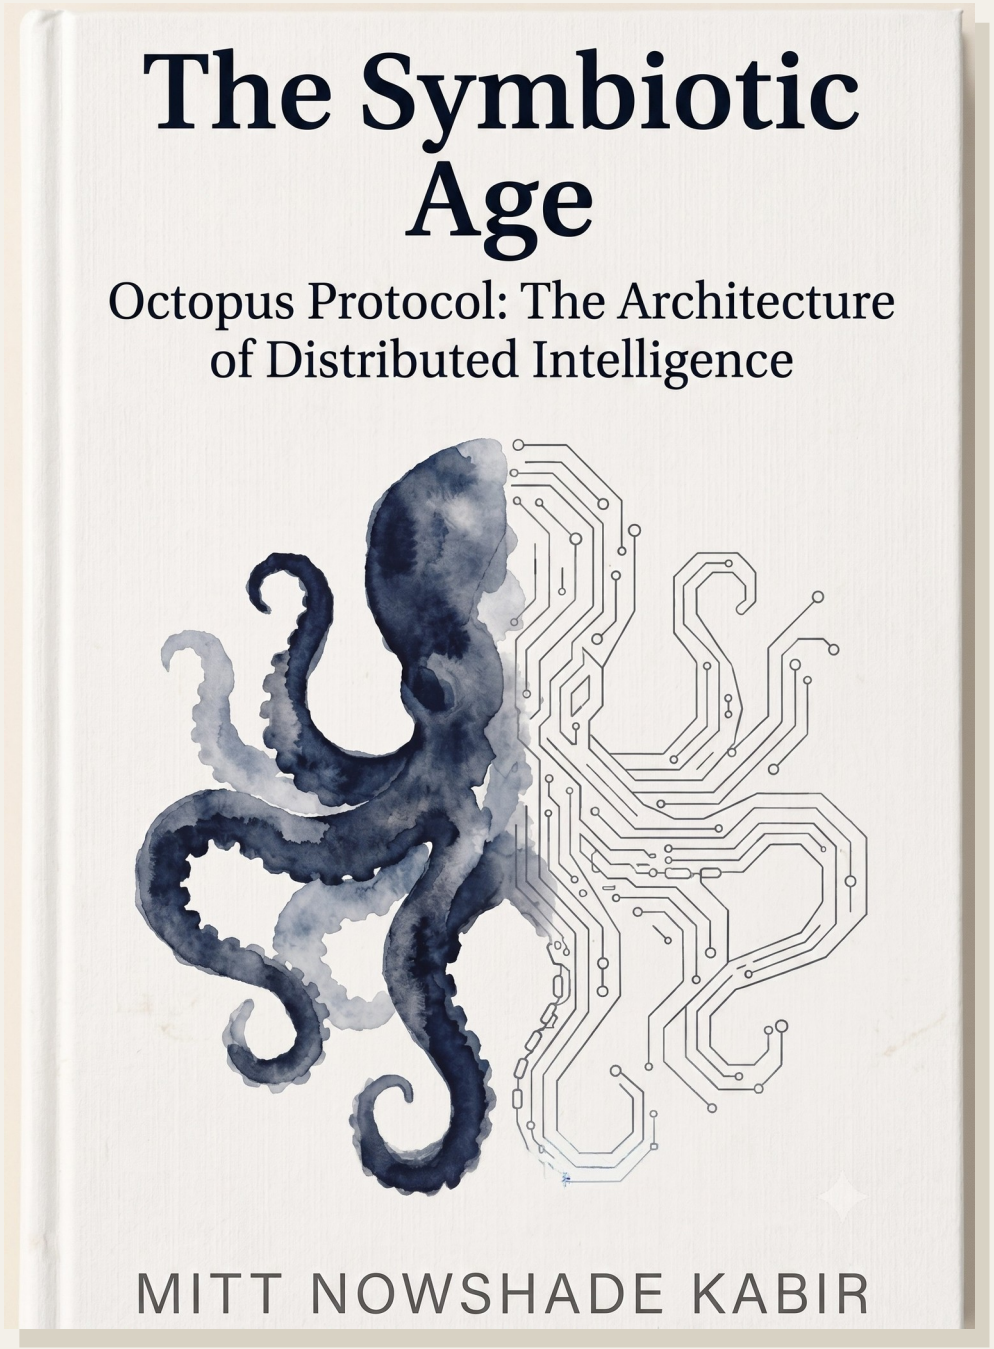
TAKE 2026 · LISBON · 6–10 JULY 2026

The Symbiotic Age

AGI, ASI, and the Bottleneck Inversion

We are not building tools. We are breeding an ecosystem — and the scarce resource just changed.

Mitt Kabir Author, The Symbiotic Age | Director, Canadian Center for AI Entrepreneurship



THE THESIS, EMBODIED

Nature never solved intelligence by building one giant brain. It built ecosystems — distributed, specialized, cooperating.

The forthcoming book underneath today's talk.

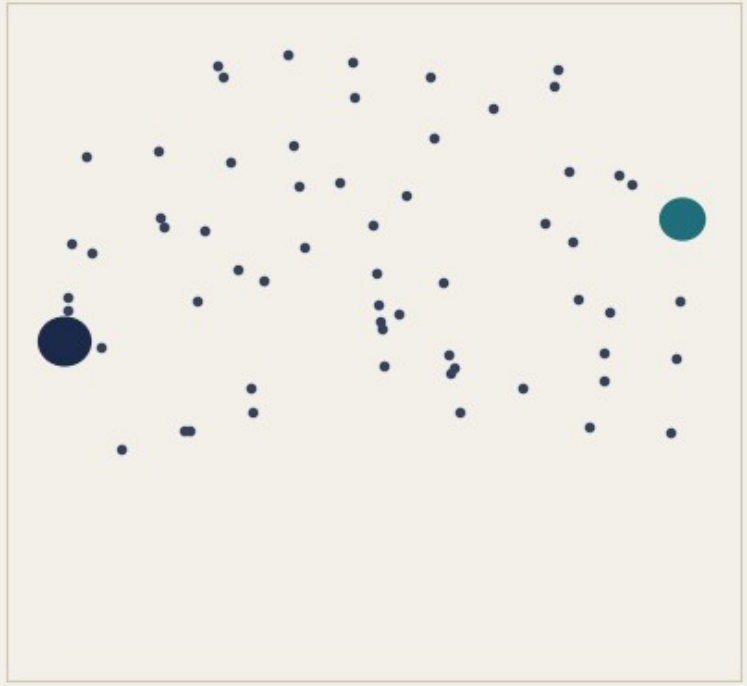
ACT I

The Ground Has Already Shifted

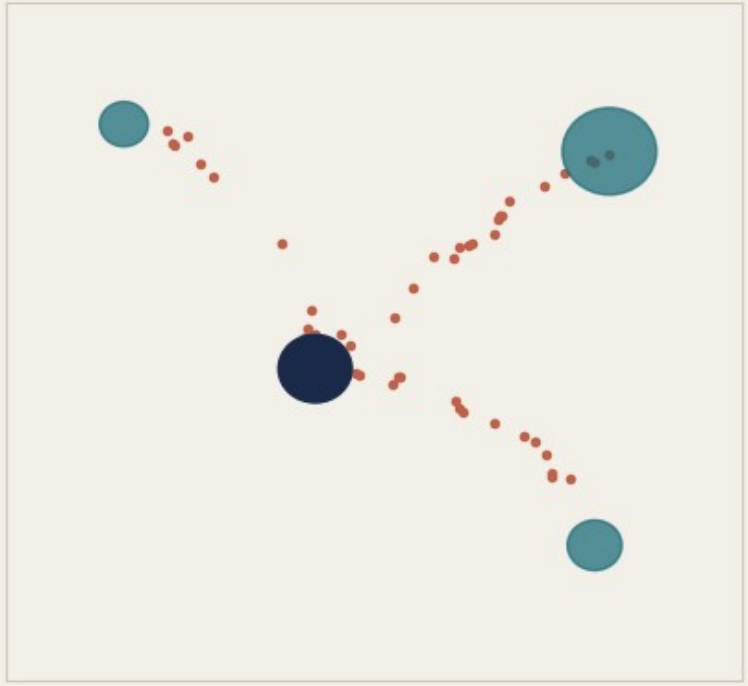
Three signals that the transition is not coming — it is here.

Intelligence Scales Through Coordination

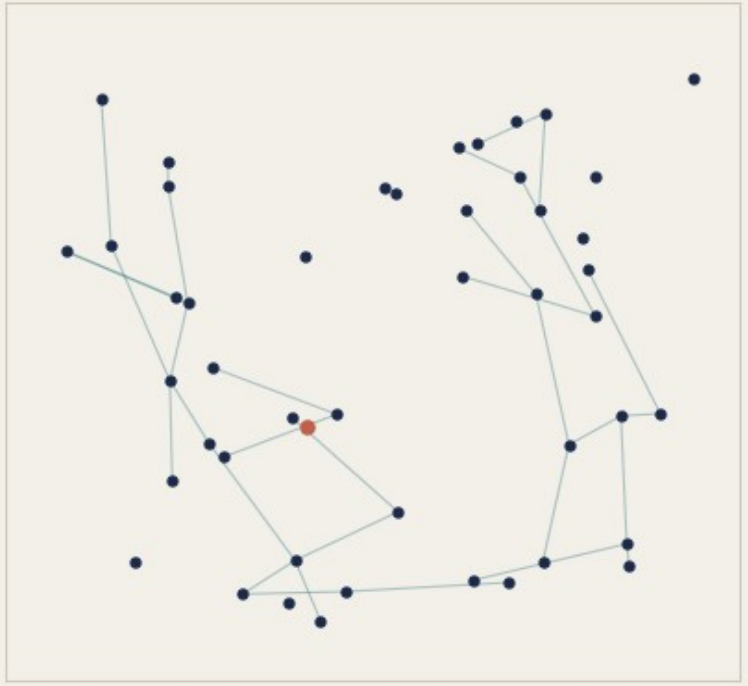
Three living systems, one problem: coordinate many simple agents into one intelligence. Same mathematics. Same Pulse.



