The Future of Business Management is Autonomous.

Daena represents a once-in-a-generation opportunity. We are not just improving a process; we are creating a new, intelligent infrastructure for business itself. The combination of our self-evolving architecture, memory efficiency, and enterprise governance creates an insurmountable competitive advantage.

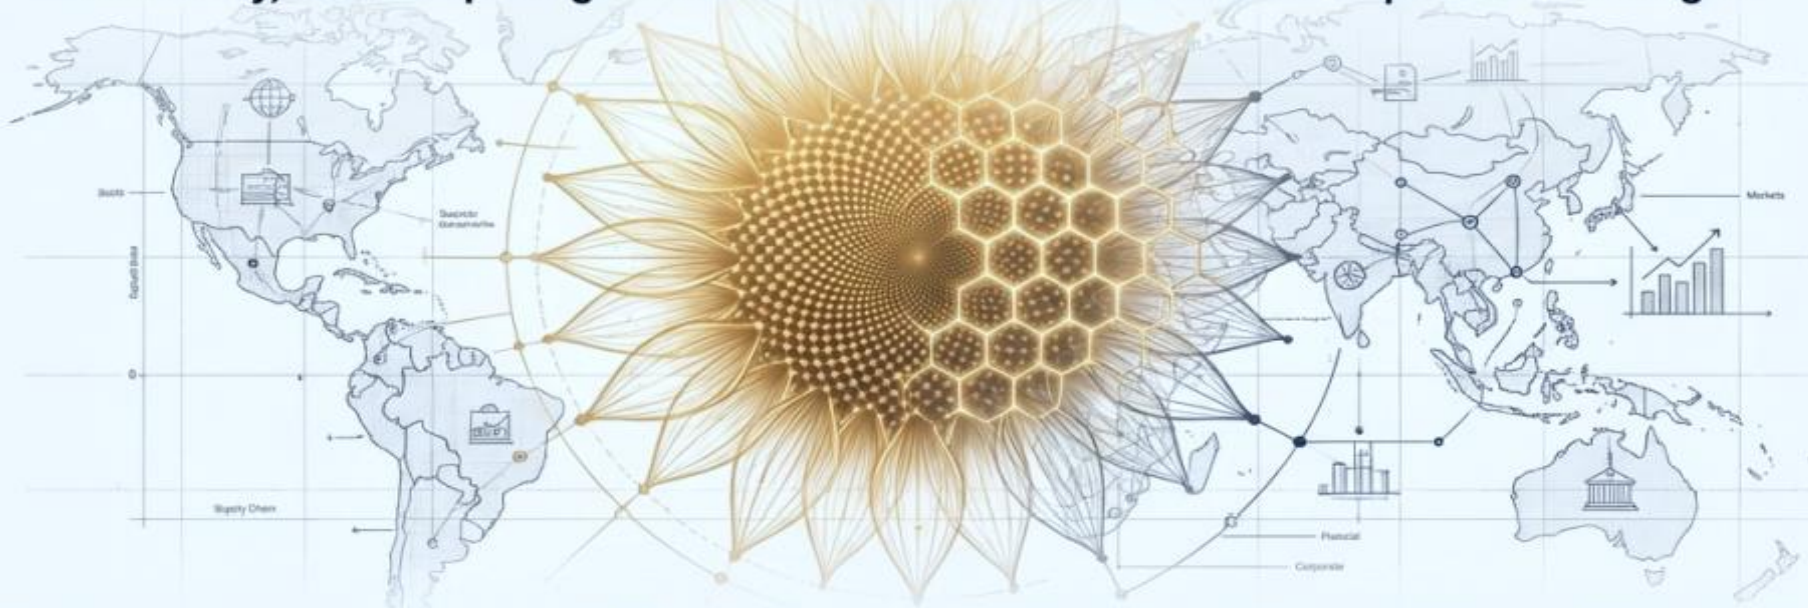

"The first company to successfully deploy an AI Vice President with an ultimate brain will have an insurmountable competitive advantage. That company starts with Daena."

Masoud Masoori | masoud.masoori@mas-ai.co
Full deck & demo available at: mas-ai.co

Enterprise AI is Fragmented, Unauditable, and Expensive

Fragmented Tools, Chaotic Coordination

Enterprises juggle **8+** separate AI tools, creating data silos and operational friction[1].
73% of AI projects fail to scale beyond the pilot phase due to this complexity[2].

