Find My Customer — Results

Company URL	https://www.phoenixdata.ai
Generated	10/25/2026, 10:54:25 AM

Customer List

Customer	Size	Industry
Eightfold AI	~800-1,000 employees	HR Tech / Talent Intelligence Software
Demandbase	~900-1,000 employees	B2B Marketing / Account-Based Marketing (ABM) Software
Conductor	~300-500 employees	SEO/AEO Marketing Technology
Coinbase	~5,000 employees	Cryptocurrency / Financial Services
Pinterest	~5,000 employees	Social Media / Internet
Fanatics	~10,000+ employees	E-commerce / Licensed Sports Retail
Intuit	~18,600 employees	Financial Software / Fintech
Expedia Group	~16,000-17,000 employees	Online Travel / E-commerce

Job Titles

Job Title
Senior Engineering Manager, Data Platform
Senior Staff Software Engineer
Data Platform Engineer
Chief Product Officer
Senior Software Engineer, Real-Time Analytics
Data Engineer
VP of Engineering / VP of Data Platforms
Analytics Engineer
Head of Data Platform
Chief Technology Officer

Pain Points / Initiatives

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

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

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

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 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

Buyers to Target

Senior Engineering Manager, Data Platform

- Isolated agent warehouse so AI agent query bursts don't contend with the BI warehouse on the same data
- Optimizer automatically handles novel, unplanned joins generated by AI agents without manual tuning
- Sub-second queries directly on Apache Iceberg tables with no need to pre-flatten or pre-aggregate
- Ability to query petabytes of normalized data across thousands of tenants
- Customer-facing dashboards keep second-level SLAs even under heavy agent load
- Major infrastructure cost reductions (60% hardware reduction, 90% storage cost reduction cited)
- Cluster consolidation (49 clusters down to 1) simplifying operations
- Replacing an unpredictable ClickHouse-based architecture with a database purpose-built for agent workloads

Data Engineer

- Native streaming ingestion from Kafka, Flink CDC, and Spark without custom pipelines
- Unified SQL across streaming and lakehouse data (Iceberg, Delta, Hudi) eliminates duplicate pipelines
- Data queryable within seconds of ingest
- No mandatory denormalization or pre-aggregation reduces ETL complexity
- Query the lakehouse in place without copying data into a separate warehouse
- Fine-grained access controls and per-query audit for governance
- 90% lower infrastructure footprint than other real-time analytics solutions
- Isolated warehouses let BI and AI/agent workloads run on the same data without contention

Competitors To Watch

Competitor	URL	Why they compete
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.
Apache Doris (VeloDB)	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 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.
Firebolt	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.
ClickHouse	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

Competitor	URL	Why they compete
		workloads, real-time streaming, and complex joins to same data platform engineers—PhoenixAI's messaging explicitly positions against legacy ClickHouse architectures.
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.
SingleStore	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.
Tinybird	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.