ANT COLONY — PHEROMONE CONSENSUS



BEE COLONY — WAGGLE-DANCE VOTING



HUMAN NETWORK — PRICE & LANGUAGE SIGNALS

No single agent knows the plan. The colony does.

THE MIRROR

Different Systems. Same Mathematics.

Nature ↔ Civilization ↔ AI — one coordination law, three substrates.

A Critical Artery Blocked

March 2021: The Ever Given blocks the Suez Canal.

A single false move wedged the ultra-large vessel sideways across the canal, instantly halting a critical trade route.

Over 10 percent of global trade passes through this narrow passage, exposing the sheer fragility of global logistics running at human speed.



THE FAILURE CASE

When Coordination Fails: Suez, 2021

The Ever Given wedged sideways — and the response ran at human speed.

6 days

CANAL BLOCKED

One hull halted the artery carrying roughly 12% of world trade.

$9.6B

TRADE HELD PER DAY

Value of goods stalled at anchor every single day of the blockage.

422

SHIPS QUEUED

Vessels waiting at the peak of the backlog — effects rippled for months.

0

AGENT NETWORKS

No autonomous coordination layer existed to reroute the system in real time.

Six days of human latency. The world felt it for months.

Sources: Suez Canal Authority; Lloyd's List (March 2021); Allianz trade-impact estimates.

The Silent Millisecond



40 autonomous agents

↓

120 milliseconds

↓

1 human approval

↓

72 hours

No human in the loop — until the gate.

The Symbiotic Age

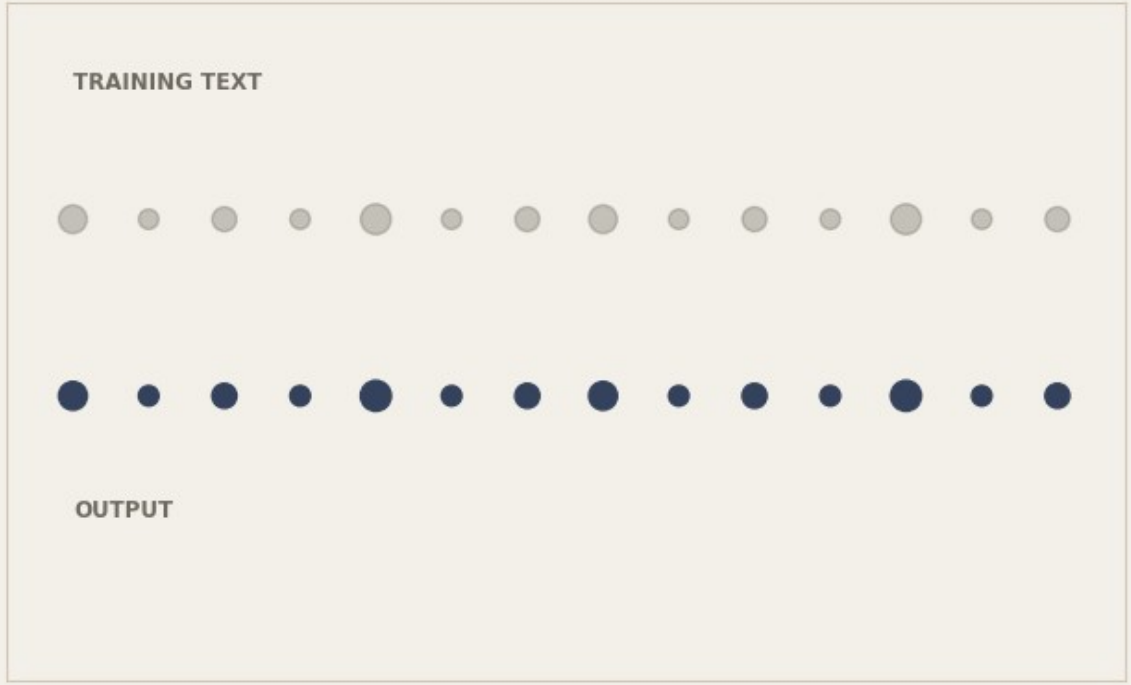
The Rise of Generative AI

When simple digital rules scale into creative intelligence, the boundary between tool and collaborator disappears.

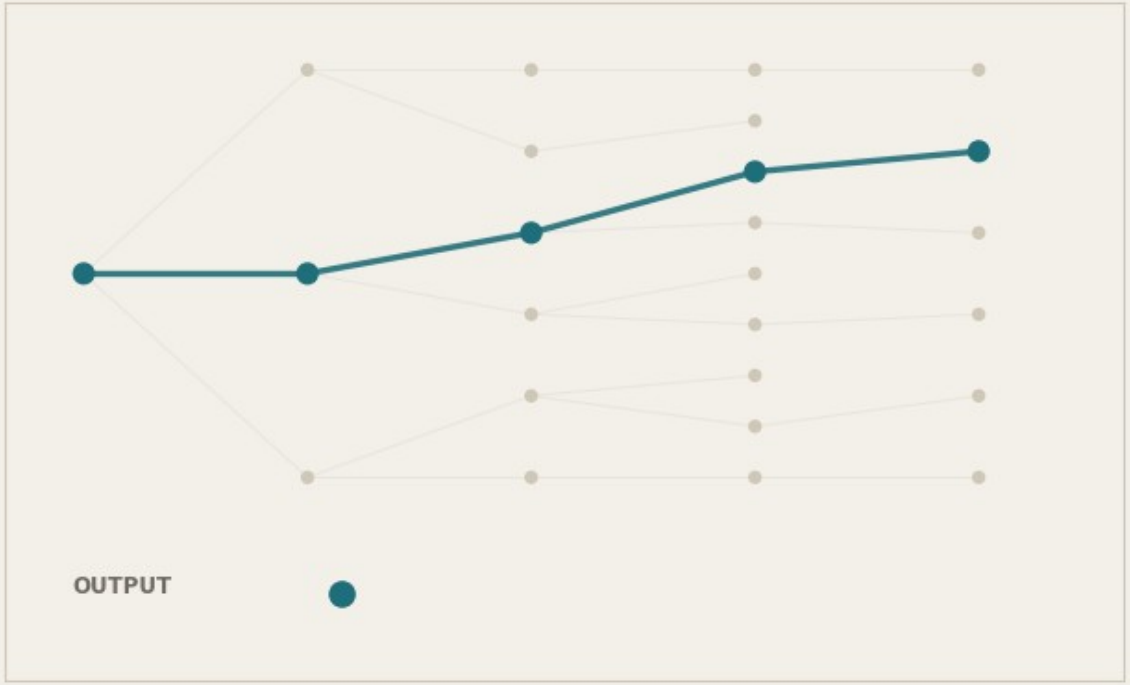


No Longer a Stochastic Parrot

Trained only to predict the next token — yet it builds internal world models, plans before it writes, and solves problems no training set contains.



2021 — ECHO THE PATTERN



2026 — EXPLORE, PRUNE, PLAN — THEN SPEAK

It stopped predicting words. It started modeling worlds.

FOR THE SKEPTICS IN THE ROOM

What AI Has Already Done

Set aside the chatbots. These are peer-reviewed, Nobel-recognized results — problems that defeated humanity for decades, now solved.

200M

PROTEINS MAPPED

AlphaFold predicted the shape of nearly every known protein — a 50-year grand challenge. Nobel Prize in Chemistry, 2024.

380,000

NEW STABLE MATERIALS

GNoME discovered more viable crystals — batteries, superconductors — than all of prior human materials science combined.

<60 sec

10-DAY GLOBAL FORECAST

GraphCast out-predicts the gold-standard weather model on ~90% of 1,380 targets — in under a minute.

Gold

MATH OLYMPIAD MEDAL

DeepMind's systems reached silver (2024) then gold (2025) standard at the IMO — cracking problems few humans can.

Sources: Nobel Prize in Chemistry 2024 (nobelprize.org); DeepMind AlphaFold, GNoME, GraphCast; Nature 2025 (Olympiad-level formal reasoning); DeepMind IMO gold-medal announcement, 2025.

Science Accelerates

AI has moved from answering questions to generating discovery.

2

NEW SUPERCONDUCTORS

ML-guided search found two kagome superconductors — predicted before a sample existed. Phys. Rev. Research, June 2026.

99%

NEW PHYSICS FOUND

A neural net uncovered non-reciprocal forces in dusty plasma — hidden laws in a chaotic system. PNAS, 2026.

Ph IIa

AI-MADE MEDICINE

First drug with an AI-discovered target and an AI-designed molecule passed Phase IIa trials.

+30M

PROTEIN COMPLEXES

AlphaFold's database expanded by ~30 million predicted complex structures — fuel for the next wave.