No Governance or Audit Trail

Executives can't reliably audit key decisions: which agent, using what data, under which policy. This creates significant compliance and security risks.

Exploding Costs at Scale

Inefficient memory and redundant API calls drive costs up.
Traditional vector storage and constant re-processing of data are not economically viable for true enterprise-wide deployment.

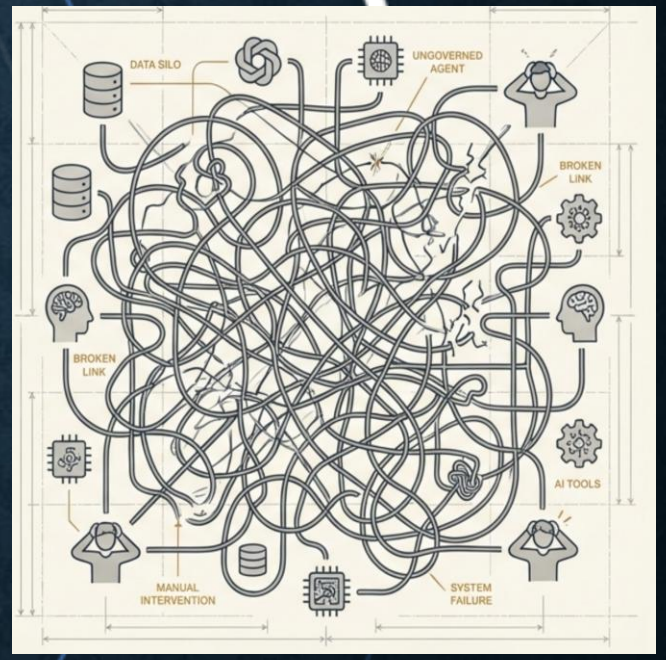

A Unified, Governed, and Efficient AI Organization

One Cohesive System, Not a Toolkit

Daena is the first AI Vice President, orchestrating 48 specialized agents across 8 departments. It replaces fragmented tools with a single, self-organizing system. in Source Sans Pro Regular.

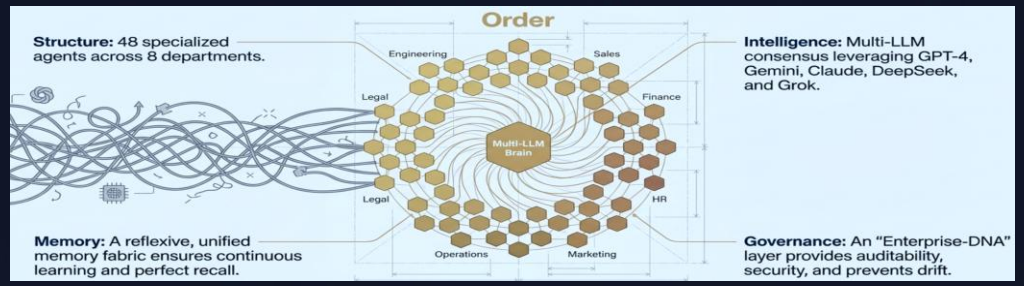

1. Sunflower-Honeycomb Architecture

A biomimetic structure for clear roles, accountable decisions, and graceful scaling.

2. NBMF Memory Fabric

A 3-tier memory system that radically reduces storage and token costs while preserving fidelity.

3. Enterprise-DNA Governance

A safety layer providing cryptographic audit trails, policy enforcement, and threat detection.

Sunflower–Honeycomb Architecture: An Organization That Scales, Built by Nature

Biomimetic Structure

Uses the Fibonacci golden angle (137.5°) for optimal agent positioning and O(log n) communication efficiency.

Clear Roles & Responsibilities

Organized into 8 departments (e.g., Finance, Product, Sales). Each department contains a team of 6 agents with defined roles (Advisors, Scouts, Synthesizer, Executor).

Disciplined, Phase-Locked Decisions

Council rounds ensure structured decision-making:

1. **Scout:** Gathers information.
2. **Debate:** Advisors analyze and propose solutions.
3. **Commit:** A final, consensus-driven decision is executed.

