



# From humans to

And what it means for the data platform

# agents

Aaron Katz

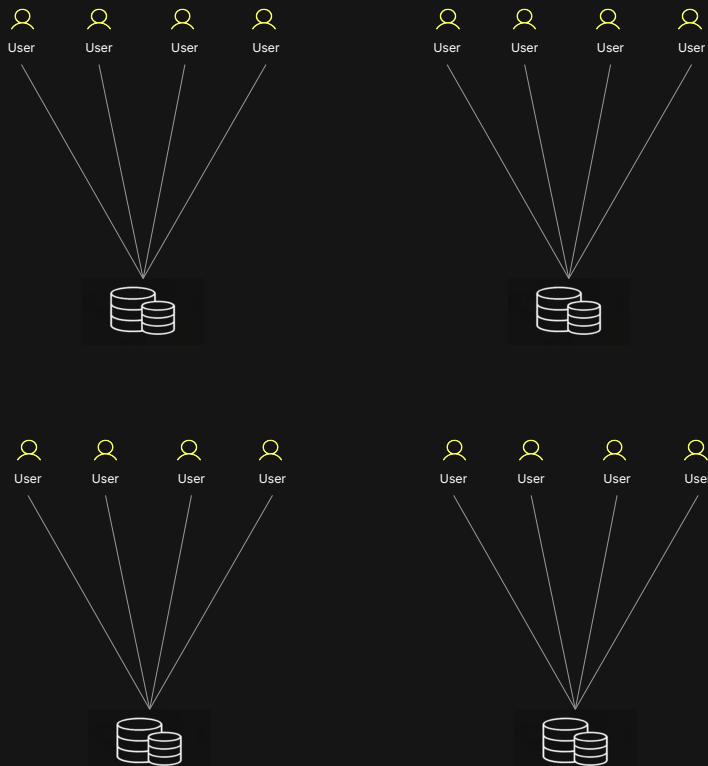
Co-Founder and CEO ClickHouse

**AI adoption will be hindered  
not by model choice and  
token costs but the **data**  
**layer****

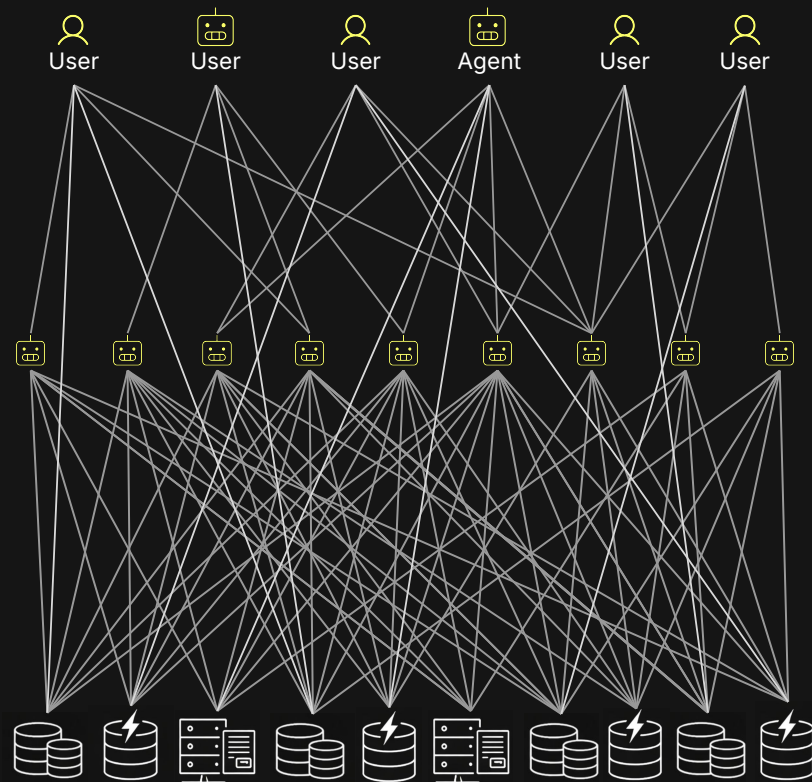
Agents need access to **data** —

at high **concurrency**,  
interactive **latencies**, and  
disruptively lower **cost**