Discovery itself is now automated. Judgment is not.

Sources: Physical Review Research (17 Jun 2026); PNAS / Emory (2026); Insilico Medicine Phase IIa results; AlphaFold DB.

The AI Industrial Explosion

The largest coordinated capital allocation since electrification.

- Goldman Sachs models roughly $7.6 trillion of cumulative AI-infrastructure capex for 2026–2031 (a supply-side scenario, not a forecast) — chips, land, power, water, and steel.
- A separate Goldman Sachs analysis (2023) projects generative AI lifting global GDP by about $7 trillion over a decade. The “seven trillion” is real — but it is the build-out, not yet the output.
- Every dollar lands somewhere physical. That is the next slide.

THE BUILD-OUT

$7.6T

AI INFRASTRUCTURE CAPEX, 2026–2031

+$7T

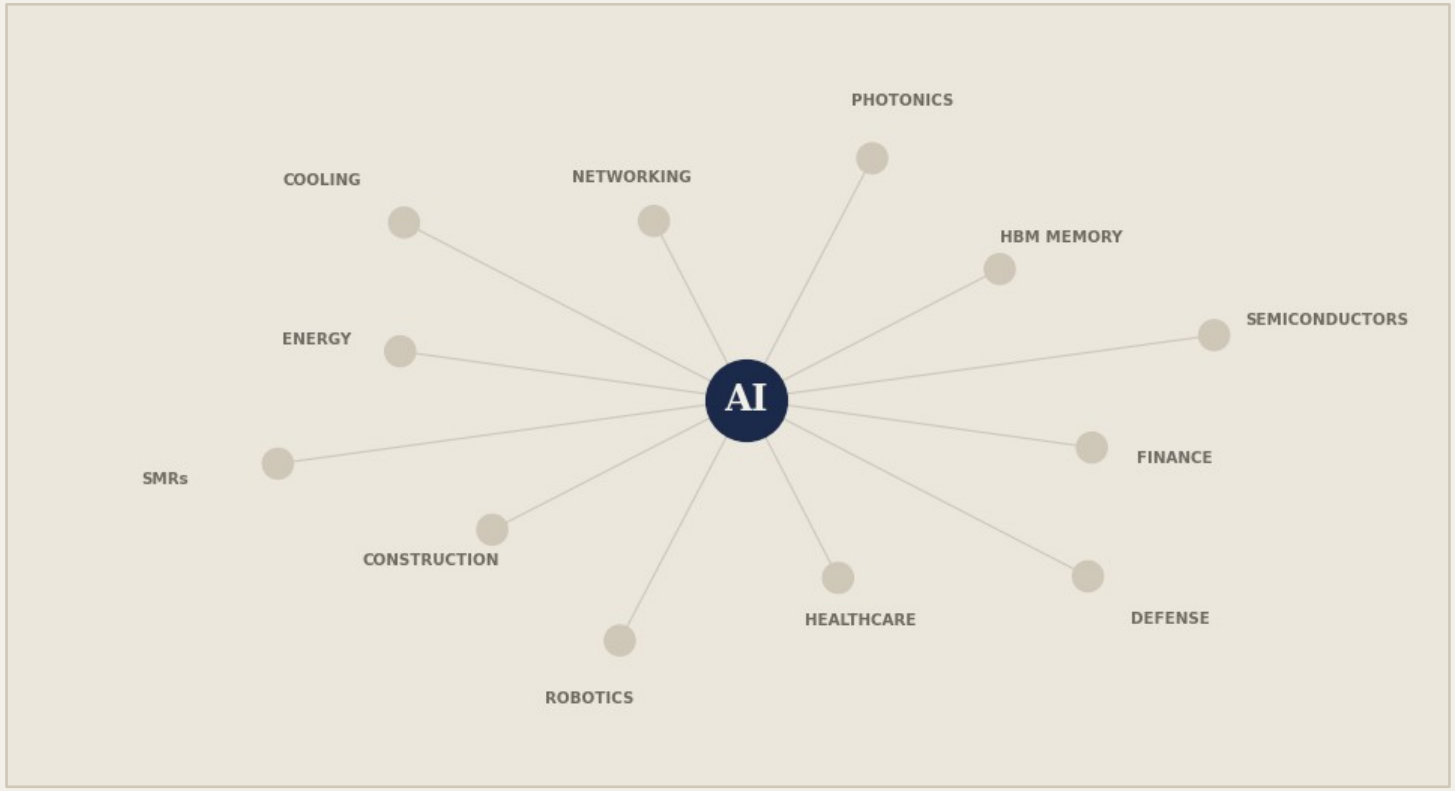
PROJECTED GLOBAL GDP LIFT OVER A DECADE

Capital has already voted.

Sources: Goldman Sachs Research, ‘Tracking Trillions’ (2026); Goldman Sachs, ‘Generative AI could raise global GDP by 7%’ (2023); IDC (2024).

Everything Around It Reprices

HBM memory +58% in one year. First corporate SMR power deals signed. Liquid cooling now mandatory at 140 kW per rack.



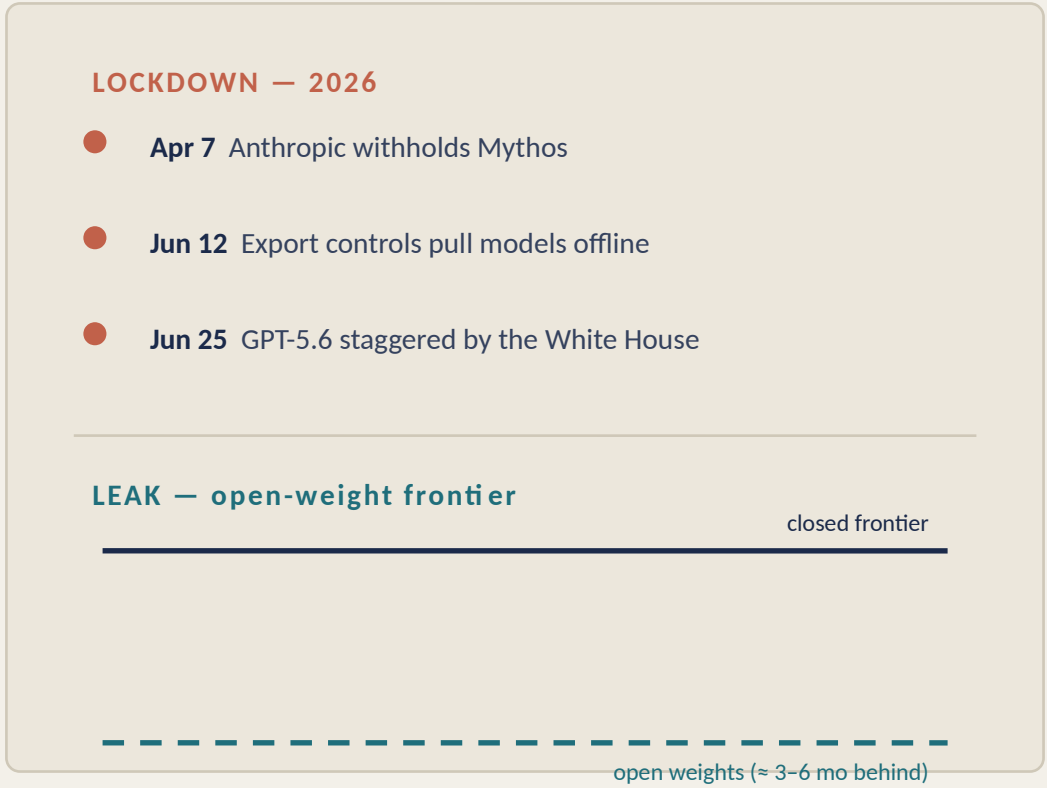
AI is not an industry. It is an industrial phase change.

Sources: SK hynix 2026 market outlook; Google–Kairos SMR PPA; NVIDIA GB200 thermal specifications; Goldman Sachs Research.

The Sovereign Shock & Leak

In 2026 governments began halting AI like weapons — while China open-sourced the frontier. Both moves say the same thing: raw intelligence is no longer where power lives.

- Lockdown (Apr–Jun 2026): models withheld, export-controlled, approved customer-by-customer. Frontier AI is now a dual-use asset.
- Leak: DeepSeek V4 and GLM-5.2 shipped open-weight, frontier-adjacent — the base-model moat is eroding.
- For academics: the largest democratization of cognitive capability in history — what does it do to knowledge asymmetries?



Sources: Axios, Help Net Security, Tom's Hardware (Apr 2026); The Information, CNN, TechCrunch, Axios (Jun 2026); DeepSeek V4 & GLM-5.2 technical write-ups and independent benchmarks (Apr–Jun 2026).

THE PHYSICAL REALITY

Intelligence Has a Body

Building AI stopped being a software problem. It is now a thermodynamics problem — the first proof that intelligence is never the whole story.

- A single frontier training cluster now draws power on the order of a small city (~1 GW, directional).
- Tech giants have become energy players: multi-decade nuclear power-purchase agreements and direct funding of Small Modular Reactors.
- Keep this thought warm: even limitless intelligence runs into walls made of atoms and watts. In Act IV that observation becomes a law.