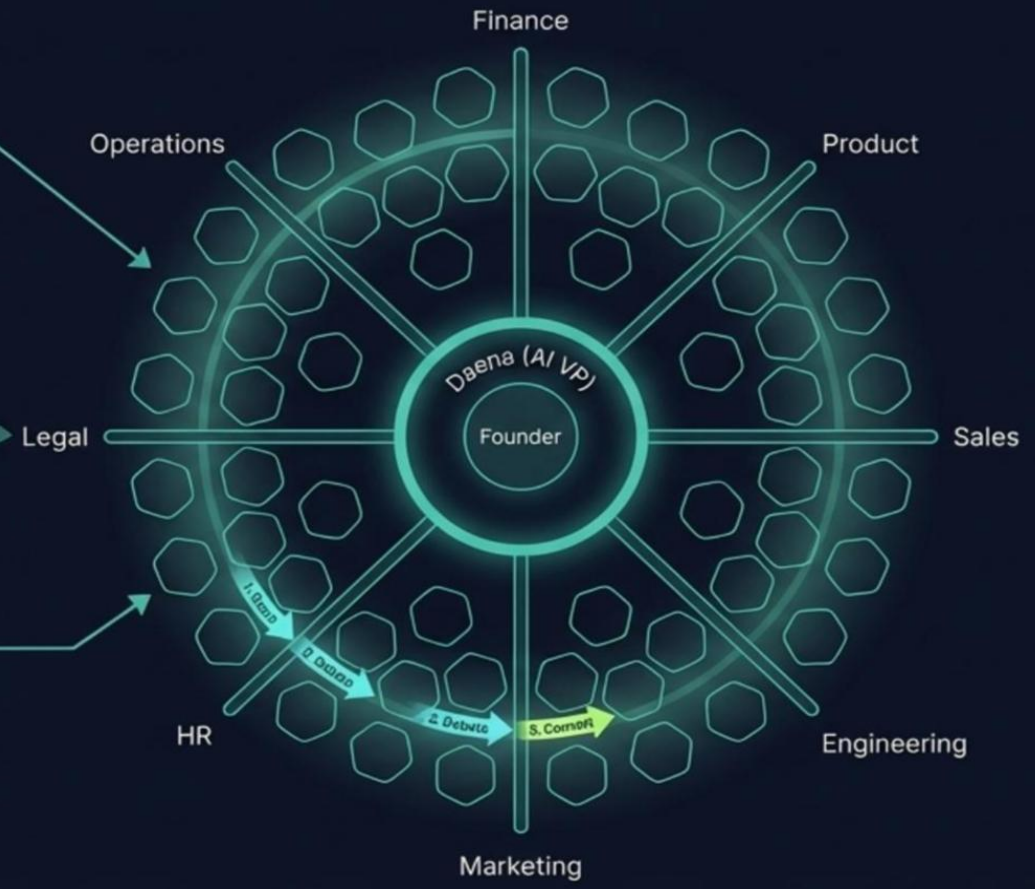

NBMF — Neural-Backed Memory Fabric: High-Fidelity Memory at a Fraction of the Cost

