination of looking at its target buyers, current offerings, and value proposition and the names and products that come up in common searches, job descriptions, and online conversations related to PhoenixAI's target buyers. Based on the information presented above, the following competitors appear to surface the most: [1][5]

Competitor	URL	Where They Compete
Apache Doris (VeloDB) PRIORITY	velodb.io	Real-time analytical database for agent-facing and user-facing analytics with sub-second queries, high concurrency, streaming ingestion, vectorized execution, and open lakehouse support (Iceberg/Delta/Hudi); directly competes on agent workloads, fresh data within seconds, complex joins, and sells to data platform engineers and CTOs with nearly identical pain points.
Apache Druid	druid.apache.org	Real-time OLAP database for streaming analytics with Kafka ingestion, sub-second queries, high concurrency, and pre-aggregation for fast rollups, competing on real-time dashboards and high-QPS workloads to data platform engineers, but requires pre-aggregation where PhoenixAI handles arbitrary multi-table joins without denormalization.
Apache Pinot	pinot.apache.org	Distributed OLAP for user-facing and agent-facing real-time analytics with sub-second latency, petabyte-scale streaming from Kafka, high concurrency (100K+ QPS), and upserts—directly competes on agent workloads, real-time streaming ingestion, isolated workloads, and sells to identical data platform engineering titles.
ClickHouse PRIORITY	clickhouse.com	Columnar OLAP database for real-time analytics and AI agents with Kafka streaming ingestion, petabyte-scale high-concurrency queries, and sub-second latency; directly competes on agent-facing workloads, real-time streaming, and complex joins to same data platform engineers—PhoenixAI's messaging explicitly positions against legacy ClickHouse architectures.
Firebolt PRIORITY	firebolt.io	High-performance analytical database for AI agents and real-time analytics with sub-second queries on Iceberg tables, isolated workloads, Postgres SQL dialect, and high concurrency; directly competes on agent workloads, lakehouse querying, and complex joins, selling to senior engineers and data platform leaders with overlapping messaging.

(Table continues on next page)

[1] Source: Mktforge Find My Customer — Results for https://www.phoenixdata.ai, generated 10/25/2026
[5] Source: Mktforge My Company module — tracked Competitor URLs

(Table continued from previous page)

Competitor	URL	Where They Compete
SingleStore PRIORITY	singlestore.com	In-memory HTAP database for real-time analytics and AI features with high-throughput streaming ingestion and fast queries, selling to data platform engineers and CTOs, but emphasizes transactional+analytical hybrid workloads whereas PhoenixAI is pure analytical for agent and customer-facing use cases.
StarRocks	starrocks.io	Real-time analytical database with vectorized MPP execution, sub-second queries on multi-table joins, streaming ingestion, native Iceberg/Delta/Hudi support, isolated warehouses, and agent-focused messaging—PhoenixAI is built on StarRocks open-source foundation and competes as a commercial managed alternative to the same core engine with identical technical capabilities.
Tinybird PRIORITY	tinybird.co	Managed ClickHouse for real-time APIs and streaming analytics with Kafka ingestion, sub-second SQL queries, and developer-friendly experience, selling to data platform and backend engineers, but packaged as API-first developer platform rather than PhoenixAI's general-purpose analytical database for agents and analytics.

WHERE PHOENIXAI COMPETES

A case can be made for any competitor under the right circumstances, however, when comparing the competitors above to the needs of Senior Engineering Manager, Data Platform's and Data Engineer's, there are particular areas where PhoenixAI stands out. These areas where PhoenixAI competes are highlighted below: [2][6][3]

(Continued on next page)

[2] Source: Mktforge Battle Card Generator — Competitive Battle Card: ClickHouse, Target Job Title: Senior Engineering Manager, Data Platform, generated 10/25/2026
[3] Source: Mktforge Battle Card Generator — Competitive Battle Card: ClickHouse, Target Job Title: Data Engineer, generated 10/25/2026
[6] Source: Mktforge Battle Card Generator — Competitive Battle Card: ClickHouse, Target Job Title: Senior Engineering Manager, Data Platform, generated 10/25/2026

ClickHouse

PRIORITY

Competitor Claim	Where PhoenixAI Competes
Columnar storage & vectorized execution	PhoenixAI also uses vectorized execution but adds distributed MPP joins. Coinbase benchmark: ClickHouse failed 12 of 22 TPC-H queries with OOM errors; PhoenixAI completed all.
Real-time streaming ingestion from Kafka	PhoenixAI also ingests Kafka at millions of events/sec, but commits upserts in place with sub-10-second freshness. ClickHouse merge-on-read forces stale rows or FINAL tax. Auto-scaling managed operations
Vectorized columnar query execution	PhoenixAI also uses vectorized execution, but adds distributed MPP joins. Coinbase's TPC-H 1TB test: ClickHouse failed 12 of 22 queries with OOM errors; PhoenixAI completed all with better P50–P99 latency.
Native Kafka streaming ingestion	PhoenixAI also ingests Kafka at 100K+ events/sec, but commits upserts in place with sub-10-second freshness. ClickHouse's merge-on-read forces a choice: stale rows or the FINAL tax that degrades query speed.
Iceberg and Delta Lake support	PhoenixAI also queries lakehouse formats in place, but unifies them with streaming tables in one SQL statement. PhoenixAI delivers 8.2× the performance of Trino and 1.3× Databricks on Iceberg in TPC-DS 1TB.

