

# SochDB

---

**The enterprise data platform for agentic AI. One system of record for embeddings, structured data, and agent memory — with the compliance, scale, and SLAs the Fortune 500 actually buys.**

Series Seed Pitch · 2026

# 1 • The \$40B Problem Inside Every Enterprise AI Program

Every Fortune-500 AI platform team is running the same broken architecture:

Layer	Vendor	Annual cost (mid-size deployment)
Vector DB	Pinecone / Weaviate Cloud	\$250K – \$1.2M
Operational DB	Postgres / Aurora	\$150K – \$400K
Cache / state	Redis Enterprise	\$80K – \$250K
Memory layer	Mem0 / custom	\$50K – \$200K
Glue + retrieval engineering	4–8 FTEs	\$1.5M – \$3M
<b>Total per program</b>		<b>\$2M – \$5M / year</b>

And after all that spend they still get: - ❌ No cross-system transactions → silent data drift - ❌ No unified audit trail → compliance friction (SOC 2, HIPAA, GDPR, EU AI Act) - ❌ Token-budget bugs in production prompts - ❌ A separate vendor invoice, contract, and DPA for every layer - ❌ Procurement & InfoSec reviews per vendor — 6–9 months *each*

*"Our biggest LLM-program cost isn't tokens. It's the four databases under it." — Director of AI Platform, Fortune 100 retailer*

## 2 • The Solution — SochDB Enterprise

**One licensed, supported, audited platform** that replaces the vector DB + relational DB + memory layer + retrieval glue with a single system of record for AI workloads.

- **Unified data plane** — SQL + vectors + graph memory + token-aware context, ACID across all of them
- **Enterprise security** — mTLS, RBAC, SSO/SAML, encryption at rest, audit log export
- **Compliance-ready** — SOC 2 Type II, GDPR, HIPAA-aligned controls, EU AI Act traceability
- **Distributed scale** — sharded HNSW, replication, S3/Azure Blob/GCS archive, billion-vector tier
- **Operational maturity** — Helm chart, Prometheus + Grafana, PodDisruptionBudgets, rolling upgrades
- **Cloud-native marketplaces** — Azure AKS App, AWS Container Product, GCP Marketplace (BYOL + metered)
- **Dedicated support** — 4-hour response, named CSM, architecture reviews

One vendor. One contract. One audit. One invoice.

### 3 · Why Enterprises Buy Now

1. **Vendor consolidation is the #1 CIO priority for 2026.** Gartner: 78% of enterprise CIOs are actively cutting their AI vendor count.
2. **The EU AI Act and U.S. executive orders demand auditable retrieval.** Enterprises need provable data lineage from prompt → context → answer. Most stacks cannot produce that today.
3. **Vector-DB pricing is in revolt.** Pinecone p1.x6 pods, Weaviate Cloud Enterprise, and Datastax Astra are all >\$1M ACV at scale — with no relational, no memory, no compliance bundled.
4. **AI platform teams are forming budgets now.** Every F500 has a centralized "AI Foundation" org and a 2026 line item for the substrate. Whoever lands now becomes the standard.

## 4 · Product — Enterprise-Grade Today

### Distributed data plane

- Shard-first ANN topology, cluster-based query routing, fan-out reduction
- HNSW + Product Quantization + external vector storage (S3 / Blob / GCS)
- WAL durability, group commit, MVCC, Serializable Snapshot Isolation
- Columnar storage with projection pushdown
- PostgreSQL wire protocol — drop-in for existing BI / SQL tooling

### Security & governance

- mTLS, JWT/OIDC, cloud IAM integration
- Role-based access control + tenant isolation
- Audit log export to SIEM (Splunk, Datadog, Sentinel)
- Secrets via Kubernetes / Key Vault / AWS Secrets Manager / GCP Secret Manager
- Encryption at rest, customer-managed keys (BYOK)

### Operations

- Production Helm chart + StatefulSet + PDB
- Prometheus metrics, OpenTelemetry traces, Grafana dashboards
- gRPC, gRPC-Web, WebSocket, Postgres-wire entry points
- Rolling upgrades, zero-downtime version migration
- WAL archiving + point-in-time recovery

### AI / agent primitives that competitors don't ship

- ContextQuery builder with hard token budgets and multi-source fusion
- TOON dense output format — measurably reduces prompt cost
- Hybrid retrieval: vector + BM25 + Reciprocal Rank Fusion
- Graph overlay for agent memory + relationship traversal
- Policy hooks with audit trails (compliance-ready)
- Tool routing for multi-agent coordination

## 5 · Proof — Benchmarks Procurement Will Care About

### Vector workload · VectorDBBench (OpenAI 50K × 1536D)

Metric	SochDB	ChromaDB	LanceDB
Recall@100	0.9899	0.9966	0.6574
Avg latency	<b>3.3 ms</b>	15.4 ms	5.6 ms
P99 latency	<b>5.9 ms</b>	22.3 ms	12.2 ms
Insert 50K vectors	<b>0.1 s</b>	76.9 s	0.4 s

~5× faster queries · ~770× faster ingestion at near-equivalent recall.