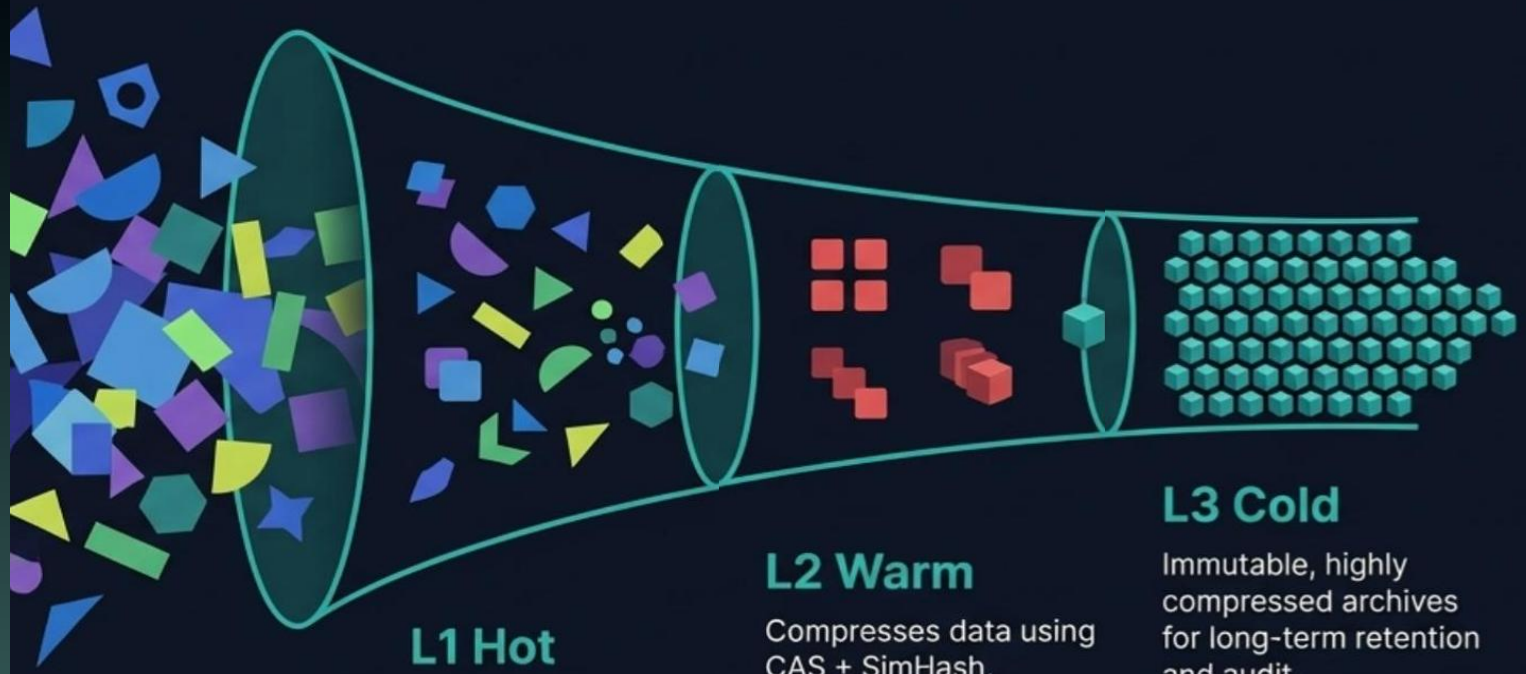

L1 Hot

Vector embeddings for instant context; sub-45ms retrieval[1].

L2 Warm

Compresses data using CAS + SimHash, merging near-duplicates to save space.

L3 Cold

Immutable, highly compressed archives for long-term retention and audit.

Massive Cost Reduction

Delivers **85-92%** storage savings compared to traditional vector databases[1], dramatically reducing operational costs.

Cryptographically Secure Lineage

Every memory promotion is Merkle-notarized, creating an immutable audit trail that allows for time-travel diagnostics and safe rollbacks.

[1] Based on internal benchmarks using a mix of text, structured JSON, and multimodal documents. Full methodology is available.