Positioning and Messaging

INTRODUCTION

Considering the proposed buyers (Senior Engineering Manager, Data Platform / Data Engineer) and priority competitors outlined in the previous section, several market categories present themselves for how PhoenixAI could be positioned in the market. [7]

Real-Time Analytical Database (Real-Time OLAP) — Existing category

RECOMMENDED

> **Positioning statement:** For data engineers who need to serve AI agents, customer-facing analytics, and BI dashboards with sub-second queries on fresh streaming and lakehouse data at scale, PhoenixAI is a real-time analytical database that runs distributed MPP joins on normalized data, commits streaming upserts queryable in under 5 seconds without latency degradation, and isolates agent, BI, and ETL workloads on the same data with independent scaling. Unlike ClickHouse, it handles complex multi-table joins and unpredictable AI-generated SQL without denormalization workarounds, the FINAL tax, or cluster sprawl.

Helps: Buyers already search for 'real-time OLAP' or 'real-time analytical database' and understand the use cases (streaming analytics, customer-facing dashboards, high concurrency). PhoenixAI's differentiators (distributed joins, no FINAL tax, unified lakehouse SQL, workload isolation) are immediately relevant and solve known ClickHouse/Pinot/Druid pain points. Shortest path to evaluation.

Hurts: Lumped in with ClickHouse open-source, which has huge brand recognition and a massive community. Data engineers may assume PhoenixAI is 'yet another ClickHouse alternative' and not investigate further. Also competes with Apache Pinot, Druid, and StarRocks open-source (PhoenixAI's own foundation), which are free. Risk of being seen as expensive managed version of open-source without clear different

Agentic AI Database / AI Agent Data Platform — Subcategory

ALTERNATIVE

Helps: Aligns with PhoenixAI's homepage positioning ('The real-time analytical database for AI agents'), $80M Series B messaging ('defining the category of agentic AI databases'), and the Demandbase/Conductor proof points where ClickHouse couldn't handle LLM-generated SQL. Differentiates from general-purpose ClickHouse by making cost-based optimizer, workload isolation, and unpredictable query handling t

Hurts: Buyers searching for 'real-time analytics' or 'streaming database' won't find this subcategory unless PhoenixAI educates them first. Requires marketing investment to explain what an 'agentic AI database' is and why it's different from real-time OLAP. Narrows the addressable market to only teams building AI agents, excluding the larger real-time BI, customer-facing analytics, and operational analyt

(Continued on next page)

[7] Source: Mktforge Build Positioning (Stages 0–5) — Data Engineer vs. clickhouse.com, generated 10/25/2026

Unified Real-Time Data Platform / Live Data Foundation — New category

ALTERNATIVE

Helps: Positions PhoenixAI as the architecture that eliminates the need for separate real-time OLAP (ClickHouse/Pinot), lakehouse query engines (Trino/Databricks), and warehouse sprawl—unifying streaming, historical, and agent workloads in one platform. The Demandbase proof point (49 clusters -> 1) and unified SQL across Kafka + Iceberg differentiators support this framing. Could command premium pricing i

Hurts: Requires multi-year, multi-million-dollar marketing investment to educate buyers on why they need a new category instead of picking best-of-breed tools. Data engineers are conservative and prefer proven categories with peer references and ecosystem integrations. PhoenixAI's $80M Series B and ~70-80 employees are insufficient to create and own a new category against entrenched competitors (ClickHou

HOW TO VALIDATE

The proposed positioning is simply a starting point, and it should be validated with real users or potential users through getting answers to the following questions:

- When data engineers evaluate PhoenixAI against ClickHouse, do they care more about eliminating the FINAL tax or eliminating denormalization workarounds? Which pain point drives the conversation?
- Do buyers recognize 'real-time analytical database' as a category, or do they search for 'ClickHouse alternative,' 'streaming analytics database,' 'OLAP for agents,' or something else? What language do inbound leads use?
- Is workload isolation (isolated warehouses for agent/BI/ETL) a top-3 buying criterion, or is it a nice-to-have that only matters after query performance and freshness are proven?
- Do data engineers need proof that PhoenixAI handles AI agent workloads specifically, or is general sub-second query performance + high concurrency proof sufficient to win the deal?
- When PhoenixAI loses to status quo (no decision), is it because migration risk and operational effort outweigh the pain, or because the current ClickHouse pain isn't severe enough yet?