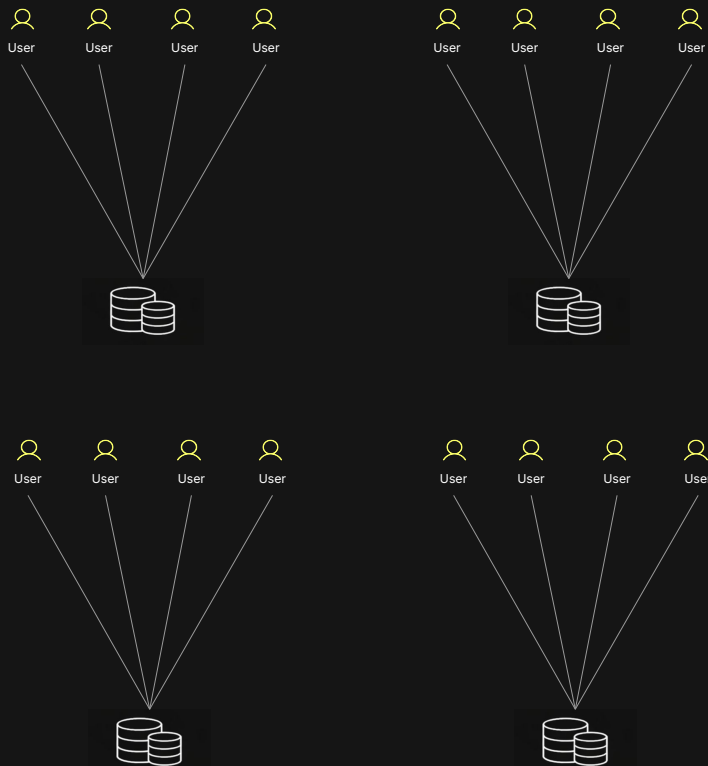
# Old world




# AI era



# Old world



# AI era



**This tectonic shift affects every part of the business**

**Developers**  
Building Applications

OLTP DATABASES



REAL-TIME ANALYTICS



**Data Engineers & Scientists**  
Supporting Business Analytics

DATA WAREHOUSING



DATA SCIENCE



**SREs**  
System Health

OBSERVABILITY



## Developers Building Applications

UNIFIED REAL-TIME & AI APPS DATA STACK



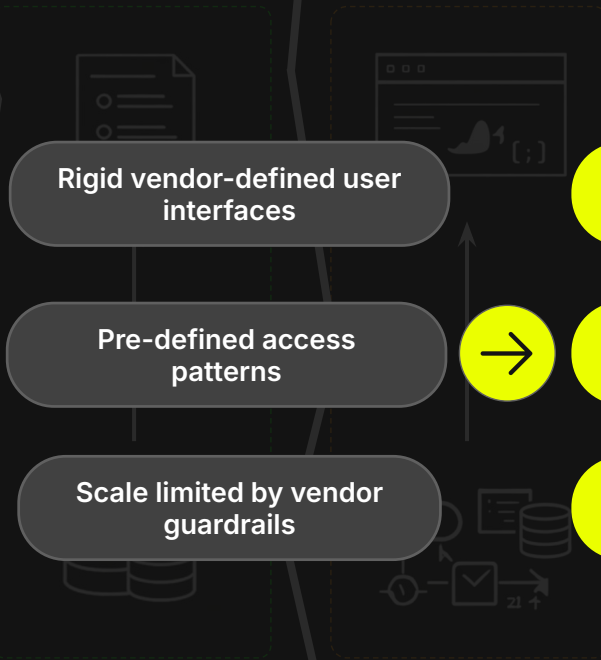
OLTP

REAL-TIME ANALYTICS

Natural Language Co-Pilots  
AI-Powered Workflows

## Old World

DATA WAREHOUSES DATA SCIENCE

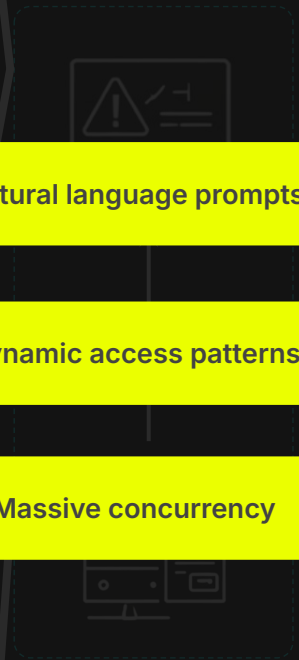


Rigid vendor-defined user  
interfaces

Pre-defined access  
patterns

Scale limited by vendor  
guardrails

## AI Era



Natural language prompts

Dynamic access patterns

Massive concurrency



## Powering real-time AI finance products

### DETAILS

**Postgres no longer scaled to build  
AI-powered spend insights and controls  
at scale.**

A completely new backend powered by ClickHouse was needed to resolve the problem.