### Agent memory · MemoryAgentBench (Ruler QA1 197K, gpt-4.1-mini)

Rank	System	Exact-Match
🥇	<b>SochDB V2</b>	<b>60.0%</b>
🥈	SochDB + HyDE	30.0%
🥉	GraphRAG	25.0%
—	Mem0 (managed)	5.0%

2× the previous best · 12× Mem0 on a public, peer-reviewed benchmark.

### Operational throughput

- 10.5K concurrent ops/sec sustained
- P99 latency < 2.2 ms
- Verified on commodity Kubernetes (3-node AKS, Standard\_D8s\_v5)

## 6 • Market — Where the Enterprise Dollars Are

Segment	2026 TAM	2030 TAM	CAGR
Vector database (enterprise)	\$2.2B	\$10.6B	22%
Enterprise database management	\$100B	\$180B+	12%
Agentic-AI infrastructure (Gartner)	\$4B	\$48B	65%
<b>Addressable wedge (intersection)</b>	<b>~\$8B</b>	<b>~\$35B</b>	—

**Beachhead ICPs (named accounts already in pipeline conversations):** - Tier-1 banks & insurers — RAG over policy + claims, audited retrieval - Healthcare payers / EHR platforms — HIPAA-aligned clinical agents - Retail & CPG (e.g., 7-Eleven, Walmart-tier) — store-ops & supply-chain agents - Federal / DoD primes — sovereign / air-gapped LLM deployments - LinkedIn-class platforms — agentic search & member-graph augmentation

## 7 · Why Enterprises Pick SochDB Over the Field

	Pinecone Enterprise	Databricks Vector	MongoDB Atlas Vector	Weaviate Enterprise	SochDB Enterprise
ACID across vectors + relational	✗	✗	⚠	✗	✓
Token-budgeted context query	✗	✗	✗	✗	✓
Graph memory built-in	✗	✗	✗	✗	✓
Postgres wire protocol	✗	✗	✗	✗	✓
BYOK / customer-managed keys	✓	✓	✓	✓	✓
Sovereign / on-prem / air-gapped	✗	✗	✗	⚠	✓
Fixed per-node pricing	✗	✗	✗	⚠	✓
Single vendor for retrieval stack	✗	⚠	⚠	✗	✓

**The wedge:** every other vendor sells *part* of the stack. SochDB sells the whole substrate, on the customer's infrastructure, under one contract.



## 8 • Business Model — Enterprise-First Pricing

Tier	Price	Target ACV
Professional	\$99 / node / month	\$30K – \$120K
Enterprise	From \$250K / year	\$250K – \$2M+
Sovereign / Air-gapped	Custom	\$1M – \$5M+

**Enterprise tier includes:** - Unlimited nodes, billion-vector capacity - 99.99% SLA, 4-hour P1 response - Dedicated CSM + solutions architect - SSO/SAML, BYOK, SOC 2 artifacts, custom DPA - Private VPC / on-prem / air-gapped deployment - Quarterly architecture reviews, prioritized roadmap input

**Marketplace metering:** \$0.05 / core-hour (Azure), \$0.15 / node-hour (AWS) — drives self-serve to PLG-converted enterprise deals.

**Revenue mix at scale (yr 3 plan):** 70% Enterprise · 20% Sovereign · 10% Pro / Marketplace.

## 9 · Go-To-Market — Top-Down Enterprise Motion

**Land** - Direct sales to AI Platform / Data Platform leaders at F500 - Cloud-marketplace co-sell (Azure, AWS, GCP) — buyers spend committed cloud credits - Design-partner program: 2 lighthouse customers per vertical (FinServ, Healthcare, Retail, Federal)

**Expand** - Per-workload land → enterprise-wide ELA in 12–18 months - Add-on SKUs: Sovereign deployment, Premium support, Custom compliance audits, Professional Services

**Channels** - Big-4 SI partnerships (Accenture, Deloitte, EY, KPMG) — implementation revenue share - Hyperscaler co-sell incentive programs (Microsoft Commercial Marketplace, AWS ISV Accelerate, GCP Partner Advantage) - Private-equity portfolio plays — standardize SochDB across portcos

**Sales motion** - 2 Enterprise AEs + 1 Sales Engineer in year 1 - Avg sales cycle: 4–6 months · Avg ACV: \$400K · NRR target: 130%

## 10 · Roadmap — Tied to Enterprise Revenue

Quarter	Milestone	Revenue unlock
Q3 2026	Distributed multi-node GA, replication, segment compaction	\$250K+ ACV deals
Q4 2026	Marketplace listings live (Azure AKS App, AWS, GCP); BYOK	Cloud co-sell pipeline
Q1 2027	<b>SOC 2 Type II</b> , mTLS/RBAC GA, audit-log SIEM connectors	FinServ + Healthcare
Q2 2027	Air-gapped / sovereign distribution, FedRAMP Ready prep	Federal / DoD
Q3 2027	Hosted SochDB Cloud (multi-tenant managed)	PLG → Enterprise pull
Q4 2027	HIPAA BAA, FedRAMP Moderate authorization in flight	Regulated verticals