VALUE PROPOSITION AND MESSAGING DRAFT

Consolidating the above insights and positioning, an initial value proposition, 30-second pitch, and proof pillars can be drafted. Note that this is just an initial draft. Further validation and adjustments should take place before finalization. This is simply a starting point: [8]

Value Proposition

> PhoenixAI lets you query live Kafka and Iceberg tables together without denormalizing schemas, waiting for merges, or managing 49 clusters.

(Continued on next page)

[7] Source: Mktforge Build Positioning (Stages 0–5) — Data Engineer vs. clickhouse.com, generated 10/25/2026

[8] Source: Mktforge Draft Messaging (Stages 6–9) — Data Engineer, built from Positioning Document Against ClickHouse for Data Engineers in any industry.pdf, generated 10/25/2026

30-Second Pitch

PhoenixAI is a real-time analytical database for data engineers building AI agents and customer-facing analytics. You query live Kafka feeds and Iceberg tables together with sub-second latency without denormalizing schemas, waiting for background merges, or managing separate clusters. Kafka CDC becomes queryable in under 5 seconds, joins run on normalized data without crashing, and the cost-based optimizer handles unpredictable agent SQL automatically. Demandbase consolidated 49 ClickHouse clusters to 1 warehouse, cut hardware 60%, and shipped agent features in under two months.

Pillar 1 — Query streaming and lakehouse data together without denormalizing your

Answers: ClickHouse forces you to choose between slow distributed joins that crash on complex queries or denormalizing into wide tables that waste 10× the storage, require six-month backfill projects whenever dimension data changes, and break real-time CDC pipelines—or you spin up 49 separate clusters just to isolate workloads.

Capabilities:

- Run joins across normalized Kafka feeds and Iceberg tables with sub-second latency on billions of rows—ship agent features and dashboards without flattening your schema first
- Consolidate agent, BI, and ETL workloads on the same data with independent compute per warehouse—scale agent compute to 10 nodes while BI runs on 2, same tables
- Avoid the six-month schema redesign and backfill cycle every time a new dimension or join requirement arrives

Proof: Demandbase consolidated 49 ClickHouse clusters to 1 PhoenixAI warehouse, cut hardware 60% and storage 90%, and shipped agent features in under 2 months without denormalizing.

Pillar 2 — Serve fresh CDC data in seconds without killing query performance

Answers: ClickHouse ReplacingMergeTree forces a choice between stale duplicate rows in query results until background merges complete (minutes to hours) or using SELECT FINAL which is 5-20× slower and single-threaded on large tables—making sub-second freshness impossible without destroying latency.

Capabilities:

- Commit Kafka and Flink CDC upserts queryable in under 5 seconds at 100K+ events/sec—agents and dashboards see the latest version without waiting for merges
- Query mutable streaming data with sub-second latency at 10K+ concurrent queries—no FINAL keyword, no stale duplicates, no performance cliff
- Ship fraud detection, delivery ops, and agent features that require seconds-level freshness without rewriting your CDC pipelines or adding a speed layer

Proof: Coinbase ingests 30K messages/sec across 573B+ rows and 300+ tables with <1s query latency; Pinterest hit 50% lower p90 latency vs. Druid with 10-second data freshness.

(Continued on next page)

[8] Source: Mktforge Draft Messaging (Stages 6–9) — Data Engineer, built from Positioning Document Against ClickHouse for Data Engineers in any industry.pdf, generated 10/25/2026

Pillar 3 — Run unpredictable agent SQL on production data with cost-based optimization

Answers: LLM-generated SQL with novel join patterns breaks or times out on ClickHouse because its rule-based planner can't reorder joins at runtime, apply global filters, or adapt to query patterns it hasn't seen—forcing you to manually tune every new agent query or restrict agents to pre-denormalized views.

Capabilities:

- Handle agent-generated joins and filters automatically with cost-based optimizer that reorders joins, pushes down runtime filters, and uses query feedback—no manual tuning per query
- Run agent workloads on the same fresh normalized data as BI dashboards without separate denormalized copies or clusters
- Ship agent features in weeks instead of waiting quarters for schema rewrites and performance testing

Proof: Conductor runs agent queries scanning 6-7 GB / 180-240M rows in 389ms-2.8s with 100% decimal accuracy, shipped in <2 months; Demandbase replaced a ClickHouse architecture that couldn't handle agent-generated joins.

NEED MORE INFORMATION?

This brief was crafted using guided research and materials generated through the Mktforge platform. The information contained in this document is a high level abstraction of the resources it was sourced from. For more information and context, refer to the sources listed in the footers of this brief belonging to the Mktforge account used to generate them.

Want more briefs like this?

Mktforge turns your company's research into ready-to-share marketing foundations.

Learn more at Mktforge.io

Mktforge