## AI-POWERED PRODUCTS

**40x**

Faster query performance

**50k**

Ramp customers served



*"What's really powered our growth is ClickHouse—it's enabled agentic AI features like Ramp Budgets and unlocked new revenue opportunities."*

Ryan Delgado, Director of Engineering

# Data Engineers & Scientists Supporting Business Analytics

UNIFIED REAL-TIME & AI APPS DATA STACK

## Old World

## AI Era

Batch latency

Real-time responses

Limited concurrency

Massive concurrency

OLTP

REAL-TIME ANALYTICS

Natural Language Co-Pilots  
AI-Powered Workflows

DATA WAREHOUSING

DATA SCIENCE

OBSERVABILITY



Conversational BI  
Agentic Analytics



Memorial Sloan Kettering  
Cancer Center.

## Conversational AI Agent for Faster Cancer Genomics Research

### DETAILS

**Memorial Sloan Kettering cBioPortal is a free cancer genomics research platform. Once the portal team added chat interface to their platform, legacy backend no longer scaled.**

cBioPortal now uses ClickHouse to enable real-time exploration of large cancer datasets and natural-language querying for thousands of daily researchers.

## AI-POWERED MEDICAL RESEARCH

**4 B+**

Genomic data points

**10x**

Faster queries



*We were inspired seeing what cancer researchers were doing with ClickHouse MCP and LibreChat.*

Ino de Bruijn, Data Visualization Team Lead

**SREs  
System Health**

UNIFIED REAL-TIME & AI APPS DATA STACK


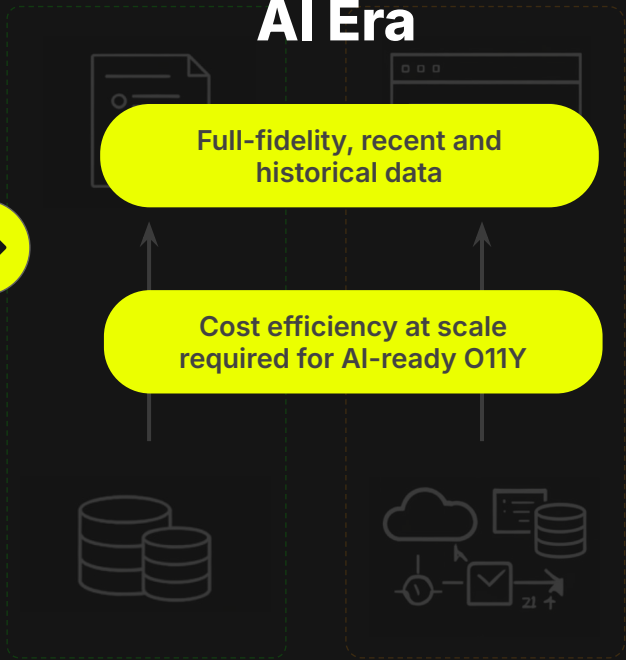
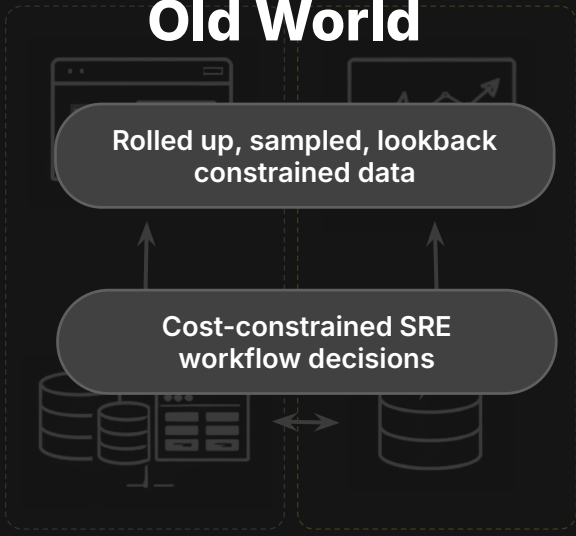
DATA WAREHOUSING

DATA SCIENCE

OBSERVABILITY

# Old World

# AI Era



Natural Language Co-Pilots  
AI-Powered Workflows

Batch-Oriented Legacy DWH  
& Data Science Workflows

AI SRE

## AI-POWERED OBSERVABILITY

# ANTHROPIC

## Developing frontier AI models

### DETAILS

Anthropic chose ClickHouse to ingest petabyte-scale observability data in real time, run fast and interactive analytics over semi-structured data, and **scale costs predictably as Claude usage skyrocketed.**

# 1M

new sign ups for Claude every day

# 43%

month-over-month web traffic growth in Feb 2026



*ClickHouse played an instrumental role in helping us develop and ship Claude 4.*

**Maruth Goyal, Member of Technical Staff at Anthropic**

# Harvey

## Powering an agentic Security Operations Center

### DETAILS

Harvey runs an **always-on AI SOC** of agents that hunt, triage, and investigate threats around the clock that reason over terabytes per day of telemetry served through optimized ClickHouse tables.