## 11 · Founders

### Sushanth Reddy — Co-founder & CEO

- 18+ years in distributed systems and platform engineering
- Currently at **LinkedIn**, San Francisco Bay Area — shipping infrastructure at hyperscale
- Recognized internally as "most valuable tech asset" on Azure-scale infra programs
- Deep expertise in agentic AI, Semantic Kernel, Azure Event Grid, distributed messaging
- M.S., Jawaharlal Nehru Technological University
- Published author on AI agents, LLM systems, and quantum computing
- [linkedin.com/in/sushanthreddy](https://www.linkedin.com/in/sushanthreddy)

**Why he wins enterprise:** shipped infrastructure consumed by 1B+ LinkedIn members. Knows what it takes to operate, secure, and harden a data system at Fortune-500 scale.







### Sai Sandeep Kantareddy — Co-founder & CTO

- Senior Applied ML Engineer at **7-Eleven**, Austin TX — running production ML across a 13,000-store footprint
- 8+ years building production ML; led 20+ projects across autonomous vehicles, medical imaging, financial document AI
- Prior: ML research at **Bayer** (+10% accuracy on medical classification) and **NXP Semiconductors** (50% model-size reduction via quantization)
- M.S. AI / Medical Image Analytics, **Arizona State University** — 4.0 GPA
- **ACL 2026 Industry Track Reviewer** · **Antler Co-founder Club** member
- 9,000+ followers in the applied-ML community
- [linkedin.com/in/saisandeepkantareddy](https://www.linkedin.com/in/saisandeepkantareddy)

**Why he wins enterprise:** operates production ML in regulated (medical, financial doc) and high-stakes retail environments. Speaks the language of every regulated buyer SochDB needs to land.

**Combined edge:** hyperscale distributed-systems credibility + production-ML credibility in regulated verticals = the exact two halves an enterprise AI buyer interviews on.

## 12 · Traction & Validation

-  Production-grade benchmarks beating ChromaDB, LanceDB, Mem0, GraphRAG on public datasets
-  Postgres wire protocol + gRPC + gRPC-Web + WebSocket gateways operational
-  Helm chart, Prometheus, Grafana, OpenTelemetry shipped
-  Security baseline live: auth interceptor, rate limiting, mTLS, audit hooks
-  Active enterprise design-partner conversations underway (Retail F100, Healthcare payer, FinServ Tier-1)
-  Co-founder pedigree: LinkedIn + 7-Eleven — direct line into two F100 buyer organizations

## 13 · The Ask

**Raising \$5M Seed** to convert the technology lead into enterprise revenue.

Use of funds	%	Outcome
Enterprise engineering — distributed scale, multi-tenant, marketplace packaging	45%	Unlocks \$250K+ ACV deals
Security & compliance — <b>SOC 2 Type II</b> , HIPAA, FedRAMP prep, pen-tests	25%	Unlocks regulated verticals
GTM — 2 Enterprise AEs, 1 SE, 1 CSM, marketplace co-sell ops	25%	Builds repeatable sales motion
G&A — finance, legal (DPAs, MSAs), compliance program	5%	Enterprise-ready paperwork

### 18-month commitments to investors

- **5 Enterprise customers signed** at avg \$400K ACV → \$2M ARR exit run-rate
- **SOC 2 Type II** issued
- **3 cloud marketplace listings** live with active co-sell motion
- **2 Big-4 SI partnerships** signed
- **NRR ≥ 120%** on initial cohort

## 14 · Why SochDB Wins the Enterprise Category

The enterprise AI substrate market is consolidating in 2026–2027. The winning vendor will be the one that:

1. **Ships the whole retrieval stack under one contract** — vector + relational + memory + context, with ACID across all of them
2. **Speaks compliance fluently** — SOC 2, HIPAA, FedRAMP, EU AI Act, BYOK, audit-log export
3. **Deploys where the customer wants** — hyperscaler marketplace, private VPC, on-prem, air-gapped
4. **Has the systems chops to be trusted** — Rust, MVCC, SSI, distributed shards, Postgres-wire compatibility
5. **Beats the dedicated specialists on their own benchmarks** — already true today

*SochDB is the **enterprise system of record for agentic AI** — one platform, one vendor, one audit, billion-vector scale.*

## Contact

- Sushanth Reddy — [sushanth53@gmail.com](mailto:sushanth53@gmail.com)
- Sai Sandeep Kantareddy - [saisandeep.kantareddy@gmail.com](mailto:saisandeep.kantareddy@gmail.com)
- [github.com/sochdb/sochdb](https://github.com/sochdb/sochdb) · [sochdb.dev](https://sochdb.dev)