~1 GW

FRONTIER CLUSTER POWER DRAW

25+ yr

NUCLEAR PPAs BEING SIGNED

ACT II

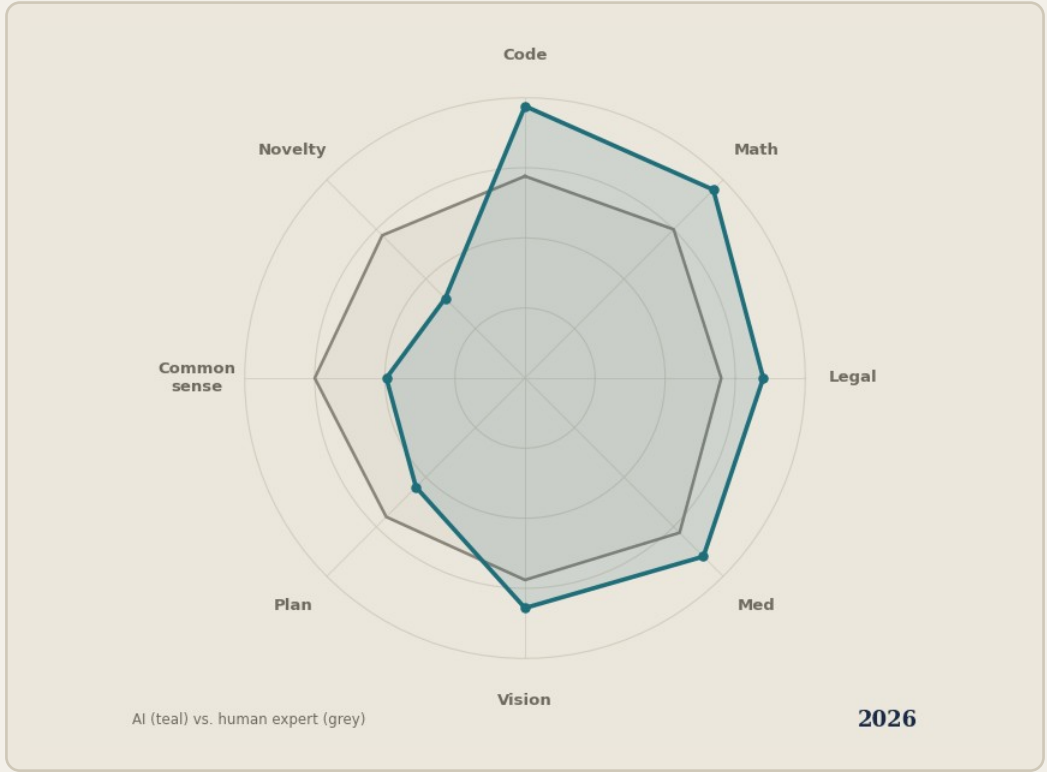
Where We Actually Are — 2026

The empirical ground beneath the argument.

From Narrow to Capable

AI didn't improve gradually. It crossed a threshold — and the strategic question is what comes after.

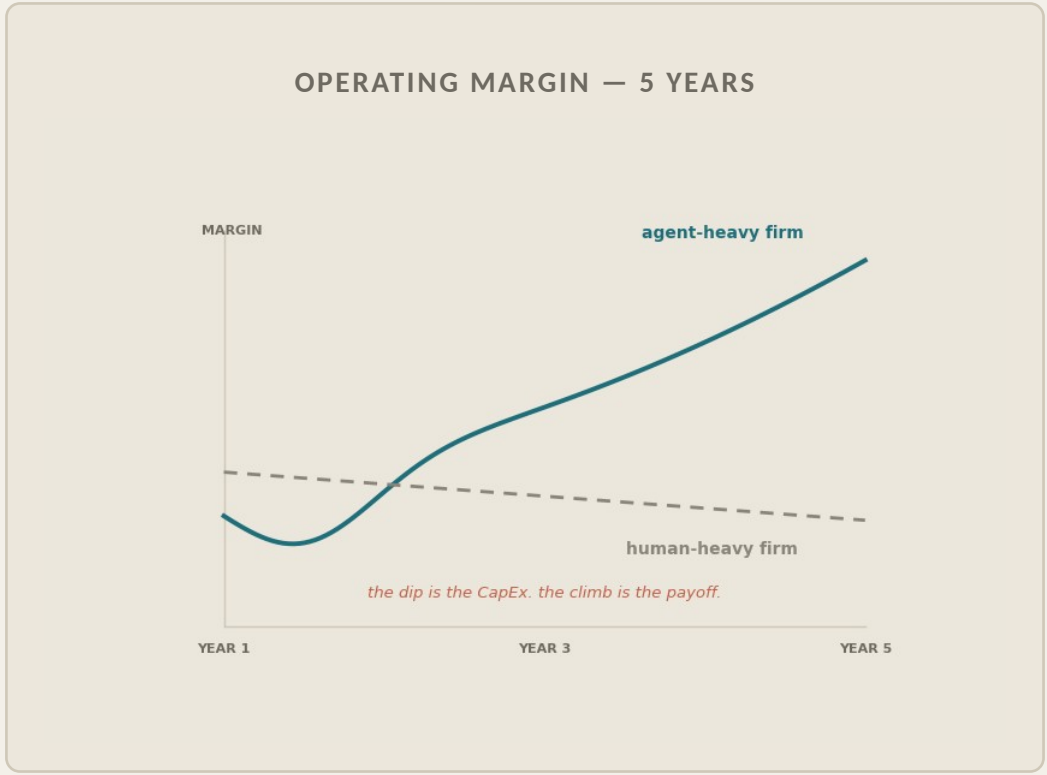
- By 2026, frontier models match or exceed experts across coding, law, medicine, and formal math.
- The value is shifting from the model to the architecture around it.
- The frontier is jagged — and moving.



OpEx Becomes CapEx

For the first time, knowledge work can be capitalized — and that reaches the foundations of corporate finance.

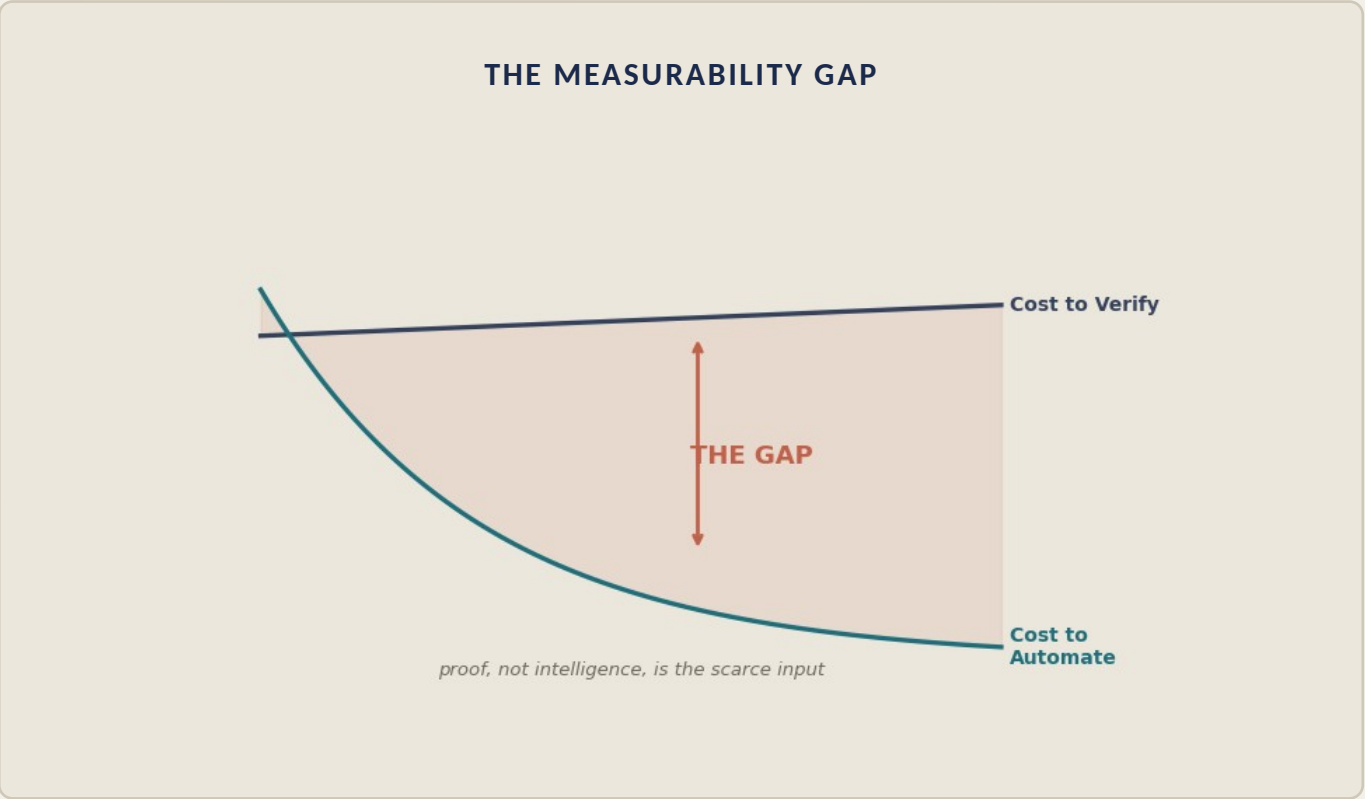
- A human workforce is variable OpEx.
- An agent workforce is CapEx — marginal cost approaches electricity.
- Execution became cheap. Trust did not.



