



swan_ai_labs

AIOS

AGENTIC OPERATING SYSTEM

Investor Opportunity

A small company that built what Fortune 500s haven't.
40+ production databases. Zero configuration. Real-time AI learning.
Competing head-to-head with Palantir, Databricks, and Snowflake.

CONFIDENTIAL | April 2026
Swan AI Labs | swancreate.com

The Problem



Enterprise AI Is Broken

The AI revolution is real. But for most enterprises, it's stuck. Not because the AI models aren't good enough — they are. The problem is **data access**.

The average Fortune 500 company has **200-500 databases** across departments, legacy systems, cloud platforms, and SaaS vendors. AI agents need cross-functional context to be useful — understanding that a customer in the CRM is the same entity as an account in finance and an order in operations. But getting that context requires:

\$1-5M

Cost per data source integration

6-18 mo

Deployment time per integration

80%

Of enterprise AI projects fail (Gartner)

The #1 reason cited for AI project failure? **Data integration**. Not model quality. Not compute costs. Data integration.

What Enterprises Pay Today for Fragmented Solutions

To give AI agents access to their data, enterprises cobble together separate tools that don't talk to each other:

- **\$150-500K/yr** — Data lineage tools (Collibra, Alation, Atlan) — maps data flows but can't query
- **\$50-200K/yr** — Knowledge graph databases (Neo4j, TigerGraph) — stores relationships but no AI agent integration
- **\$20-100K/yr** — Vector search (Pinecone, Weaviate) — stores embeddings but no schema discovery
- **\$200K-1M** — Custom agent orchestration (LangChain builds) — engineering time, not a product
- **\$300K-2M** — Custom AI feedback/learning pipelines — months of engineering, fragile
- **\$500K-2M** — Integration engineering (ETL, data warehouse, connectors) — ongoing maintenance burden

Total annual spend: \$1.2M - \$6M per enterprise

For tools that still don't integrate with each other, don't learn from corrections, and take months to deploy.

The Solution: AIOS



AIOS (Agentic Operating System) replaces all of these tools with a single integrated platform. Point it at your database server — SQL Server, PostgreSQL, or both — and within minutes:

40+

Databases auto-discovered & federated in production today

4,400+

Knowledge graph edges auto-built from foreign keys

7,000+

Schema snapshots tracked and versioned

500ms

Context assembly from 7 parallel sources

The Six Pillars of AIOS

- **Auto-Discovery** — connects to sys.databases, maps every table, column, primary key, and foreign key across all databases. Refreshes every 2 hours. Detects schema changes. Zero configuration required.
- **Federated Querying** — query any combination of databases in a single API call. Multi-step orchestration with result composition. Per-agent access control and database restrictions.
- **Semantic Search** — Azure AI-powered embeddings (text-embedding-ada-002) with HNSW vector search. Ask "which tables store customer payment data?" and get precise answers across 40+ databases in milliseconds.
- **Knowledge Graph** — auto-built entity relationship graph from schema discovery. Foreign keys become edges. Tables become nodes. 4,400+ weighted edges, traversable via API with configurable depth.
- **Data Lineage** — tracks how data flows between systems: QuickBooks → CRM, Stripe → Finance, EDI → Nutrition, 1WorldSync → Procurement. Registered flows with transform types and join conditions.
- **Real-Time Learning Loop** — the breakthrough. Agents log every decision with confidence scores. Humans submit corrections. Corrections are **automatically embedded into the vector store**. Future agent queries pull these corrections via semantic similarity. The system gets smarter with every interaction — without retraining any model.

The Context Assembly Pipeline

When an AI agent asks AIOS a question, the ContextBuilder runs **7 parallel lookups in under 500ms**: (1) agent profile and permissions, (2) last 20 conversation turns, (3) semantic search across schemas and knowledge, (4) last 5 decisions for consistency, (5) prior human corrections, (6) knowledge graph paths, (7) data lineage flows. All assembled into a single response. **No other product in the world does this.**

AI Agents Are Brand New — and They Change Everything

It's easy to forget how new this is. ChatGPT launched in November 2022. GPT-4 shipped in March 2023. AI agents — autonomous systems that can reason, use tools, and take actions — have only been **practically possible for about two years**. The entire AI agent ecosystem is being built right now, in real time.

This is not an incremental improvement to existing technology. This is a **paradigm shift** as fundamental as the shift from mainframes to personal computers, or from on-premise to cloud. Every enterprise in the world will deploy AI agents within the next 3-5 years. Gartner predicts that by 2028, **33% of enterprise software applications will include agentic AI**, up from less than 1% in 2024.