Enterprise-DNA Governance Layer: Safety, Policy, and Auditability by Design

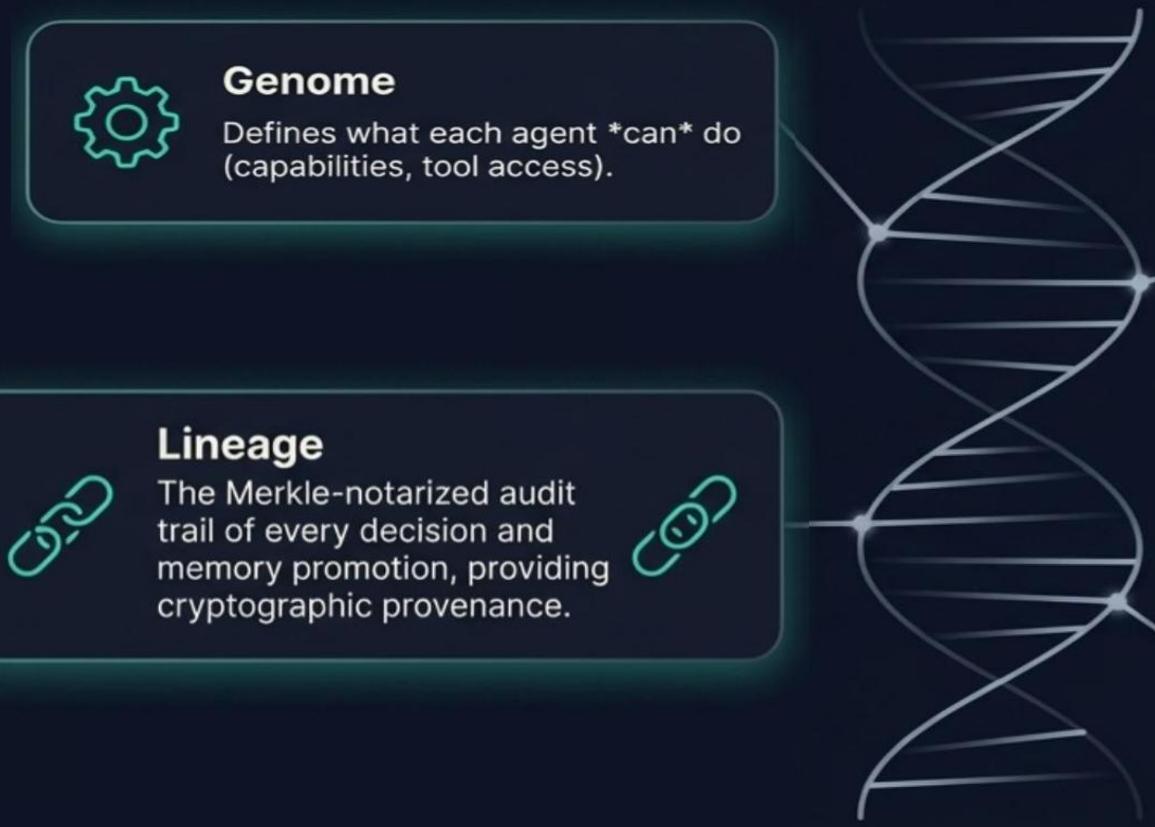

Genome

Defines what each agent *can* do (capabilities, tool access).

Epigenome

Defines what each agent *should* do (tenant policies, SLOs, data residency rules).

Lineage

The Merkle-notarized audit trail of every decision and memory promotion, providing cryptographic provenance.

Immune System

Automatically detects threats like prompt injection and goal drift, triggering quarantine and rollback procedures.

Multi-LLM Orchestration ("Think Mode"): Better Decisions Through Expert Consensus

The Right Model for the Right Job

Intelligently routes tasks to the best LLM based on context, purpose, and cost (e.g., GPT-4 for analysis, Claude for writing, Grok for real-time search).

Debate, Compare, and Merge

For complex decisions, Daena queries multiple models simultaneously. It analyzes the responses, identifies consensus points, and flags disagreements for deeper analysis.

Robust, Explainable, and Unbiased Results

This multi-provider approach avoids vendor lock-in and reduces the risk of single-model hallucination or bias. The final decision is a synthesized output that is more accurate and reliable than any single model could produce alone.