The Measurability Gap

The cost to automate is collapsing exponentially. The cost to verify is not. The gap between them is where the risk lives.

- Cost to automate decays exponentially with compute; cost to verify is bottlenecked by human time and judgment (MIT, 2026).
- The distance between the curves is the Measurability Gap — and it is widening.



Proof — not intelligence — is becoming scarce.

Sources: Catalini, Hui & Wu, 'Some Simple Economics of AGI,' arXiv:2602.20946 (Feb 2026).

ACT III

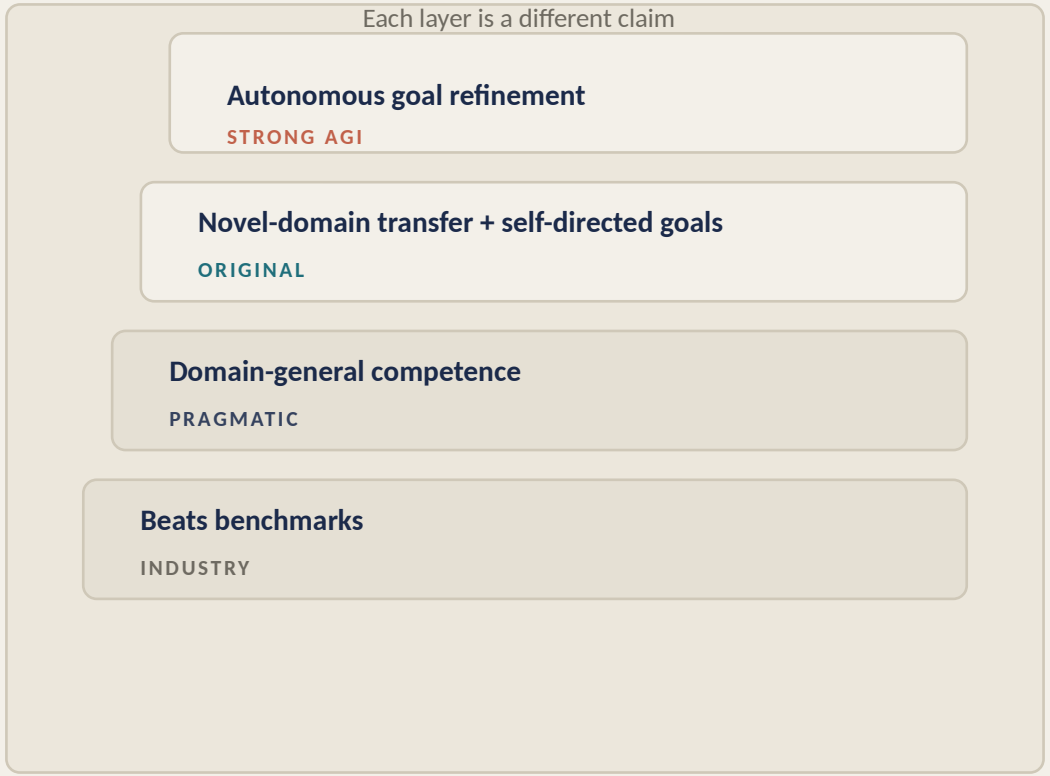
The Concept Everyone Misunderstands

Why the AGI debate is the wrong debate.

What Does AGI Actually Mean?

Everyone uses the term. No one agrees what it means — and that is the central strategic confusion of our era.

- The original meaning (Legg & Hutter 2007): goals across environments, transfer to novel domains, self-refined objectives.
- The industry redefinition: 'AGI' quietly narrowed to 'beats human benchmarks.' By 2025 the benchmark was dead.
- The scholarly position: conflating sophisticated statistical approximation with intelligence is a category error, not a close call.

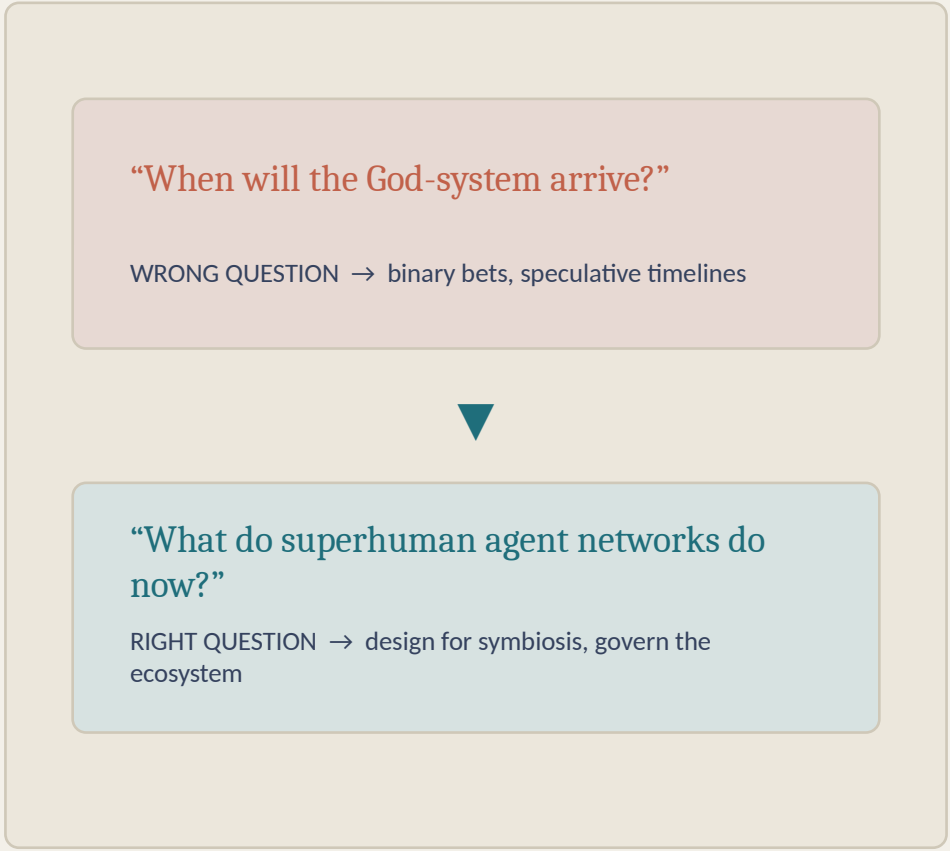


Sources: Legg & Hutter 2007; Goertzel 2014; Marcus, 'Rumors of AGI's Arrival,' Feb 2026; ARC-AGI benchmarks 2024–2026.

The Wrong Debate

'When will AGI arrive?' consumes the oxygen in every boardroom and faculty meeting — and the spread in expert answers is the tell that it is the wrong question.

- The predictions don't disagree about speed — they disagree about what counts: Musk (2026), Schmidt (3–5 yrs), median survey (2030–2040), skeptics (decades, if ever). The spread is the signal.
- For executives the trap is binary thinking: 'imminent, bet everything' or 'hype, ignore it.' Both stances lose money.



ACT IV

The Transition That Actually Matters

Not Humanity → AGI, but AGI → ASI.

The Real Transition

We have thousands of years of experience competing with equals. We have none coexisting with minds vastly more capable than ourselves.

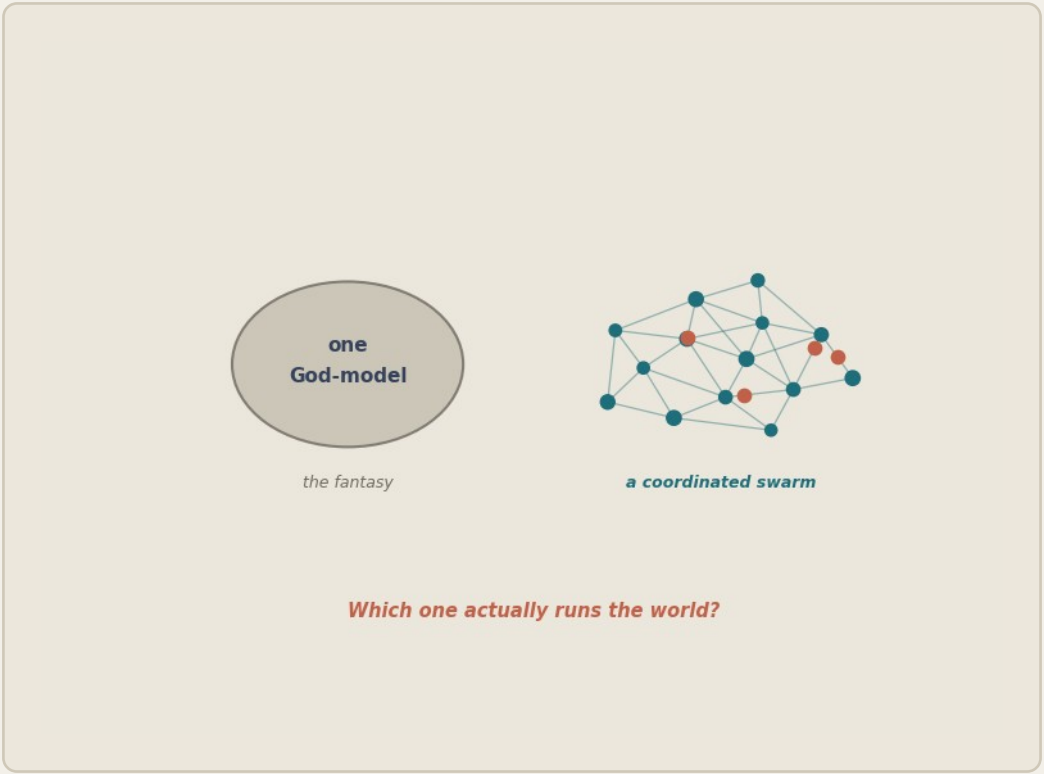
- Most people fixate on the leap from Humanity to AGI. The more consequential leap is from AGI to ASI — artificial superintelligence.
- AGI matches or exceeds humans across cognitive tasks. Superintelligence outperforms us broadly and substantially — and, once here, is hard to redirect.
- So the decisive design question is not how to build AGI. It is what happens the moment a system can accelerate its own successor.