## AI-POWERED SECURITY

**+5.7x**

Detection coverage

**~90%**

Weekly alert volume  
dropped

“

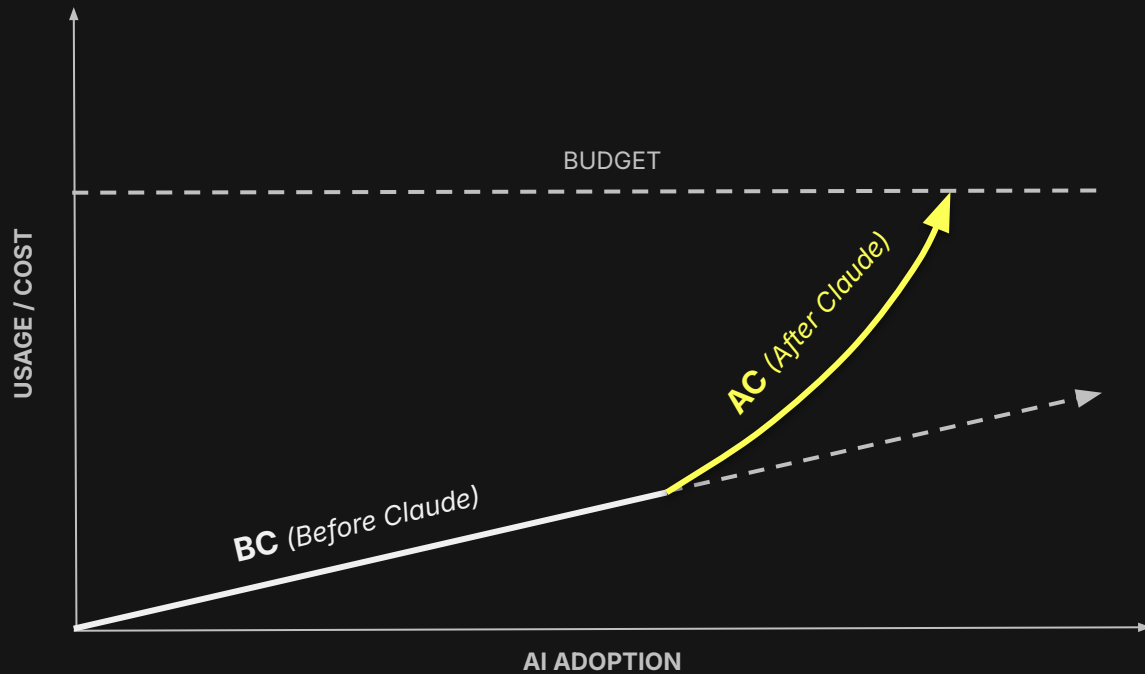
*Invest in your log warehouse before you invest in your agents.*

*“A typical agent investigation issues dozens of queries; the difference between 200ms and 2s per query is an agent that explores 3 hypotheses **versus one that explores 30.**”*

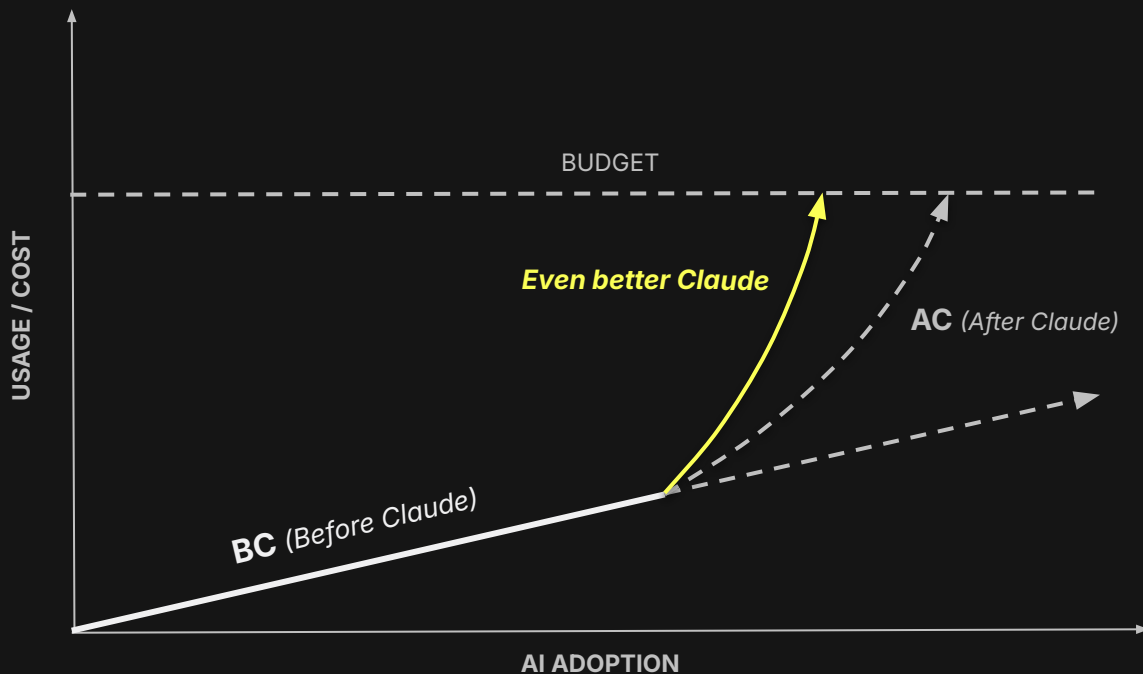
*Mike Parowski, Security Engineer at Harvey*