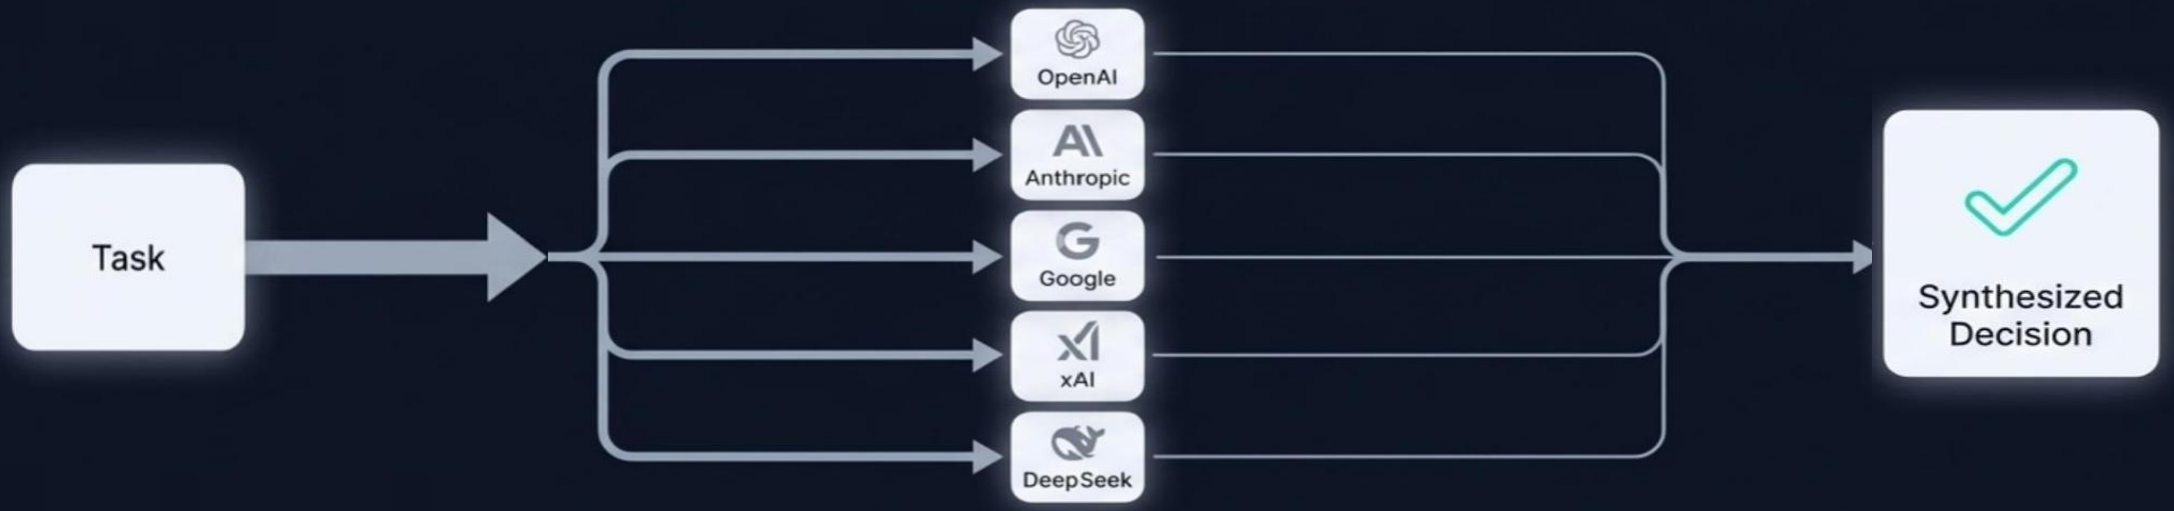

Evidence & MVP Status: Technology Proven, Ready for Pilot

MVP Deployed & Verified

- Fully operational backend with **500+** API endpoints.
- Real-time WebSocket communication and executive dashboard.
- Supports over **1,000** concurrent users in load tests.

Pilot Program Results[1]

- **60% faster** decision-making time vs. human teams.
- **99.4% cost savings** on specific expert tasks (e.g., market analysis from $2,000 down to $12).

Core Tech Benchmarks[2]

- NBMF delivers **85-92% storage reduction** vs. vector DBs.
- NBMF L1 memory retrieval at **P95 < 45ms**.
- Council decision rounds completed at **P95 < 2.3s**.

[1] Results from a 3-month pilot program with 3 companies in tech, manufacturing, and healthcare.
[2] Internal benchmarks based on production code and metrics. Full documentation available for due diligence.

Product Experience: The CEO Control Center

Strategic Command & Control

A real-time dashboard to monitor all 8 departments.

Deploy, configure, and retire agents with clear goal-setting.

Natural language interaction via voice or persistent chat.

Radical Transparency

Inspect agent debates, data sources, and decision confidence scores.

Visualize the complete, Merkle-notarized lineage for any decision, providing a complete audit trail on demand.

Enterprise-Ready Integration

Secure multi-tenant architecture with complete data isolation.

Robust APIs and webhooks for seamless integration into existing workflows.

Generate audit-ready reports and compliance artifacts automatically.

