End-to-End Delivery Plan

v1.0 · April 2026 · gsiso.ai

The execution manual behind the GTM roadmap.
Sprints, hires, budget, and decision gates for the first 18 months.
Written for Gaurav Sisodia and the founding team.

SEED RAISE	TEAM	TARGET
$15M	14 hires	$10M ARR · M18

Gaurav Sisodia, gsiso.ai · Confidential — Not for distribution

gsiso.ai — End-to-End Delivery Plan

v1.0 · April 2026 · Author: Gaurav Sisodia · Confidential

This document is the build plan, not the go-to-market plan. The Roadmap & GTM brief describes what gsiso will ship and why the market exists. This plan describes how we ship it: which engineers, which sprints, which deliverables, what we spend, and the exact conditions under which we proceed, pivot, or sell.

§	Section	Summary
§1	Executive Summary	Wedge, 18-month outcomes, team, seed raise
§2	Tech Stack (Locked)	Full layer table + build-vs-buy philosophy
§3	Phased Roadmap — 18 Months	Five phases with metrics, deliverables, team size
§4	Phase 1 Sprint Plan	8 sprints × 2 weeks, Physical AI Bridge MVP
§5	Team Plan	14 hires in priority order with compensation bands
§6	18-Month Budget	Line-item budget, monthly burn curve
§7	Design Partners — Target List	Three anchors across pharma, manufacturing, capital markets
§8	Risk Register — Top 10	Likelihood, impact, mitigation for each risk
§9	Decision Gates	Go/no-go at M6, M12, M18 with pivot options
§10	What We Explicitly Don't Build	Scope boundaries and rationale
§11	Closing: The 18-Month Bet	The precise thesis statement

§1 Executive Summary

The Wedge

gsiso.ai enters on two dateable, defensible differentiators:

- Physical AI Bridge — a native ROS 2 + MCP + VLA orchestration layer for humanoid robots, cobots, and lab automation. No hyperscaler, no open-source framework, and no orchestration startup ships this as of April 2026. The gap is confirmed by the State of Robotics 2026 report: the software bridge between LLM agents and physical systems is the largest under-served layer in the entire AI stack.
- EU AI Act Compliance Certification — first orchestration fabric to achieve CE marking for high-risk AI orchestration. EU AI Act enforcement for high-risk AI begins August 2026. Compliance adds $8–$15M to each large enterprise implementation. First-mover certification compounds with every month of lead time.

18-Month Outcomes (May 2026 → October 2027)

Outcome	Target	Gate
Design partners in production	3 (pharma, manufacturing, capital markets)	M12
First paid enterprise contract	≥$500K ACV	M14
ARR milestone	$10M in signed contracts	M18
Series A close	$30–50M	M18
Team size	14 hires	M14
Seed raise	$15M (2-month buffer past M18)	M0

Why This Plan Is Different from the Roadmap Doc

The Roadmap doc sets quarterly milestones at the GTM level: design partners signed, physical bridge GA, CE marking. This plan drills to execution: sprint-by-sprint deliverables, week-2 demo scripts, compensation bands, line-item budget, and explicit go/no-go gates with pivot options. Anyone reading both should be able to pick up a Jira board and start tomorrow.

§2 Tech Stack (Locked)

All choices below are locked for Phase 1 and Phase 2. Re-evaluation points are noted where applicable.