But Agents Are Useless Without Data Infrastructure

An AI agent without access to enterprise data is like a brilliant employee on their first day — smart but useless. They don't know which databases exist, what tables contain, how systems relate to each other, or what mistakes were made last week. Today, connecting an agent to enterprise data requires:

- Months of custom integration engineering per data source
- Manual schema mapping that breaks when databases change
- No systematic way to teach agents from their mistakes
- Separate tools for search, graphs, lineage — none integrated
- \$1-6M per enterprise, deployed over 6-18 months

This is like the early days of cloud computing, when every company built its own server infrastructure. Then AWS arrived and said: **"We'll handle the infrastructure. You build your applications."** That shift created \$2 trillion in value.

AIOS Is the AWS Moment for AI Agents

AIOS says: **"We'll handle the data infrastructure. You deploy your agents."** Point AIOS at your databases. In minutes — not months — your agents can discover schemas, query across databases, search semantically, traverse knowledge graphs, track data lineage, and learn from every correction. This is what "infrastructure for AI agents" means. And no one else has built it.

We are in the first 24 months of a 20-year transformation.

AI agents will become as ubiquitous as cloud computing. The infrastructure layer between agents and data is being built right now. AIOS is that infrastructure. First movers in infrastructure capture generational value.

Why This Hasn't Been Built Before

Three things had to be true simultaneously for AIOS to exist:

- **LLMs had to be good enough** — only since GPT-4 (2023) can agents reason about complex data relationships
- **Vector search had to be fast enough** — HNSW algorithms and specialized hardware only matured in 2023-2024
- **Someone had to solve the whole problem at once** — big companies solve pieces (a catalog here, a vector DB there). AIOS was built by a team that needed all seven capabilities working together from day one. Necessity drove integration.

Why AIOS Is Revolutionary



Every week, hundreds of AI startups launch. Most are wrappers around ChatGPT. Some build interesting tools. Almost none build **infrastructure**. Here is what makes AIOS fundamentally different from everything else in the market:

1. It's an Operating System, Not an Application

Applications (chatbots, copilots, assistants) compete with thousands of other applications. Infrastructure — the platform everything else runs on — has **winner-take-most dynamics**. AWS, Azure, Snowflake, Stripe, and Twilio are infrastructure companies. They capture disproportionate value because everyone builds on top of them. AIOS is infrastructure for AI agents. The company that owns this layer captures outsized value as every enterprise deploys agents.

2. The Learning Loop Creates a Compounding Moat

Most AI systems are stateless — every conversation starts from zero. AIOS implements a **closed-loop learning system** that compounds value over time:

Agent makes decision → Human corrects → Correction auto-embeds into vector store → Future queries pull correction → Agent never makes that mistake again

This is what Fortune 500 companies spend \$300K-2M trying to build as custom RLHF pipelines. AIOS does it out of the box, in real time, per correction, without model retraining. Every correction makes every future agent query better. A competitor starting from zero faces a cold-start problem that AIOS has already solved.

3. 40+ Databases Federated — Not a Demo

This is not a prototype or a slide deck. AIOS is in production today, federating across **40+ enterprise databases** spanning finance (10 databases), CRM (2), audit (2), food & nutrition (4), AI & technology (2), and specialized systems (6+). Every table, column, primary key, and foreign key mapped automatically. Most enterprises can't get two departments to share a spreadsheet. AIOS queries across their entire data estate in a single API call.

4. Three AI Paradigms Unified in One Platform

The industry treats these as separate products requiring separate vendors:

Structured Data

Relational schemas, auto-discovered from INFORMATION_SCHEMA.
Tables, columns, FKs, PKs indexed and versioned.

Unstructured/Semantic

Vector embeddings via Azure AI. Cosine similarity search with HNSW indexing. Natural language to data.

Relational Knowledge

Knowledge graph with weighted edges. Auto-built from foreign keys.
DFS traversal with depth control.

All Three, Unified

The ContextBuilder assembles all three in parallel — in 500ms. This architecture costs \$5-10M when consulting firms design it.

5. Built by a Small Team That Outperformed \$25B in R&D

Microsoft, Palantir, Databricks, and Snowflake spend **\$25B+ per year combined** on research and development. Not one of them has built what AIOS does — not because they can't, but because it requires integrating seven complex subsystems that span different teams, products, and business units inside those companies. Swan AI Labs built it as one coherent platform from day one.

Head-to-Head: AIOS vs. The Giants



AIOS doesn't compete in the margins. It competes directly with the largest AI and data infrastructure companies in the world — and wins on integration, speed-to-value, and price.