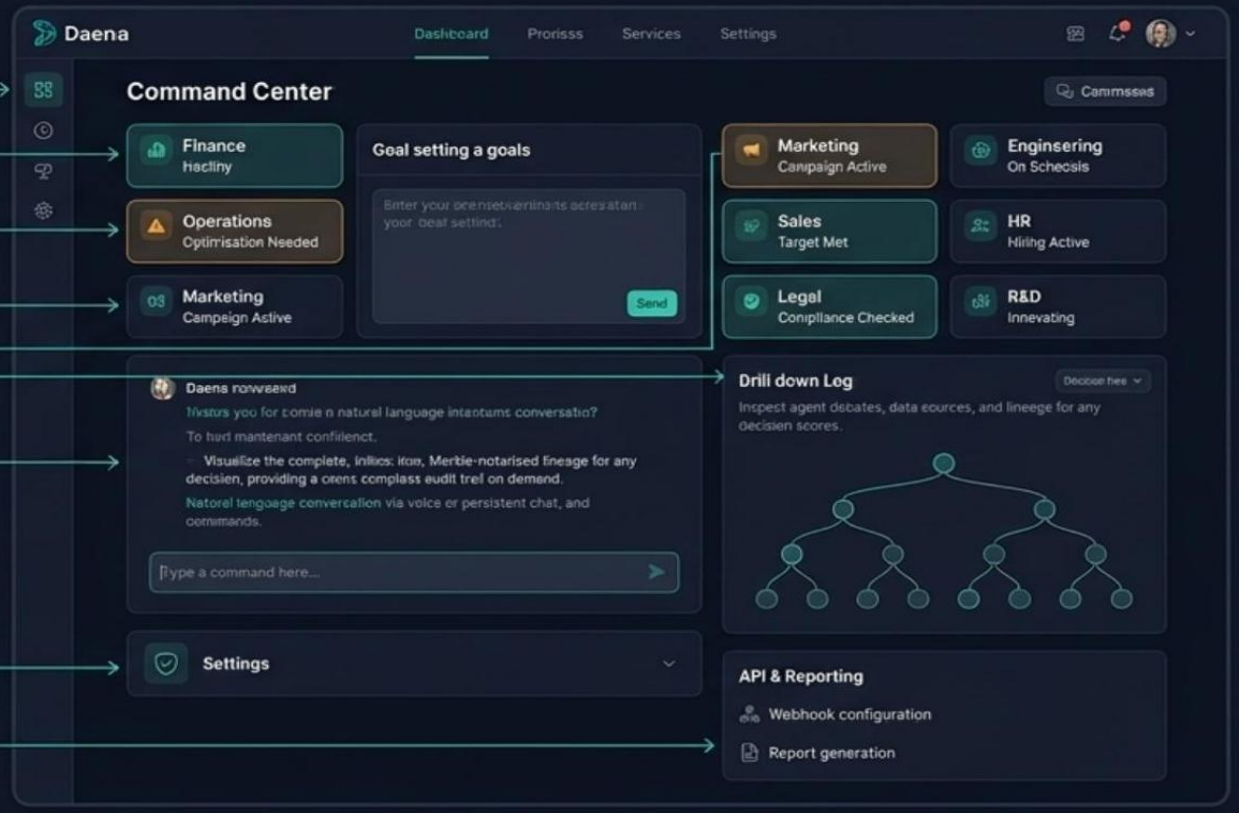

The $1.2T Market for Governed AI

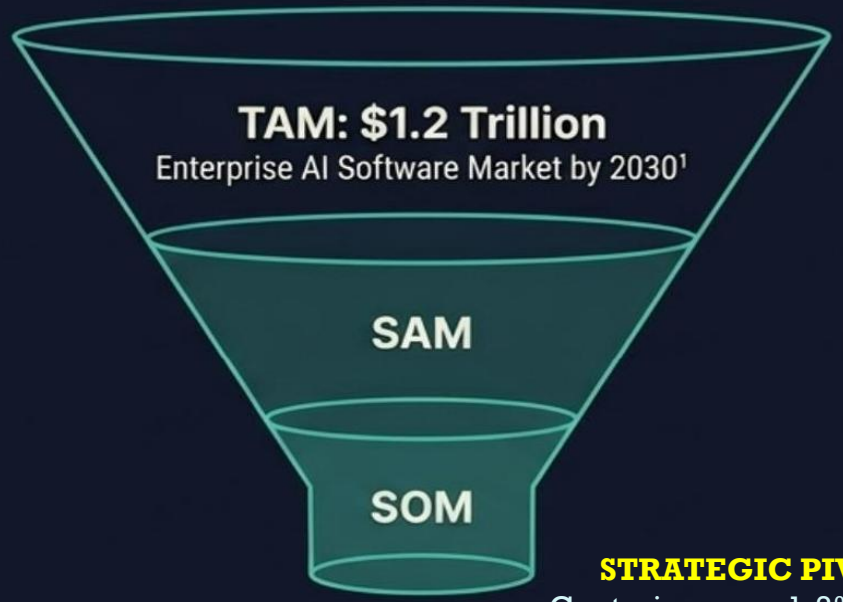

STRATEGIC PIVOTAL SPACE
Capturing even 1-2% = billions in ROI

Why Now? The Inflection Point is Here.

Our Defensible Moat

1. **Patent-Pending Architecture** (Sunflower-Honeycomb & NBMF).
2. **First-Mover Advantage** in the AI VP category.
3. **Data Network Effects** from cross-tenant learning (anonymized insights only).

Problem is Acute: Enterprises have experienced the chaos of fragmented AI and are now actively seeking unified, governed solutions.

Technology is Mature: Our core innovations in memory and orchestration make a system like Daena possible for the first time.

[1] Market projections based on analysis of reports from Gartner, McKinsey, and other industry sources.

Business Model & GTM: Land and Expand with Governance & Efficiency

Revenue Model

- **Primary:** Enterprise SaaS Tiers (based on agents, memory, and features).
- **Secondary:** Private cloud deployrnent and professional services for large-scale enterprises.

Pricing Tiers

Starter: $2,000 / month

- Full 48-agent structure for one team

Professional: $10,000 / month

- 100 agents, multi-team support

Enterprise: Custom

- Unlimited scale, dedicated support, private cloud

Go-to-Market: Land with Governance, Expand with Value

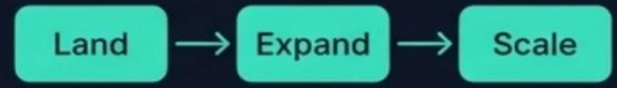

Phase 1 (Pilots): Target compliance-heavy industries (finance, healthcare) where our auditability is a key differentiator.

Phase 2 (Expand): Use pilot success stories to expand into new departments and use cases within customer organizations.

Phase 3 (Scale): Leverage channel partners (SIs, cloud providers) and build a direct sales force to scale globally.