Concerns about **rising**  
**costs of AI** are real

# Agents **supercharge** usage



# As agents improve, usage goes **exponential**



# Tokens costs aren't the problem

## Cost per query is

### LLM TOKENS

Cheap and getting cheaper.

EARLY 2023

**\$30**

/ million tokens



TODAY

**\$1**

/ million tokens

30x cheaper.

### DATABASE QUERIES

Flat, or creeping up.

EARLY 2023

**\$5**

/ TB scanned



TODAY

**\$6.25**

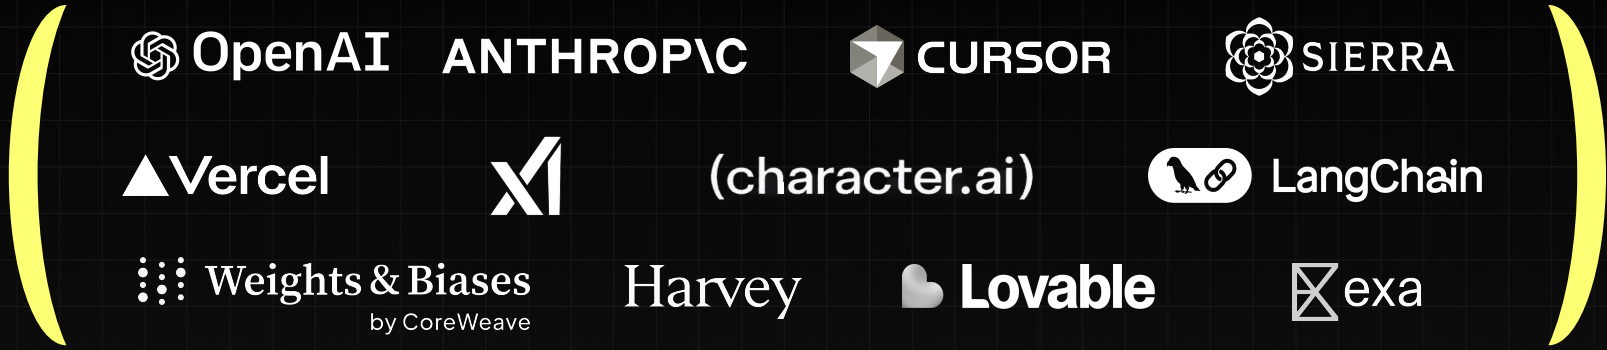
/ TB scanned

Multiplied by **agent-level traffic**.

**Optimizations outside of  
data platforms buy time -  
not an answer**

**Data platforms  
need to adapt**

# Industry is moving too fast for incremental changes



3 Million apps and more than 25k apps deployed daily.

 **Lovable**

"When you're digging into logs during an incident, every second matters."

(character.ai)

"We needed a completely new data platform to build and ship Claude 4."

**ANTHROPIC**

# Questions?

# Appendix

# Customer Examples

## REAL-TIME ANALYTICS & OBSERVABILITY

# OpenAI

## Powering real-time visibility across model research

### DETAILS

OpenAI chose ClickHouse for its open-source flexibility, horizontal scalability, and SQL-native design that could handle petabyte-scale log ingestion without vendor lock-in and lower operational overhead.