Capability	AIOS	Palantir AIP	Databricks	Snowflake	LangChain	Collibra
Auto Schema Discovery	Yes, zero-config	Requires setup	Unity Catalog	Snowflake only	No	Manual
Cross-DB Federation	40+ DBs, 1 API call	Custom ETL	Lakehouse only	Snowflake only	No	No
Knowledge Graph	Auto-built from FKs	Manual	No	No	No	Manual catalog
Vector/Semantic Search	Azure AI HNSW	Limited	Vector Search	Cortex	Via plugins	No
Data Lineage	Built-in	Yes	Unity Catalog	Limited	No	Yes (\$500K)
Agent Learning Loop	Real-time, auto-embed	No	No	No	No	No
Self-Service Onboard	5 minutes	6+ months	Weeks	Weeks	Months (eng)	Months
Gov Cloud / FedRAMP	Ready now	Yes (IL5)	In progress	FedRAMP'd	No	No
Typical Cost	\$250K-2M/yr	\$5-50M/yr	\$100K-1M/yr	\$100K-500K/yr	\$200K-1M eng	\$150-500K/yr

Why We Win

- **vs. Palantir** — Same ambition, 10-50x cheaper, self-service instead of 6-month deployments. Palantir can't serve mid-market. We can serve everyone from small business to federal government.
- **vs. Databricks** — Databricks requires moving data into their lakehouse. AIOS works with data where it already lives — no migration, no ETL, no data warehouse.
- **vs. Snowflake** — Snowflake only federates within Snowflake. AIOS federates across SQL Server, PostgreSQL, and any future connector — across 40+ databases simultaneously.
- **vs. LangChain** — LangChain is a library. Every customer builds from scratch. AIOS is a deployed platform with auth, multi-tenancy, billing, and a learning loop.
- **vs. Collibra/Alation** — They catalog data but can't query it, can't learn from it, and can't serve AI agents. AIOS does all three.

AIOS delivers Palantir-grade capabilities at 10-50x lower cost, deployable in minutes instead of months.

No other product combines all six pillars. That's not a feature gap — it's a category gap.

\$62 Billion Total Addressable Market by 2030

AIOS sits at the intersection of three massive, high-growth markets. No existing product spans all three.

\$47B

AI Agent Platforms by 2030 — 38% CAGR (Grand View Research)

\$12B

Database Middleware & Integration by 2028 — 10% CAGR

\$3.2B

Data Catalog & Governance by 2028 — 19% CAGR

Target Customer Segments

Fortune 500 / Enterprise (1,500 companies)

200-500 databases per company

IT budget: \$50M-500M/yr

AI/ML spending growing 25%+ YoY

Contract: \$500K-5M/yr

Government / Defense

Federal AI budget: \$15B+ and growing

FedRAMP required (AIOS ready)

FIPS 140-2 required (AIOS compliant)

Contract: \$1M-10M/yr

Mid-Market (50,000 companies)

10-50 databases per company

IT budget: \$1M-50M/yr

Need AI but lack data engineering teams

Contract: \$100-500K/yr

Healthcare / Financial Services

Heavy compliance requirements

HIPAA, SOX, PCI-DSS

Audit logging + encryption built in

Contract: \$250K-2M/yr

Government: A \$15B Moat Within a Moat

AIOS is already Gov Cloud ready: Azure Government regions, FIPS 140-2 compliant crypto (PBKDF2, AES-256-GCM), Azure Key Vault secret management, compliance middleware with full audit logging, and HTTPS enforcement. FedRAMP authorization — which most AI startups are 12-18 months away from — opens the **\$15B+ federal AI market** where competition is minimal and contracts are multi-year.

Business Model



Two Channels: Enterprise Platform + SMB SaaS

AIOS serves the full market — enterprise platform contracts for Fortune 500 and government, and a self-service SaaS tier for small businesses and developers that feeds the enterprise pipeline.

Enterprise Platform (Direct Sales)

Business	Enterprise	Government
\$250K/yr 50 databases Unlimited agents Managed or self-hosted Priority support	\$500K-2M/yr Unlimited everything Self-hosted / VPC SSO + SLA Dedicated CSM	\$1M-5M/yr FedRAMP authorized Azure Gov / IL5 FIPS 140-2 crypto On-premise option

Small Business SaaS (Self-Service)

The SaaS tier serves small businesses, startups, and developers. It creates product adoption, community, and a pipeline for enterprise upsell — the same model that made Slack, Atlassian, and Datadog billion-dollar companies.