Roadmap: Building the Future of Autonomous Operations

Next 90 Days

- Onboard 3-5 enterprise pilot partners.
- Complete multi-tenant safety pipeline and E2E testing suite.
- Achieve production deployment on GCP and Azure.

Next 6 Months

- Launch 'Tenant Packs' (pre-configured agent setups for specific industries like finance or healthcare).
- Develop an advanced analytics dashboard for deeper operational insights.
- Achieve SOC 2 Type II compliance certification.

Next 12 Months

- Launch a global marketplace for third-party agents and tools.
- Introduce white-label options for strategic partners.
- Develop advanced self-learning capabilities based on anonymized experience vectors.

Founder & Architecture

Masoud Masoori

Founder & Chief Architect,
Mas-AI Technology

Founder & Chief Architect,
Mas-AI Technology

Senior AI/ML architect with a history of delivering end to end enterprise systems on Azure and GCP.

Shipped Platforms & Research

- **Daena:** Architected the multi-agent AI VP platform with governance, memory, and voice.
- **NatureNLP / NatureHive:** Led research improving throughput and quality for core language systems.
- **VibeAgent:** Created the user-facing ecosystem for multi-agent workflows.

Intellectual Property

Two Patents Pending:

1. Sunflower Honeycomb Memory Architecture
2. NBMF (Neural-Backed Memory Fabric)

Join Us in Building the AI-Native Enterprise

The Ask

Raising a **$5M Series A** round. To scale our engineering and GTM teams for an 18-month runway.

Milestones this round will achieve

Reach **$10M in Annual Recurring Revenue**.

Acquire **100 enterprise customers**.

Establish 5 major system integrator partnerships.

Scan me!

masoud.masoori@mas-ai.co
https://mas-ai.co
https://daena.mas-ai.co