**40%**

reduction in CPU usage

**50%**

log volume spike overnight seamlessly handled



*"The scale is absolutely mind boggling. The cloud-nativeness and open-sourceness of ClickHouse really saved us when we were in a pinch."*

**Akshay Nanavati, Engineering Manager, OpenAI**

## AI-POWERED OBSERVABILITY

# ANTHROPIC

## Developing frontier AI models faster and more securely with ClickHouse

### DETAILS

Anthropic chose ClickHouse to ingest petabyte-scale observability data in real time, run fast and interactive analytics over semi-structured data, and scale costs predictably as Claude usage skyrocketed.

# 1M

new sign ups for Claude every day

# 43%

month-over-month web traffic growth in Feb 2026



*ClickHouse played an instrumental role in helping us develop and ship Claude 4.*

**Maruth Goyal, Member of Technical Staff at Anthropic**

# character.ai

**Reducing observability costs by 50% while scaling logs 10x**

## DETAILS

Character.AI runs a centralized observability platform on ClickHouse Cloud, ingesting billions of logs daily and giving engineers real-time visibility across thousands of GPUs and Kubernetes clusters.

## OBSERVABILITY

**450 PB**

Raw log data ingested monthly

**20x**

Better compression



*When you're digging into logs during an incident, every second matters. With ClickHouse, it was just a case of how fast I could blink. The performance is real.*

Mustafa Yildirim, Site Reliability Engineer



## Powering real-time AI finance products with ClickHouse

### DETAILS

Ramp moved from Postgres to ClickHouse to build AI-powered spend insights and controls.

### AI-POWERED PRODUCTS

**40x**

Faster query performance

**50k**

Ramp customers served



*"What's really powered our growth is ClickHouse—it's enabled agentic AI features like Ramp Budgets and unlocked new revenue opportunities."*

Ryan Delgado, Director of Engineering



Memorial Sloan Kettering  
Cancer Center

## Conversational AI Agent for Faster Cancer Genomics Research

### DETAILS

MSK's cBioPortal uses ClickHouse to enable real-time exploration of large cancer datasets and natural-language querying via MCP and LibreChat for thousands of daily researchers.

## AI-POWERED MEDICAL RESEARCH

**4 B+**

Genomic data points

**10x**

Faster queries



*We were inspired seeing what cancer researchers were doing with ClickHouse MCP and LibreChat.*

**Ino de Bruijn, Data Visualization Team Lead**

# Harvey

## Powering an agentic Security Operations Center with ClickHouse

### DETAILS

Harvey runs an always-on AI SOC of agents that hunt, triage, and investigate threats around the clock that reason over terabytes per day of telemetry served through optimized ClickHouse tables.

### AI-POWERED SECURITY ANALYTICS

**+5.7x**

Detection coverage

**~90%**

Weekly alert volume dropped

“

*Invest in your log warehouse before you invest in your agents.*

*“A typical agent investigation issues dozens of queries; the difference between 200ms and 2s per query is an agent that explores three hypotheses **versus one that explores 30.**”*

*Mike Parowski, Security Engineer at Harvey*

# Appendix

# ClickHouse

## OUR MISSION

Make **all data**  
instantly  
available to:

✓ Every person

✓ Every agent

✓ Every application

# Requirements for the **AI era**



## Speed

Real-time queries at sub-second latency.



## Scale

Massive event and behavioral data - no ceiling.



## Simplicity

Fewer systems. One platform agents can talk to.



## Cost-efficiency

Predictable and low.



## SQL

A language agents already know. No glue code.



## APIs, not UIs

AI is programmatic. Every capability must be callable.

SAME REQUIREMENTS, 15 YEARS EARLIER

# We didn't build ClickHouse for AI

## We built it for the same requirements

Designed to power customer-facing, multi-tenant, real-time analytics — for millions of concurrent users.



### Fast and cost-efficient

Sub-second queries on billions of rows.  
10–20× compression. Fewer nodes,  
lower cost, faster results.



### Open source

Standard SQL, Apache 2.0, 150+  
integrations, strong community. No  
vendor lock-in.



### Secure and compliant

Transparent code auditing, robust  
access controls, encryption, industry  
standards.



### Proven at scale

Petabytes of data. Storage and compute  
scale independently.

# ClickHouse is the world's **fastest** data platform.

At any scale.

For every workload.