Layer	Choice	Rationale	Build or Buy
Agent runtime / orchestration	Python 3.12 + FastAPI 0.115 + Ray 2.x + LiteLLM gateway	Ray for distributed agent scheduling at scale; LiteLLM for multi-model routing across OpenAI, Anthropic, Gemini, local vLLM	Build: scheduler, policy enforcement, agent lifecycle Buy/integrate: LiteLLM, vLLM, Ray cluster
Agent memory & state	PostgreSQL 16 + pgvector 0.7 + Redis 7 + S3 Drizzle ORM	Postgres+pgvector for long-term semantic memory (tenant-isolated); Redis for hot agent state; S3 for artifacts + audit cold archive	Buy all components. Build: schema design, tenant isolation, memory namespace policy
Identity & cryptography	Ed25519 per-agent keys + DID:web + Sigstore/Cosign + HashiCorp Vault	Ed25519 for DID keypairs; DID:web for discoverable identities; Cosign for tamper-evident receipts; Vault for key lifecycle management	Build: DID minting, receipt chaining, kill-switch propagation Buy: Vault, Cosign
Policy engine	Open Policy Agent (OPA) + Rego + custom Policy Studio UI	OPA provides WASM-compilable, version-controlled policy-as-code. Policy Studio is the commercial moat on top. Evaluate Cedar (AWS-native) in Phase 2 if required by customer	Build: Policy Studio UI, Rego authoring interface, policy compilation pipeline Buy: OPA runtime
Physical AI Bridge	ROS 2 Jazzy + WebRTC + MCP servers per robot class + NVIDIA Isaac Sim	ROS 2 Jazzy (LTS through 2027) as robot primitive layer; WebRTC for real-time telemetry; custom MCP servers for humanoid, cobot, AMR, lab-arm; Isaac Sim for digital twin gate	Build: all MCP server adapters, sim gate pipeline, WebRTC ingest, safe-stop proof primitive Buy: ROS 2, Isaac Sim license, hardware
Control plane (console)	Next.js 15 + React 19 + Tailwind CSS 4 + shadcn/ui + Clerk + Stripe + tRPC	Next.js App Router for console; shadcn/ui for accessible primitives; Clerk for enterprise IdP auth; Stripe for billing metering; tRPC for type-safe APIs	Build: console product, Policy Studio, agent dashboard, fleet telemetry views Buy: Clerk, Stripe, tRPC, shadcn
Observability	OpenTelemetry + ClickHouse + Grafana + Sentry	OpenTelemetry for instrumentation; ClickHouse for agent event stream and trust receipt ledger (columnar analytics); Grafana for ops dashboards; Sentry for error tracking	Build: agent event schema, trust receipt indexer Buy: ClickHouse, Grafana, Sentry
Infrastructure	AWS EKS (us-east-1 + eu-west-1) + Cloudflare + Terraform + GitHub Actions + Docker/OCI	EKS as primary compute; eu-west-1 mandatory for EU AI Act data residency; Cloudflare for edge + R2 fallback; Terraform IaC; GitHub Actions CI/CD; OCI agent packaging	Build: Terraform modules, CI/CD pipelines, OCI agent spec Buy: AWS, Cloudflare, GitHub

Layer	Choice	Rationale	Build or Buy
Compliance tooling	Vanta + custom AI Act evidence pipeline	Vanta for SOC2 Type II automation. Custom AI Act evidence pipeline is our differentiator: automated Annex VIII documentation, conformity assessment packaging, Notified Body submission tooling	Build: AI Act evidence pipeline entirely in-house (this is the moat) Buy: Vanta
Developer SDKs	Python + TypeScript (Phase 1–3) Go (Phase 3+)	Python targets robotics/data science engineers; TypeScript targets web/full-stack consumers. Go deferred to Phase 3 when infrastructure-adjacent use cases emerge	Build: all SDKs. Publish on PyPI and npm. Open-source SDK clients Apache 2.0 at Phase 4

Build-vs-Buy Philosophy

We build the moat. The trust ledger (Merkle-chained, ed25519-signed audit receipts), the Physical AI Bridge and its per-robot-class MCP servers, the Policy Studio UI and Rego authoring interface, the EU AI Act evidence pipeline, and the agent DID minting service. These are proprietary because no vendor ships them, they are hard to replicate, and they are what customers pay for.

We buy everything commoditized. LLM inference (OpenAI, Anthropic, Google APIs — LiteLLM makes us model-neutral); auth (Clerk handles OIDC, SAML, MFA, enterprise IdP); billing (Stripe handles metering, invoicing, dunning); databases and infra; observability tooling. Building any of these would absorb engineering time with no differentiation return.

Why not build our own vector DB? pgvector with Postgres 16 covers Phase 1–2 workloads with strong tenant isolation, ACID semantics, and no new operational surface area. We revisit at Phase 3 if p99 semantic search latency exceeds 200ms at tenant scale.

Why not build our own LLM? Our value proposition is model-neutral governance, not model capability. Building a model positions us as a competitor to our model-provider partners and destroys the neutrality story that is our primary differentiator for multi-cloud enterprises.

Long-term model. Open-source the Agent Mesh OS core at Phase 4 (Apache 2.0). Trust ledger, Policy Studio, Physical AI Bridge adapters, and the AI Act evidence pipeline remain proprietary commercial features — the Red Hat model applied to agent orchestration.