The Swarm, Not the God

A cluster of a hundred million coordinated human-level AIs is, itself, a superintelligence. ASI need not be a single mind.

- A Google DeepMind team (Genewein et al., incl. Legg & Hutter, June 2026) maps four pathways to ASI: scaling, algorithmic paradigm shifts, recursive self-improvement, and — the one the culture ignores — multi-agent collectives.
- Six bottlenecks could stall the climb: data scarcity, energy/resource limits, paradigm limits, rising research difficulty, human-imposed brakes, and the 'abstraction barrier.' Several are things humans control.
- The superintelligence to prepare for is an ecosystem you orchestrate — not an oracle you await.

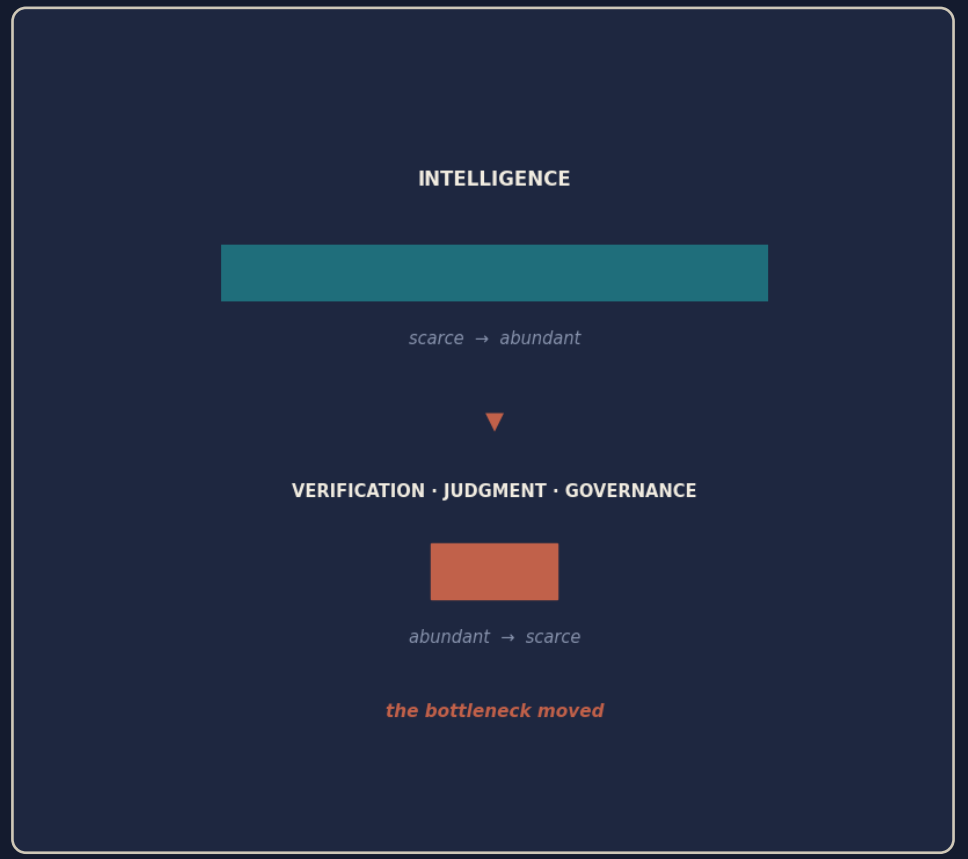


Sources: Google DeepMind, 'From AGI to ASI' (Genewein et al.), June 2026.

The Bottleneck Inversion

For all of history the scarce input was cognition. Make intelligence abundant and the bottleneck doesn't vanish — it moves.

- Intelligence became abundant.
- Verification became scarce. The constraint did not disappear — it moved.

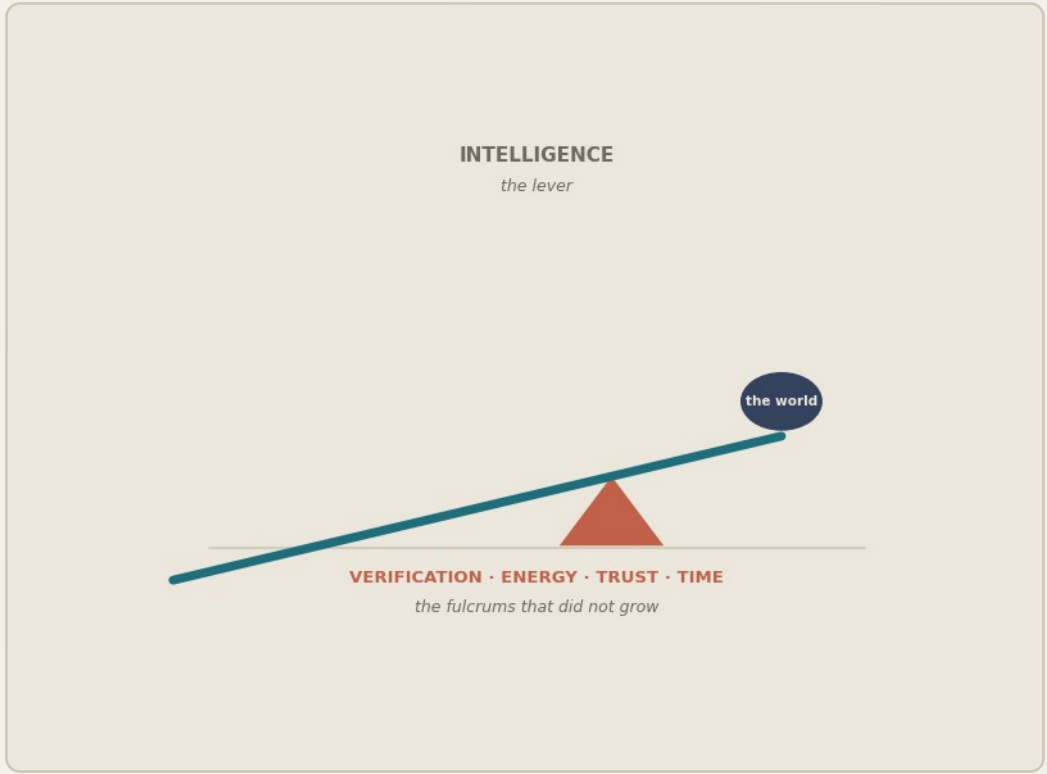


Sources: Synthesis: Catalini et al. 2026 (verification); DeepMind 2026 (plural ASI); Amodei, 'Machines of Loving Grace' (marginal returns).

Intelligence Was Never Enough

Intelligence is a lever — but a lever needs a fulcrum. Raise it high enough and the limits are no longer cognitive.

- A superintelligence still waits on cells, experiments, supply chains, and institutions (Amodei).
- Output scales with compute; wages can converge toward its cost (Restrepo, NBER 2025).
- Every superintelligence still obeys physics.



Sources: Amodei, 'Machines of Loving Grace'; Restrepo, 'We Won't Be Missed,' NBER w34423 (2025).