**1B+**

INGESTED ROWS / SEC

**PB+**

SCALE

**<1s**

QUERY LATENCY

**1000+**

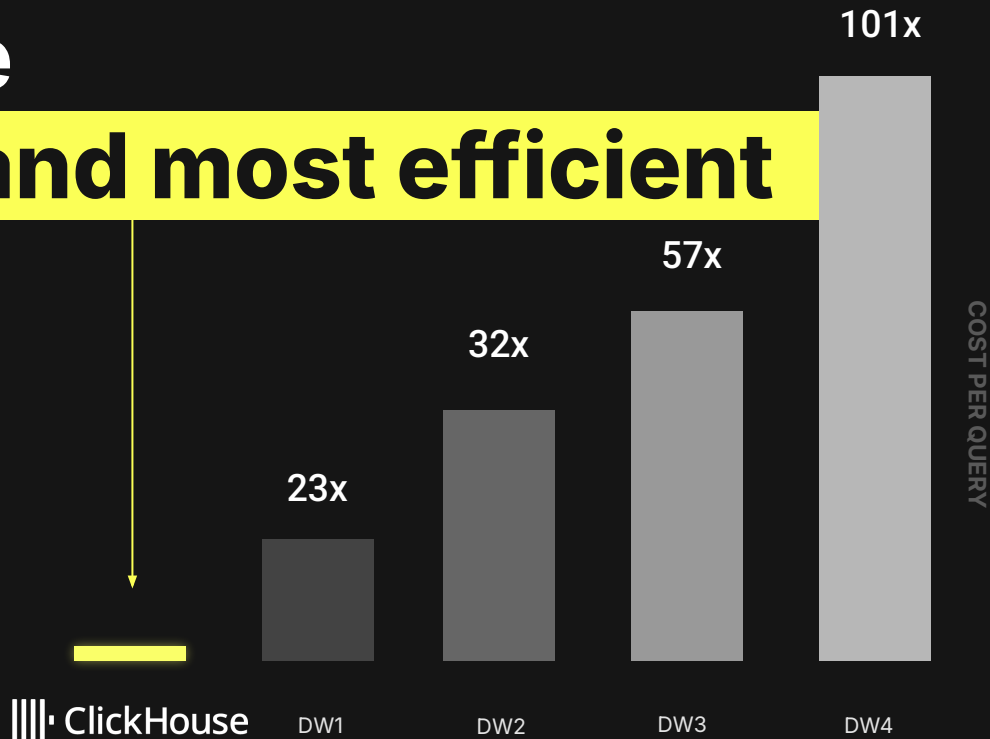
NODE CONCURRENCY

# ClickHouse is the world's fastest data platform.

# and most efficient

At any scale.

For every workload.



# Customer growth

4,000+ in the last 4 years. 1,000 in the last 3 months.