§3 Phased Roadmap — 18 Months

Phase	Dates	North-Star Metric	What Ships	Team End
Phase 0 Foundation	M0–M2 May–Jun 2026	3 signed LOIs	Dev infra, repo structure, agent identity service alpha, Vault setup, 3 L	5
Phase 1 Physical AI Bridge MVP	M3–M6 Jul–Oct 2026	LLM→ROS 2 action loop in production for 3 robot classes	Humanoid (1X NEO or Figure 02), cobot (UR10e), lab arm (OT-2) drive	8
Phase 2 Closed Beta	M7–M10 Nov 2026–Feb 2027	3 design partners live in production	Policy Studio GA; Agent Mesh OS core stable; 3 design partners on pr	11
Phase 3 Compliance + First Paid	M11–M14 Mar–Jun 2027	First paid enterprise contract ≥$500K ACV	EU AI Act conformity assessment initiated with Notified Body; first paid	13
Phase 4 Open Source + Series A	M15–M18 Jul–Oct 2027	Series A closed + $10M ARR	Agent Mesh OS core open-sourced (Apache 2.0); Series A close $30–	14

Phase 0. The first two months are not about code. They are about hiring the founding engineers, standing up dev infrastructure, locking the Phase 1 spec, and getting three signed LOIs. LOIs are non-binding but commit a named contact at the partner organization to co-design the pilot. No LOI = no Phase 1 integration work.

Phase 1 — highest risk. ROS 2 integration with production robot hardware has more unknown unknowns than any other phase. The sprint plan in §4 runs the lab-arm integration (OT-2 — lowest mechanical risk) first, then cobot (UR10e), then humanoid last in Sprint 7 after the bridge has proven itself across two simpler embodiments.

Phase 2 — demonstration to evidence. The design partner relationship moves from LOI to pilot with contractual success metrics. The console becomes a real product. SOC2 Type II audit starts — a six-month process that must begin by M7 to complete before Series A.

Phase 3 — the commercial moat. The EU AI Act conformity assessment dialogue requires: (a) Physical AI Bridge in production, (b) Trust Ledger with six months of live audit receipts from design partners, and (c) Notified Body engagement under contract. Series A conversations start the moment we can show a Notified Body engagement letter.

§4 Phase 1 Sprint Plan — Physical AI Bridge MVP

Eight sprints of two weeks each, spanning M3–M6 (July–October 2026). Goal: LLM agents controlling three robot classes via MCP bridge, running through Isaac Sim gate, with Trust Ledger alpha generating signed audit receipts for every robot command.

Sprint	Dates	Goal	Deliverables	Demo
S1	Jul 1–11	MCP server skeleton + ROS 2 adapter scaffolding	OPA policy engine integrated; first MCP server stub compiled; ROS 2 Jazzy in Docker; GitHub Actions CI green	Engineer walkthrough: policy evaluation of a dummy robot command returning allow/deny in <50ms
S2	Jul 14–25	Trust Ledger alpha — ed25519 signing pipeline live	Agent DID minting service; ed25519 key generation per agent; first signed audit receipt written to Postgres; Merkle chain root committed to S3	Mint an agent DID, issue robot command, retrieve signed receipt with verified Merkle chain
S3	Jul 28–Aug 8	First LLM→ROS 2 action loop in Isaac Sim	Claude (via LiteLLM) instructs simulated UR10e to pick-and-place a vial in NVIDIA Isaac Sim; OPA policy check passes; receipt signed	60s video: Claude agent issuing pick-and-place to simulated UR10e, receipt shown in terminal
S4	Aug 11–22	OT-2 lab arm — first physical robot	OT-2 MCP server complete; LLM agent executes liquid dispensing protocol via ROS 2 bridge; WebRTC telemetry feed live in console stub	GPT-4o agent executes 3-step pipetting protocol on OT-2; telemetry in browser; audit receipt verified
S5	Aug 25–Sep 5	Safe-stop proof + kill switch propagation	Safe-stop proof token (triple-signed: agent DID + operator key + hardware attestation) generated; kill switch tested; p99 safe-stop ≤120ms	Live: issue kill switch from console, OT-2 halts, revocation receipt logged; latency measured
S6	Sep 8–19	UR10e physical cobot integration	UR10e MCP server; physical UR10e arm running LLM-directed pick-and-place; Isaac Sim gate live; fleet telemetry for OT-2 + UR10e in console	Sim gate rejects unsafe trajectory, then approves safe version; physical execution shown
S7	Sep 22–Oct 3	Humanoid integration — 1X NEO or Figure 02	Humanoid MCP server (GR00T N1 dispatch); LLM→humanoid action sequence in Isaac Sim; WebRTC at 25Hz; human gate before physical command	Claude instructs simulated humanoid to carry sample rack; human gate approval required; receipt chain shown
S8	Oct 6–17	Integration hardening + alpha handoff	Three robot classes integrated; Trust Ledger running continuously; Policy Studio alpha UI deployed; partner onboarding guide written	Live partner demo: agent receives lab task, dispatches to three robot classes, receipt chain records every action

