


Your agents' memory lives in your own bucket.

Graph, vectors and search in one open-source engine, running on the storage you already pay for — in your cloud, or on your own metal.

In the gold rush, the ones who sold the tools got rich.

In the AI rush, we sell the technology the tools are built on.



WHY NOW

Object storage became a database primitive 21 months ago.

THE WINDOW OPENED IN 2024

Before 2024 an atomic commit on S3 needed an outside coordinator — DynamoDB, etcd, ZooKeeper. Now the write itself is the coordinator.

AUG 2024 S3 ships If-None-Match

NOV 2024 compare-and-swap on ETag

DATA WAREHOUSE

Snowflake

S3 as primary storage, compute billed apart.

STREAMING

WarpStream

Kafka straight to S3. Bought by Confluent, 2024.

SERVERLESS POSTGRES

Neon

Pages in object storage. Bought by Databricks, 2025.

VECTOR SEARCH

turbopuffer

\$70 per TB against \$3,600 on RAM and SSD.

Every database category already moved to object storage. Graphs are the last one still billing by the gigabyte of RAM.

THE PROBLEM

Priced by RAM, right as the workload became a million idle graphs.

01

There is nothing cheap to start with

Neo4j Aura's cheapest paid plan is \$65 a month for a single gigabyte of memory, with a one-gigabyte minimum cluster. Below that there is only a free tier.

02

The workload inverted underneath the pricing

Agent fleets go from fewer than 15 at the average Fortune 500 today to over 150,000 by 2028. Most of them are asleep most of the time.

03

Regulated buyers are locked out entirely

Sovereign workloads cannot ship their graph to somebody else's cloud. And the open embedded option vanished when Kùzu was archived in October 2025.

\$66,560

1 TB of graph in Neo4j AuraDB memory

VS


\$23.55

the same terabyte in Amazon S3

≈2,800×

at published list prices

We made the bucket the database.



Priced like storage

You pay for bytes at rest, not for reserved memory. A sleeping agent costs nothing.

Runs anywhere

Same engine in your cloud, in your data centre, or air-gapped with no outbound connection.

Nothing to leave behind

Open source, open formats, your storage. Walking away costs you nothing.

One engine, three deployments, one query across all three.

```
namidb - agent memory

$ pip install namidb

import namidb
mem = namidb.Client("s3://acme/agents/42")

# one query across meaning and relationships
mem.cypher("""
    CALL search.hybrid('memory', $q, 20)
    YIELD node AS fact
    MATCH (fact)-[:LEARNED_FROM*1..3]→(e:Event)
    RETURN e.what, e.when ORDER BY e.when DESC
""", q=question)

# the memory is just files you own
$ aws s3 sync s3://acme ./backup

no servers · no cluster · no on-call
```

Embedded

SHIPPED

Import the library, point it at storage, run Cypher inside your process.
The free front door.

Server

SHIPPED

One Rust binary. Speaks HTTP and Bolt, so existing graph drivers connect unmodified. Runs air-gapped.

Cloud

CLOSED BETA

A namespace per tenant or per agent, scaling to zero when nobody queries. This is where the revenue is.

TRACTION

The engineers who build the world's infrastructure found us first

COMMUNITY

Weeks after the repository went public, its stars already came from inside these companies

Cloudflare

Microsoft

Intel

Alibaba

Tencent

Red Hat

Dropbox

AT&T

Fly.io

Weaviate

STRATEGIC PARTNER

GRD Radical Corporation C.L.

25 years, 8 countries, 3,000+ organisations protected across LATAM, its own 24/7 SOC with an RFC 2350 certified CERT and ISO 27001. A standing government supplier — and the channel NamiDB sells sovereign through.

25 YEARS · 8 COUNTRIES · 3,000+ ORGANISATIONS · ISO 27001

FRAMEWORK AGREEMENT SIGNED

Universidad Regional Autónoma de los Andes

Signed with the university as an institution, not a pilot with one lab: academic research, products and startups built on NamiDB — real workloads, and the people who carry the technology to market with us.

INSTITUTION-WIDE · RESEARCH · PRODUCTS · TALENT PIPELINE

The engine is free; the Cloud is the business.

Free	\$0
a few GB to start, self-hosted or managed	
Starter	\$25 / mo
more GB included, then pay for what you use	
Growth	\$600 / mo
team workspaces, SSO, support SLA	
Enterprise	custom
on-premise, air-gapped, sovereign, OEM	

OPEN SOURCE

The engine

Open code from day one under BSL 1.1, converting to Apache 2.0. Developers adopt it without asking permission and carry it into work.

REVENUE 01

NamiDB Cloud

A subscription today, moving to usage-based billing as the Cloud matures. An idle workspace costs almost nothing.

REVENUE 02

Sovereign & enterprise licence

On-premise and air-gapped deployments for government and regulated buyers, plus OEM licensing for companies embedding the engine.

THE MARKET

We're after a \$3.6B market compounding at 25% a year.

TAM

\$3.6B

Global graph database market in 2026, growing to \$14–25B by 2031–2035 at a 24–27% CAGR.

SAM

\$2.0B

The cloud share of that market in North America and Europe — cloud is 71–74% of category spend.

SOM

\$40–100M

The reachable share of that segment on a five-year horizon, where three bottom-up paths converge.

WHERE THE SPEND SITS

Cloud is 71–74% of category spend

and it is the segment compounding fastest, toward \$14–25B by 2031–2035.

WHY THE SOM HOLDS

Three independent bottom-up paths — top-down share, enterprise accounts and a developer funnel — land in the same band. Appendix 15.

Sources: Mordor Intelligence, Fortune Business Insights, Precedence Research (2025–2026).

Developers adopt it; sovereign buyers cannot get it anywhere else.