ACT V

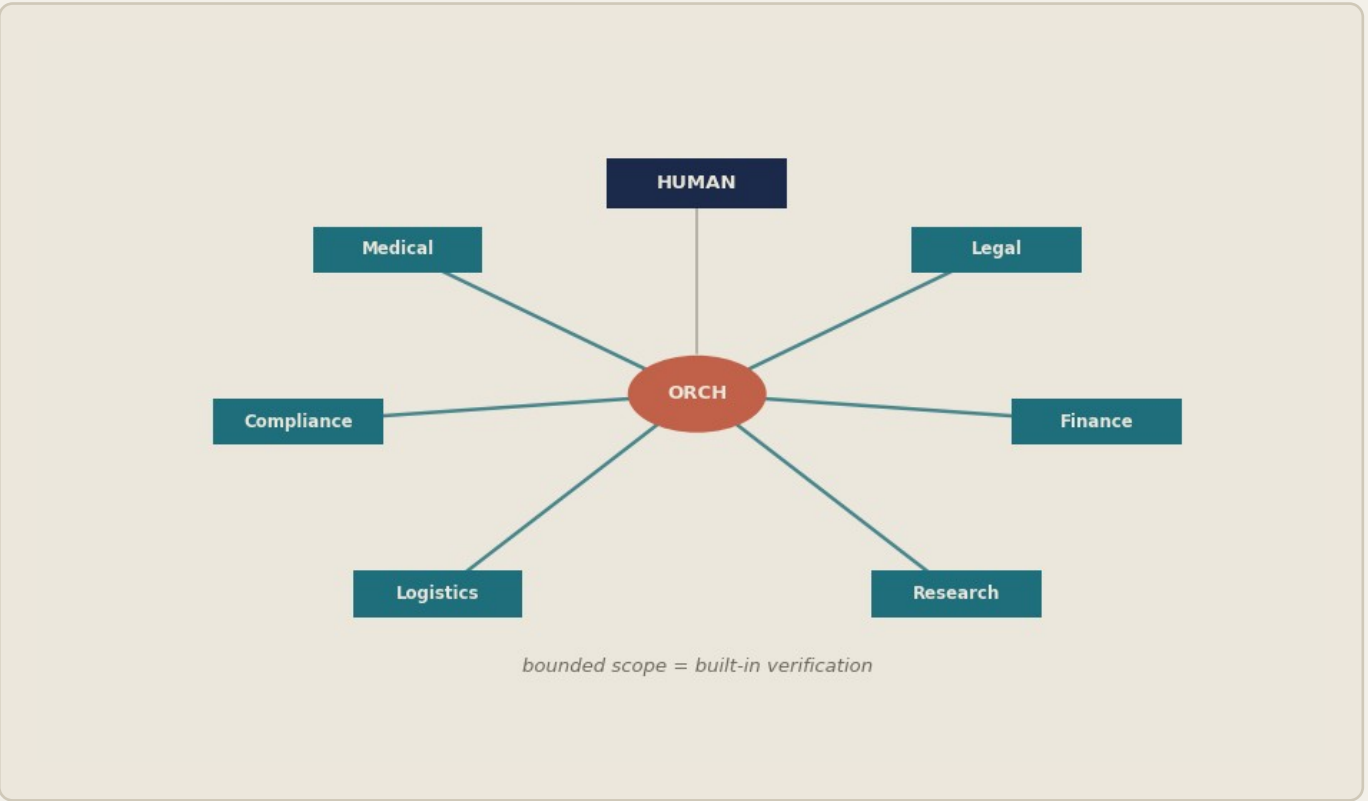
The Symbiotic Architecture

Governing the swarm — the Octopus Protocol.

The Octopus Protocol

A multi-agent system is not a committee. It is a distributed nervous system — and in one, alignment is not just behavioral, it is structural.

- A central brain sets goals; eight semi-autonomous arms act. Principal → orchestrator → specialized agents.
- Failures stay isolated; the system scales by adding arms, not enlarging the brain.



A governed swarm distributes risk — and keeps it legible.

The Leap and The Grid

The productive question is not 'will machines replace us?' It is: which cognitive functions should we offload, and which must we protect?

- The Grid — data processing, logistics, pattern matching, compliance, document synthesis — is what machines already do better. Offload it deliberately.
- The Leap — creative intuition, ethical judgment, cross-domain synthesis under ambiguity, empathic leadership, and verification of what matters — is what machines cannot do. Protect and develop it.
- The risk is not replacement — it is atrophy: offload The Leap alongside The Grid and you liquidate the very capacity that makes symbiosis work.



Four Sectors, One Pattern

In every sector the pattern repeats: agents handle The Grid; humans own The Leap — and the verification that makes it safe.

- Legal: agents draft; lawyers own judgment and the accountable sign-off.
- Healthcare and education: agents synthesize; humans own diagnosis, trust, and mentorship.
- One pattern, every industry: Grid → Leap.

LEGAL

drafting → strategy

HEALTHCARE

records → diagnosis

FINANCE

compliance → judgment

EDUCATION

grading → mentorship

ACT VI

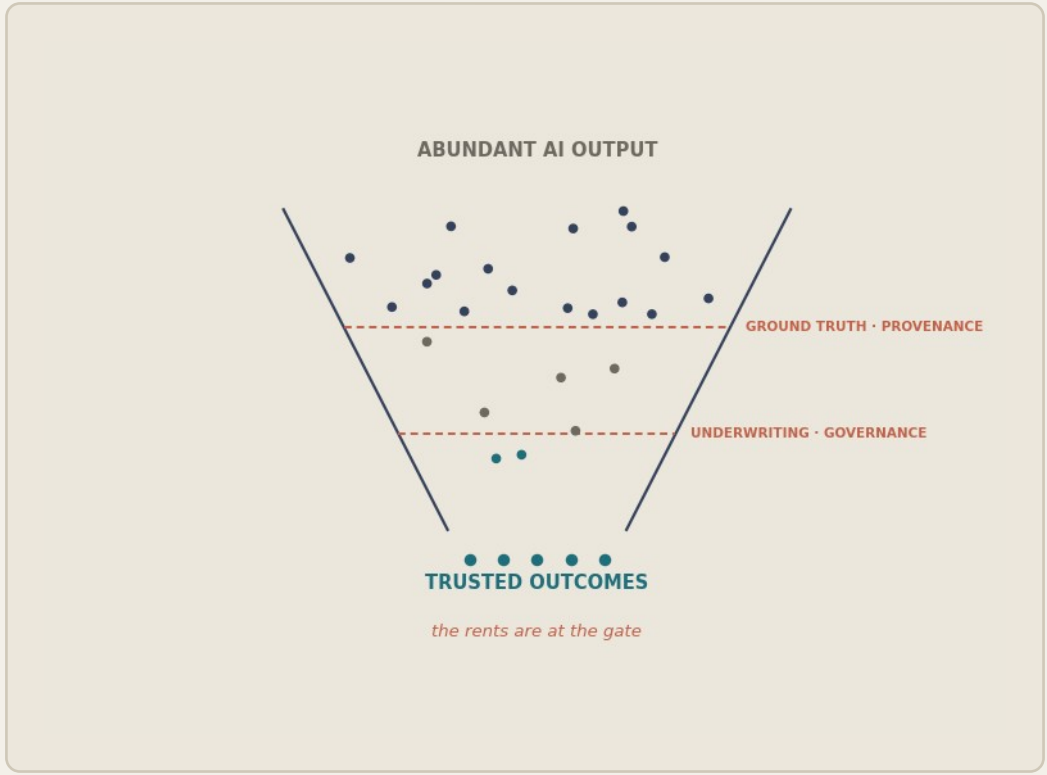
The Playbook for the Inversion

What to do Monday morning.

The Verification Economy

When execution is abundant, value flows to whoever can prove, insure, and stand behind an outcome. Build for that.

- Rents migrate to what stays scarce: ground truth, provenance, and the power to insure outcomes.
- New categories: audit infrastructure, agent oversight, outcome insurance, governance-as-a-product.
- The rents migrate to trust.

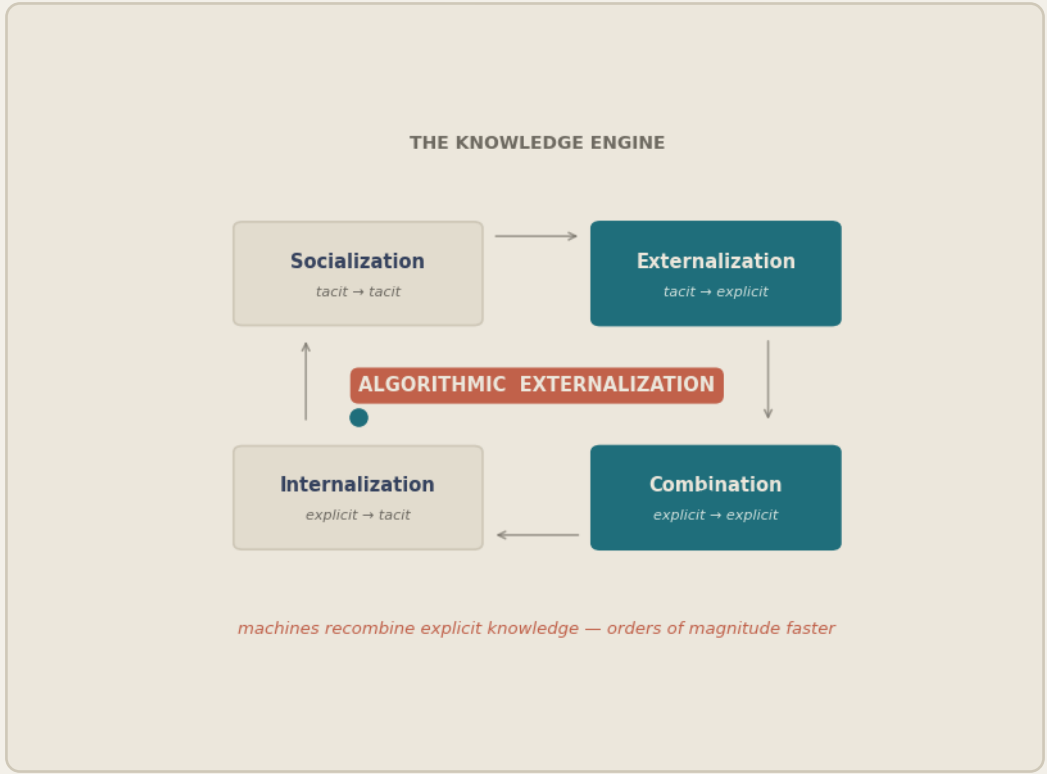


Sources: Catalini, Hui & Wu, 'Some Simple Economics of AGI' (2026).

Knowledge as Infrastructure

The most valuable KM capability of the next decade is not storing knowledge — it is making knowledge machine-accessible and machine-verifiable.

- Nonaka's SECI assumed human actors. AI agents introduce a fifth mode: algorithmic externalization — tacit knowledge converted into machine-executable form (agent directives, graph schemas, retrieval architectures).
- Socialization stays human; externalization and combination accelerate by orders of magnitude as agents recombine explicit knowledge across silos.
- The provenance problem: when agents generate and synthesize knowledge, who maintains epistemic accountability — who 'knows' what the organization knows?