Sprint convention. Each sprint ends with a Friday demo (30 minutes, internal). Demo artifacts — video recordings, terminal logs, receipt chain exports — are committed to the repo. No sprint advances without a passing demo. Failed demos extend the sprint by one week. This constraint is non-negotiable.

§5 Team Plan — 14 Hires in Priority Order

All compensation in USD annual base salary. Equity is grant-date %; vesting 4-year/1-year cliff. Hiring authority: Gaurav Sisodia approves all offers.

#	Role	Month	Band (Base)	Why Now
1	Founding Engineer — Distributed Systems	M0	$180–220K	Builds Agent Mesh OS scheduler, Ray cluster, LiteLLM routing. Must be on
2	Founding Engineer — Robotics / ROS 2	M0	$180–220K	Owns Physical AI Bridge. Must have shipped production ROS 2 — not just c
3	Founding Engineer — Security / Cryptography	M1	$190–230K	Owns DID minting, ed25519 pipeline, Vault integration, Trust Ledger. Target
4	Design Engineer — Console + Policy Studio UI	M1	$150–185K	Builds control plane console and Policy Studio. Must be strong in Next.js 15,
5	ML Engineer — Agent Evals + Safety	M2	$170–210K	Builds eval harness, outcome scoring for self-evolving workflows, safety regr
6	Compliance Lead — ex-auditor, AI Act specialist	M3	$160–200K	Owns EU AI Act evidence pipeline, Notified Body engagement, SOC2 Type
7	Founding PM — ex-enterprise robotics	M3	$160–200K	Owns design partner relationships, sprint prioritization, Phase 2 roadmap. Ta
8	Senior Engineer — Data / Observability	M4	$160–195K	Owns ClickHouse cluster, OpenTelemetry instrumentation, Grafana dashboa
9	Senior Engineer — Robotics Integration #2	M5	$160–195K	Expands Physical AI Bridge to AMR (MiR250) and drone (MAVLink/PX4); su
10	DX Engineer — SDKs + Docs	M7	$145–175K	Owns Python SDK, TypeScript SDK, developer docs site, API reference. Ta
11	Head of Sales — Enterprise AE	M9	$160–200K + OTE	First quota-carrying enterprise rep. Owns closing first paid contracts in Phas
12	Solutions Engineer	M10	$145–175K	Technical support for sales: POC design, integration scoping, post-close onb
13	Senior Engineer — Backend Scaling	M12	$160–195K	Prepares platform for Series A scale: scheduler horizontal scaling, multi-tena
14	Head of Finance / Ops	M14	$160–200K	Owns Series A data room, financial model, cap table, ops build-out. Target:

Hiring philosophy. Founding engineers (#1–5) receive 0.5%–1.5% equity, reflecting pre-PMF risk. Later hires (#6–14) receive 0.1%–0.4%, scaled by seniority and timing. Compensation bands are the bands — we do not compete with Meta or Google on cash; we compete on mission and equity upside. Hiring is founder-sourced first before engaging recruiters. Recruiter fees budgeted at 15–20% of first-year salary for positions 6–14.

§6 18-Month Budget

All figures in USD. Projected over 18 months (May 2026 – October 2027).