MOTION 01 · developers adopt it before anyone signs anything



MOTION 02 · sovereign and public sector

Runs air-gapped on the buyer's own hardware — no external dependency, no vendor in the path
EU public procurement now applies an open-source-first principle to cloud and AI purchases

Everyone bills you for RAM. We bill you for bytes at rest.

	Neo4j / AuraDB	Amazon Neptune	Memgraph	NamiDB
Where the data lives	Block storage with a RAM page cache	AWS-managed distributed volume	RAM — the whole graph must fit	Storage you already own
What you are billed for	GB of memory per month, \$65–\$146	Instance hours, NCUs, GB and I/O	Memory capacity, by their own wording	Bytes at rest and queries served
Cheapest paid plan	\$65 / month, 1 GB minimum	\$250 / month, smallest instance	Contact sales	\$25 / month, with a free tier
Cost of an idle database	Full price of the reserved memory	Full price of the instance you left on	Full price of the reserved memory	Storage only — compute goes to zero
Runs air-gapped on your hardware	Enterprise licence, still RAM-priced	No	Enterprise licence, still RAM-priced	Yes — same engine, no connection out
Licence	Open core: GPLv3 or commercial	Proprietary, AWS only	BSL 1.1 plus enterprise licence	BSL 1.1, converts to Apache 2.0

We are not trying to win a single-machine benchmark. We are changing what a graph database costs to own.

Three founders: the engine, the market, the paperwork



Matías Fonseca

CEO · NAMIDB, INC.

Founder of Fonles Studios (Ecuador, Peru, El Salvador), maker of Alhelí — a legal AI that holds up to 60,000 PDF pages in a single context, sold to 500+ clients. Writes the NamiDB engine and its 32 public RFCs.

500+ CLIENTS · 60,000-PAGE CONTEXT · 32 RFCs



Claudio Cortés

FOUNDER · GO-TO-MARKET

CEO of GRD Radical Corporation C.L. — 25 years, 8 countries, 3,000+ organisations protected in LATAM with its own 24/7 SOC and a standing government practice. Previously 3Com, 2001–2009.

3,000+ ORGS · 8 COUNTRIES · 25 YEARS



Darío Fonseca

FOUNDER · LEGAL

20 years in law and public administration, having led teams of over 100 people inside government. Co-founder of Fonles Studios. Owns NamiDB's corporate structure, licensing and multi-jurisdiction operations.

20 YEARS · 100+ PEOPLE LED · MULTI-JURISDICTION

We're raising \$1M to put the Cloud in the market

Pre-seed · post-money SAFE · standard terms · 18 months of runway

USE OF FUNDS

Marketing & distribution

50%

The engine is built; almost nobody knows it exists. Brand, developer community, public-sector channel and a commercial lead.

Product & engineering

35%

Cloud to public launch, usage-based billing, the sovereign on-premise distribution, and the first engineering hires.

Legal & operations

15%

Compliance, agreements, public-sector certifications and intellectual property.

WHAT THIS BUYS, IN 18 MONTHS

The Cloud goes public

Out of closed beta and onto a public price list, with usage-based billing behind it.

The first sovereign deployment

An air-gapped install running inside a public-sector buyer, with the certification pack to repeat it.

The default memory under agents

Shipped integrations beneath the agent-memory frameworks, so choosing NamiDB is a one-line change.

The graph is the shape of how things relate.

We're giving it a database worthy of the decade ahead.


GET IN TOUCH

Matias Fonseca • CEO

matias.fonseca@namidb.com

namidb.com

github.com/namidb/namidb



Why a bucket can be a database.

Storage

An LSM tree over object storage, coordinated by compare-and-swap on the manifest. Columnar SSTs derived from Parquet, sidecar bloom filters, row-group pruning, and range GETs instead of whole-object reads.

Query

A cost-based optimiser with cardinality estimation and join ordering, worst-case optimal joins via leapfrog triejoin, and factorised intermediate results that avoid materialising cross products.

Retrieval

DiskANN vector index with int8 quantisation and native multi-tenant pre-filtering, BM25 full-text with real IDF, and reciprocal rank fusion to combine both in a single query.

Operations

Single writer per namespace with lock-free concurrent reads, background compaction with per-level byte budgets, snapshot-aware orphan sweeps, and 32 public RFCs documenting every decision.

THE OBJECTION, ANSWERED · THIS IS NOT MEMORY LATENCY

RAM answers in microseconds; object storage answers in milliseconds. For agent memory the bottleneck is the model's own inference, measured in hundreds of milliseconds — not the read. Caching recovers the hot path, and turbopuffer already proved the trade commercially in vector search.

Three independent paths to the same order of magnitude.

TOP-DOWN

0.5–1% of the cloud graph market

\$50–100M ARR

A modest share of the North America and Europe cloud segment as it compounds toward 2031.

ENTERPRISE BOTTOM-UP

750 accounts at \$50k

\$37.5M ARR

1% of the ~75,000 companies above 250 employees in the EU and 500 in the US.

DEVELOPER BOTTOM-UP

2% OSS adoption, 1–2% paid conversion

\$15–88M ARR

Of 7.3M cloud-native AI developers, applying open-source commercial conversion benchmarks.

The convergence is the argument: a top-down share estimate, an enterprise account count and a developer funnel — built from different data — land within the same band.

Market size: Mordor Intelligence, Fortune Business Insights, Precedence Research (2025–2026); MarketsandMarkets reports a smaller figure on a narrower definition. Cloud share: Mordor 71.5% (2025), Fortune 73.8% (2026). Developers: SlashData, May 2025; CNCF and SlashData, Q1 2026. Enterprises: Eurostat, December 2025; US SBA, December 2024. Agents: Gartner, April 2026. Sovereign cloud: Gartner, February 2026.