ANTHROPIC

ramp

Weights & Biases

Meta

CoreWeave

CURSOR

mercado  
libre

Vercel

GitLab

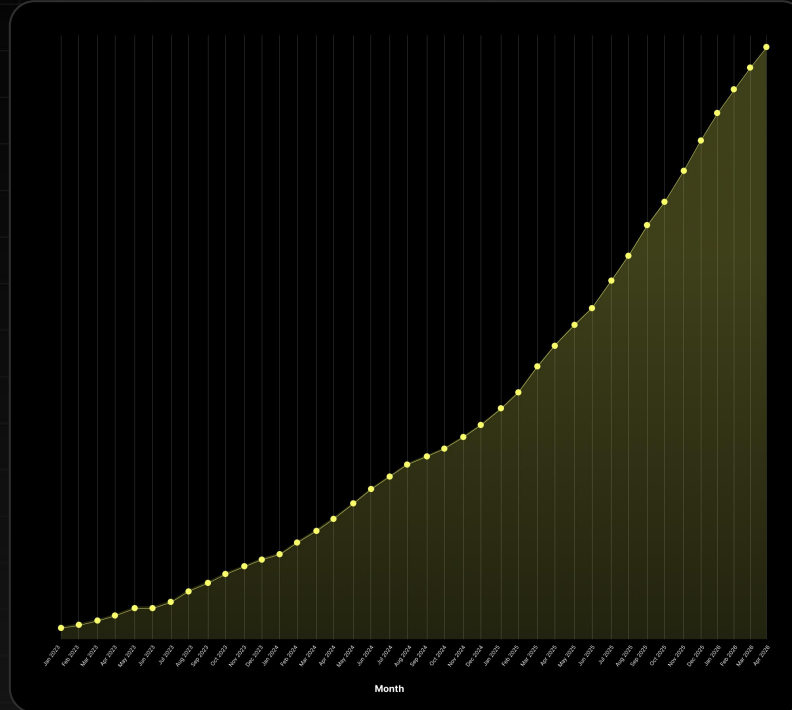
Lovable

Hewlett Packard  
Enterprise

braze

DOORDASH

VISA



twilio

HubSpot

SIERRA

lyft

Tesla

Temporal

ROBLOX

CISCO

LangChain

Property  
Finder

solarwinds

Memorial Sloan Kettering  
Cancer Center

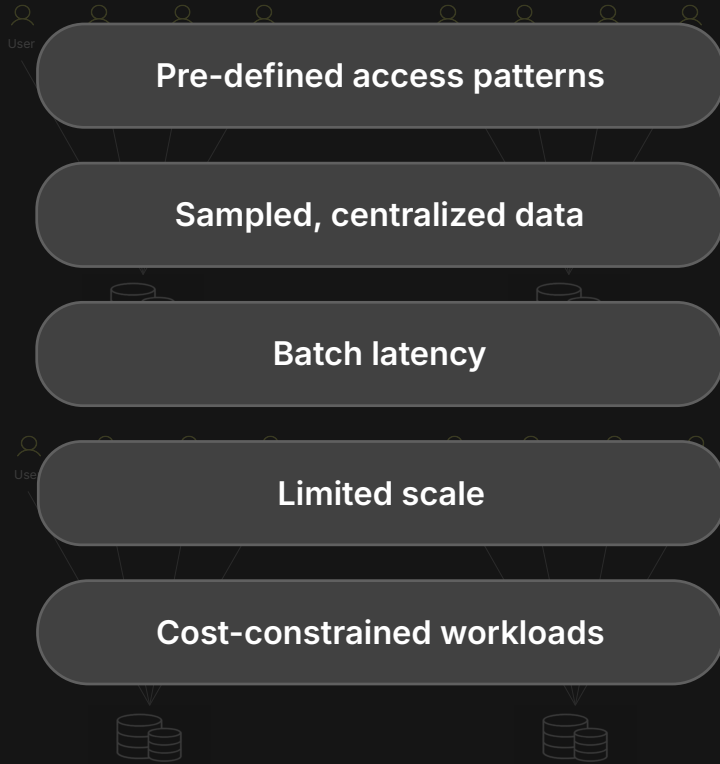
Vantage

Nansen

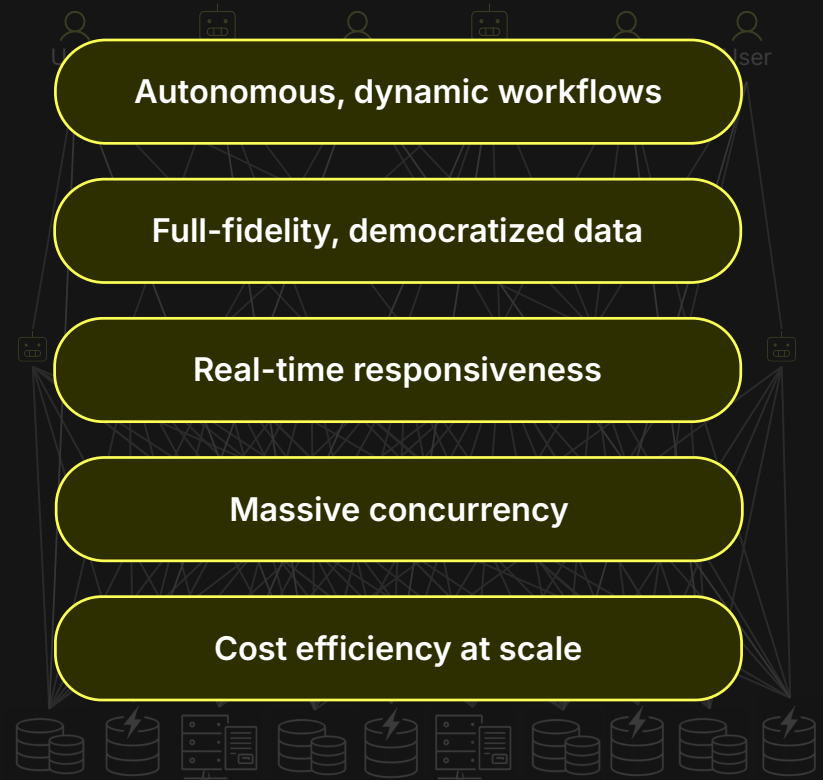
attentive

ClickHouse

# Old world



# AI era



# Old world



*Trade-offs:*

**Scale**

**VS**

**Speed**

**Volume**

**VS**

**Cost**

**Simplicity**

**VS**

**Flexibility**



# AI era



*Baseline requirements:*

**Scale**

**AND**

**Speed**

**Volume**

**AND**

**Cost**

**Simplicity**

**AND**

**Flexibility**