Line Item	Total 18M	Monthly (Avg)	Notes
Salaries + benefits (25% load)	$8,100,000	~$450K → $650K	14 hires phased per §5. Founder salary $175K from M0. Benefits
Infrastructure / cloud	$900,000	~$50K	AWS EKS (us-east-1 + eu-west-1), Cloudflare R2, S3, ClickHouse
Robot hardware — internal lab (CapEx)	$1,200,000	—	1× humanoid (~$150K), 2× UR10e (~$80K ea), 2× OT-2 (~$12K e
Legal / compliance / audits	$500,000	~$28K	IP counsel + patent filings, Series A legal, Notified Body engagem
Office / ops	$400,000	~$22K	Colocated lab space with robot-safe flooring and power. Internet,
Marketing / events	$300,000	~$17K	ROSCon 2026, NeurIPS, ICRA, EU AI Act practitioner events. Tec
Contingency (5.5%)	$700,000	—	Covers hardware delays, additional Notified Body fees, emergenc
Total projected burn	**$12,100,000**	**~$672K**	Peak monthly burn ~$800K in Phase 3–4 as sales team ramps.
Recommended seed raise	**$15,000,000**	**—**	$2.9M buffer beyond projected burn = ~4 months of Phase 4 burn

Monthly burn curve (conceptual). Phase 0 (M0–M2): ~$200K/mo — founding salaries, legal, dev tooling. Phase 1 (M3–M6): ~$400K/mo — hardware CapEx; robotics team fully staffed. Phase 2 (M7–M10): ~$550K/mo — compliance, observability, DX hires. Phase 3 (M11–M14): ~$700K/mo — sales team added; Notified Body fees; Series A legal. Phase 4 (M15–M18): ~$800K/mo — peak headcount; open-source launch; Series A capital provides next round.

§7 Design Partners — Target List

Three anchors across pharma, manufacturing, and capital markets. Contract shape: 12-month paid or fee-waived pilot ($150K–$250K for paid contracts), defined success metrics, joint press release at GA, customer reference rights at Series A.

Pharma / Biotech

Target	Entry Point	Rationale	Success Metric
Recursion Pharmaceuticals	Head of Platform Engineering	AI-native pharma, heavy ML infrastructure, GPU fleet management, existing robotic lab platforms	LLM agent + OT-2 completes 10 compound prep protocols/week with full audit trail; zero manual re-runs required
Insitro	CTO	ML-first drug discovery; familiar with lab automation; smaller decision unit than Big Pharma	Insitro lab automation pipeline runs through gsiso Physical AI Bridge; audit receipts satisfy internal compliance review
Novartis (Digital & Tech)	Head of AI Lab, Basel	Scale validates enterprise readiness; EU HQ validates AI Act story directly; large compliance budget	Lead-to-IND agent pack pilot with physical lab arm; at least one EU AI Act Annex III requirement demonstrably satisfied

Manufacturing / Industrial

Target	Entry Point	Rationale	Success Metric
Siemens Digital Industries	Head of Industrial AI	Existing ROS 2 toolchain; open to vendor-neutral orchestration; EU-based compliance budget	UR10e cobot fleet at one Siemens plant managed through gsiso; VLA policy update executed via sim gate without downtime
Fanuc America	Chief Robotics Officer	Largest cobot installed base in North America; no existing AI orchestration layer	Five Fanuc CR-series cobots receive LLM-directed task scheduling via MCP bridge; throughput improvement measurable in 30 days
Boston Dynamics (Hyundai)	VP Software Platforms	Atlas pilot at Hyundai Georgia plant live; software orchestration is the gap; BD cannot build this fast enough internally	Spot or Atlas fleet telemetry ingested; at least one LLM-directed inspection workflow running in production

Capital Markets

Target	Entry Point	Rationale	Success Metric
Citadel Securities	Head of Technology Infrastructure	Multi-cloud, compliance-obsessed, willing to pay for audit-quality agent traces; high ACV potential	Research synthesis swarm producing MiFID II-compliant audit trails; agent actions traceable to DID in < 5 seconds
Two Sigma	Head of Systems	Known ML infrastructure sophistication; open to vendor-neutral tooling; self-contained decision-making	Internal agent workflow runs across two model providers simultaneously under gsiso governance; no single-provider lock-in demonstrated