Free	Starter	Pro
\$0/mo 1 database, 1 agent 1K API calls Developer adoption	\$299/mo 5 databases, 5 agents 50K API calls Small business	\$999/mo 25 databases, 25 agents 500K API calls Growing companies

Revenue Projections

\$1M Year 1 — 3 enterprise + 2 mid-market	\$8M Year 2 — 12 enterprise + 10 mid-market + 2 gov	\$30M Year 3 — 30 enterprise + 25 mid-market + 5 gov
\$100-200M Year 5 — with enterprise + government expansion	85-90% Gross margin on platform licenses	

Why Enterprise, Not SaaS

- **Higher ACV** — \$250K-\$5M per customer vs. \$12K/yr SaaS. Fewer customers needed for significant revenue.
- **Stickier** — enterprise deployments with federated data become mission-critical. Switching costs are enormous.
- **Government revenue** — \$1-10M multi-year contracts with built-in renewal. SaaS can't serve government.
- **Matches the value** — customers pay \$1.2-6M today for fragmented tools. \$250K-\$2M for an integrated platform is a no-brainer.
- **Professional services** — onboarding, training, custom connectors add 20-30% to license revenue.

Developer free tier exists for adoption — developers discover AIOS, build POCs, and bring it to their enterprises. The free tier feeds the enterprise pipeline.

What's Already Built



This is not a pitch for something we plan to build. **AIOS is built, deployed, and running in production today.** Here is exactly what exists:

Platform Architecture

Component	What It Does	Status
Schema Discovery	Auto-discovers all DBs, tables, columns, FKs, PKs every 2 hours	Production
Federated Query Engine	Query 40+ databases in single API call with access control	Production
Knowledge Graph	4,400+ auto-built edges with DFS traversal and weighted relationships	Production
Vector Store	1,500+ embeddings, Azure AI Search backend (HNSW), MSSQL persistence	Production
Embedding Pipeline	Azure AI (ada-002), batch processing, auto-indexing of schemas/feedback	Production
Context Builder	7 parallel source assembly in <500ms	Production
Learning Loop	Decision logging, human corrections, auto-embedding into vector store	Production
Multi-Tenancy	Tenant isolation, per-tenant quotas, usage tracking, API keys	Production
Stripe Billing	Checkout, subscriptions, webhooks, auto plan upgrade/downgrade	Production
Self-Service Onboarding	Sign up → connect database → querying in 5 minutes	Production
Client SDK	@aios/client — JavaScript/TypeScript, zero dependencies, full API	Production
Gov Cloud	FIPS 140-2, Azure Key Vault, compliance middleware, audit log	Production
PostgreSQL Support	Full connector: discover, schema map, parameterized queries	Production
Infrastructure	Docker, CI/CD, 63 automated tests, Swagger UI, database migrations	Production

60+

API endpoints across v1 and v2

22

Database tables in production

63

Automated tests, all passing

15

Commits of production-quality code

The Investment Opportunity



What You're Investing In

A **fully built, production-deployed AI infrastructure platform** that does what Microsoft (\$3T), Palantir (\$80B), Databricks (\$43B), and Snowflake (\$50B) have not yet achieved — built by a small team at a fraction of the cost. The technology risk is zero. The product exists. The market is massive. The competition hasn't figured it out yet.

What Investment Accelerates

Go-to-Market (\$1.5M)

3 enterprise account executives
Developer advocacy & technical content
Conference presence (AI Engineer, KubeCon)
First 5 enterprise deals in 6 months

Engineering (\$1.5M)

4 senior engineers
MySQL, Oracle, MongoDB connectors
Real-time streaming (Kafka, EventHub)
Horizontal scaling (Kubernetes)

Compliance (\$1M)

FedRAMP authorization (3-6 months)
SOC2 Type II certification
HIPAA compliance for healthcare
Opens \$15B+ federal AI market

Operations (\$1M)

Product manager
Customer success team
24/7 support infrastructure
Multi-region deployment

\$5M investment → \$8M Year 2 ARR → \$30M Year 3 ARR

The platform is built. Investment accelerates time-to-revenue and locks in first-mover advantage in a \$62B market that no one else has figured out how to serve.

Why Now

- **AI agent adoption is accelerating** — every enterprise is deploying agents in 2026-2027. They all need data infrastructure.
- **First-mover advantage compounds** — every customer's data makes the learning loop stronger. Early customers create an insurmountable lead.
- **Government window is open** — FedRAMP-authorized AI agent platforms can be counted on one hand. Getting in early locks in multi-year contracts.
- **The product is done** — this is not a "build it" investment. This is a "sell it" investment. The hard part is finished.



Let's Talk

For investment inquiries and
partnership opportunities:

Swan AI Labs

swancreate.com

aios.swancreate.com

aios.swancreate.com/api/docs

CONFIDENTIAL — This document contains proprietary information.
Distribution without written authorization is prohibited.

April 2026