§8 Risk Register — Top 10

#	Risk	Likelihood	Impact	Mitigation
1	Hyperscaler bundles governance for free	High	High	Ship Physical AI Bridge + EU AI Act conformity assessment first. Multi-cloud enterpris
2	Figure / 1X humanoid hardware slips	Medium	Medium	Multi-vendor from day one: UR10e cobot (S6) and OT-2 (S4) do not depend on humar
3	Open-source governance gets hostile fork	Low	High	Register 'gsiso' and 'Agent Mesh OS' trademarks before open-source launch at M15.
4	EU AI Act enforcement delayed	Medium	Medium	Pivot pitch to SOC2 Type II + ISO 42001 if EU enforcement slips. Both are buyer-read
5	Safety incident at design partner site	Low	Existential	Isaac Sim gate mandatory before every physical policy update — not optional, not byp
6	Key founding hire quits	Medium	High	Four-year vesting / one-year cliff. Co-founder equity offer to engineers #1 and #2 if the
7	LLM price collapse hurts margins	High	Low	Our value is governance, identity, and physical AI integration — not inference. LiteLLM
8	Model provider deplatforms us	Low	Medium	Multi-model from day one. No single provider handles >60% of any customer's inferen
9	Design partner pilot fails publicly	Medium	High	Define pilot success metrics narrowly and conservatively in writing before integration b
10	Series A market closes	Medium	High	Reach $10M ARR with 3 named production customers before opening Series A conve

§9 Decision Gates

Three explicit go/no-go decision points. Gaurav Sisodia makes the call with input from the founding team. Pivot options are listed so every decision has known downside paths.

Gate 1 — M6 (End of Phase 1 · October 2026)

Proceed if all are true:

- Three robot classes integrated and demonstrated to design partners
- Trust Ledger generating signed receipts for every robot command, verifiable by independent party
- At least one design partner converts from LOI to active Phase 2 pilot
- Monthly burn ≤$420K

Pivot options if criteria not met:

- Drop humanoid from Phase 2; focus on cobot + lab arm — still defensible
- If no design partner converts: extend LOI engagement one month; if still no conversion, reassess ICP toward manufacturing over pharma

Gate 2 — M12 (Mid-Phase 3 · April 2027)

Proceed if all are true:

- At least 2 design partners running production workloads through gsiso
- Policy Studio GA and publicly demoed
- Notified Body engagement letter signed or conformity dialogue actively underway
- SOC2 Type II audit in progress
- Monthly burn ≤$700K

Pivot options if criteria not met:

- If only 1 design partner in production: delay Series A, use M13–M14 to close a second partner
- If Notified Body engagement cannot be initiated: substitute ISO 42001 third-party certification as compliance narrative
- If burn exceeds $700K: defer Head of Sales (position 11) by 2 months

Gate 3 — M18 (End of Phase 4 · October 2027)

Proceed if all are true:

- Series A closed ($30–50M)
- $10M ARR in signed contracts
- 3 named production customers with reference rights
- Agent Mesh OS core open-sourced on Apache 2.0

Pivot options if criteria not met:

- If Series A closes at <$30M: accept if runway to $25M ARR is credible; otherwise bridge with $3–5M from existing seed investors

- If $10M ARR not reached: evaluate acqui-hire by hyperscaler at $150–300M (10–20× on $15M seed for investors)
- If open-source launch delayed: not blocking for Series A; push to M21

§10 What We Explicitly Don't Build

Our own LLM. LLM training and inference is a capital-intensive commodity race. Our value is governance over LLMs, not LLM capability. Building a model makes us a competitor to OpenAI, Anthropic, and Google — providers whose API relationships are central to our multi-model neutrality story.

Our own vector database. pgvector on Postgres 16 covers Phase 1–3 workloads. Purpose-built vector databases add operational complexity with no differentiation return. Revisit at Phase 3 if benchmarks demand it.

Our own cloud or bare-metal infrastructure. Building a proprietary cloud is a decade-long, multi-billion-dollar undertaking. We run on AWS with Terraform IaC ensuring we're not locked to AWS primitives.

A generic 'build your own agent' UI. Competitors like xpander.ai, Beam, and Relevance AI own the low-code agent builder space. Our console is a governance and operations console, not an agent authoring tool.

A consumer application. All product decisions are made for regulated enterprise buyers. Consumer applications require different distribution, pricing, support, and regulatory considerations.

A marketplace before M18. The marketplace requires a developer community that has not yet been built. Open-source at M15 seeds the community; marketplace launches post-Series A when there is organic supply.

Non-robotic physical verticals before M18. Drones, warehouse AMRs, and smart building IoT are genuine Physical AI Bridge opportunities — for Phase 5, post-Series A. The cobot/humanoid/lab-arm focus is sufficient to prove the category and close the first three design partners.

§11 Closing: The 18-Month Bet

The thesis is specific. By October 2027, gsiso.ai will have:

- Connected LLM agents to physical robot fleets — humanoids, cobots, lab arms — in production environments, generating cryptographically signed audit receipts for every physical action. No company will have done this at enterprise scale before us.
- Initiated the first EU AI Act conformity assessment dialogue for a multi-agent orchestration platform. The Notified Body engagement letter will be public. European regulated enterprises will have a procurement shortcut that does not exist from any hyperscaler.
- Named three production customers across pharma, manufacturing, and capital markets — not pilots, not demos, but running workloads with signed reference rights.

> **The bet, precisely stated: Three regulated-industry design partners in production + Physical AI Bridge GA + first EU AI Act conformity dialogue by M12 = category ownership. Every engineering, hiring, and budget decision in this document is oriented toward that single outcome.**

If all three are true by M12, we own the category. The Physical AI Bridge plus EU AI Act certification is a combination no incumbent can replicate in 90 days — it requires robotics engineering talent and a Notified Body certification process that runs on calendar time, not engineering sprints.

If we reach M12 with two of the three, we raise Series A on what we have and use the capital to close the third. If we reach M12 with none — no design partner in production and no Notified Body engagement — we sell to a hyperscaler. The Physical AI Bridge capability and compliance engineering team are worth $150–$300M to Microsoft, AWS, or Google even in that scenario. The sell outcome is a floor that is acceptable to investors at a $15M seed raise.

The window is real. The competition is running. We ship.

Sources

1. Fortune Business Insights — Agentic AI Market Size: https://www.fortunebusinessinsights.com/agentic-ai-market-114233
2. Robotics Center — State of Robotics 2026: https://www.roboticscenter.ai/state-of-robotics-2026
3. Trilateral Research — EU AI Act Implementation Timeline: https://trilateralresearch.com/responsible-ai/eu-ai-act-implementation-timeline-mapping-your-models-to-the-new-risk-tiers
4. Fifthrow — AI Agent Orchestration Goes Enterprise (April 2026): https://www.fifthrow.com/blog/ai-agent-orchestration-goes-enterprise-the-april-2026-playbook-for-systematic-innovation-risk-and-value-at-scale
5. VentureBeat — Most Enterprises Can't Stop Stage Three AI Agent Threats: https://venturebeat.com/security/most-enterprises-cant-stop-stage-three-ai-agent-threats-venturebeat-survey-finds
6. Futurum Group — Who Will Win the Agent Orchestration Layer Battle: https://futurumgroup.com/press-release/who-will-win-the-agent-orchestration-layer-battle/
7. MetricStream — NIST's AI Agent Standards Initiative: https://www.metricstream.com/blog/nists-ai-agent-standards-initiative.html
8. AI World Journal — Introducing the AI Kill Switch for Agents: https://aiworldjournal.com/introducing-the-ai-kill-switch-for-agents/
9. NVIDIA — Isaac GR00T N1 Open Humanoid Foundation Model: https://nvidianews.nvidia.com/news/nvidia-isaac-gr00t-n1-open-humanoid-robot-foundation-model-simulation-frameworks
10. youngju.dev — Humanoid Robots 2026 Complete Guide: https://www.youngju.dev/blog/ai/2026-03-03-humanoid-robots-2026-complete-guide.en

11. CMSWire — Sierra AI $10B Valuation:
https://www.cmswire.com/customer-experience/sierra-ais-10b-valuation-marks-a-turning-point-for-conversational-ai/
12. TechCrunch — Cognition AI $10.2B Raise:
https://techcrunch.com/2025/09/08/cognition-ai-defies-turbulence-with-a-400m-raise-at-10-2b-valuation/
13. StackOne — MCP: Where It's Been, Where It's Going: https://www.stackone.com/blog/mcp-where-its-been-where-its-going/
14. Stellagent — A2A Protocol Google Agent-to-Agent: https://stellagent.ai/insights/a2a-protocol-google-agent-to-agent
15. Ringly.io — AI Agent Statistics 2026: https://www.ringly.io/blog/ai-agent-statistics-2026