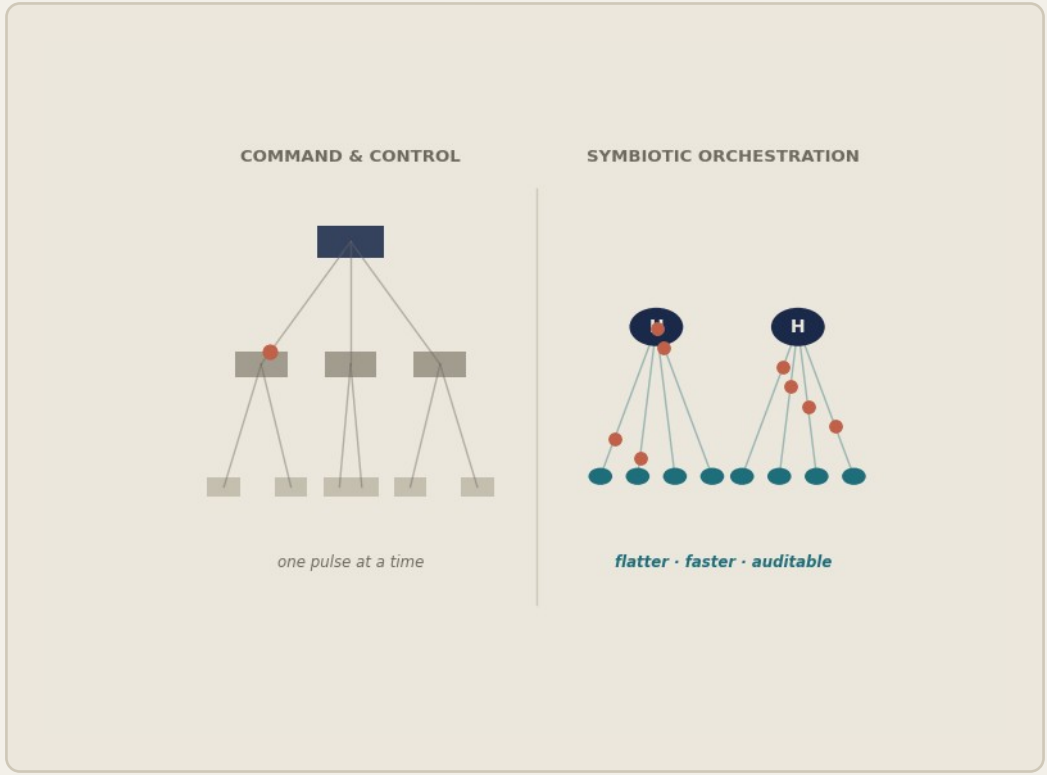


Sources: Nonaka & Takeuchi 1995; Kirchner & Scarso 2026; Zhang et al. 2025 (JKM); Yan et al. 2026 (VINE).

Redesigning the Enterprise

The firms that win the next decade are not those with the best AI. They are those with the best human-AI integration — auditable by design.

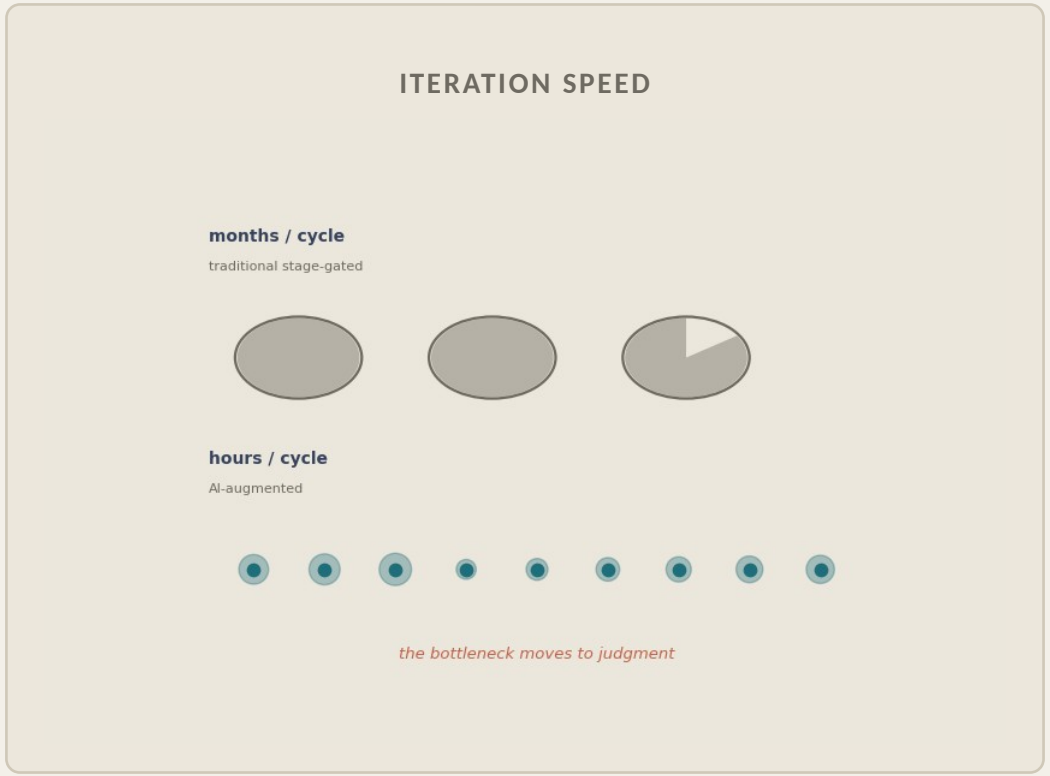
- Principle 1: humans own the goal and the ethics; agents own the execution and the optimization.
- Principle 2: every agent has a named human accountability owner. No orphan autonomy.
- Teece's dynamic capabilities, updated: sensing becomes AI-augmented market intelligence; seizing becomes agent-enabled prototyping; reconfiguring becomes agent-assisted org redesign.



The Leverage Explosion

When AI can prototype, test, and iterate in minutes, the scarce resource is no longer the idea — **it is the judgment to know which ideas are worth pursuing.**

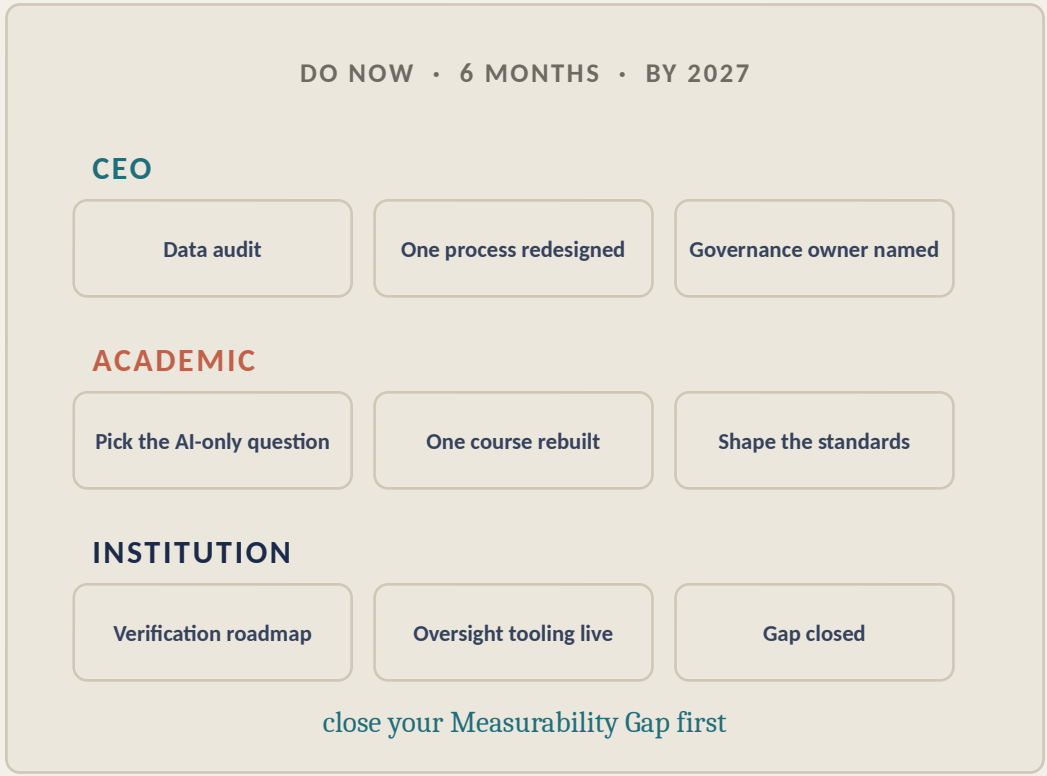
- Explore → Build → Test → Learn collapses from months to hours.
- Five humans plus agent networks execute like fifty.
- Ideas become abundant. Judgment becomes priceless.



Your Next 12 Months

The leaders who thrive are not those who adopt AI fastest — they are those who integrate it most wisely.

- CEOs: redesign one core process as human-agent orchestration; appoint a governance owner, not an evangelist.
- Academics: pursue the questions only AI-assisted research can answer; enter the governance conversation now.
- Close your Measurability Gap before it prices you.



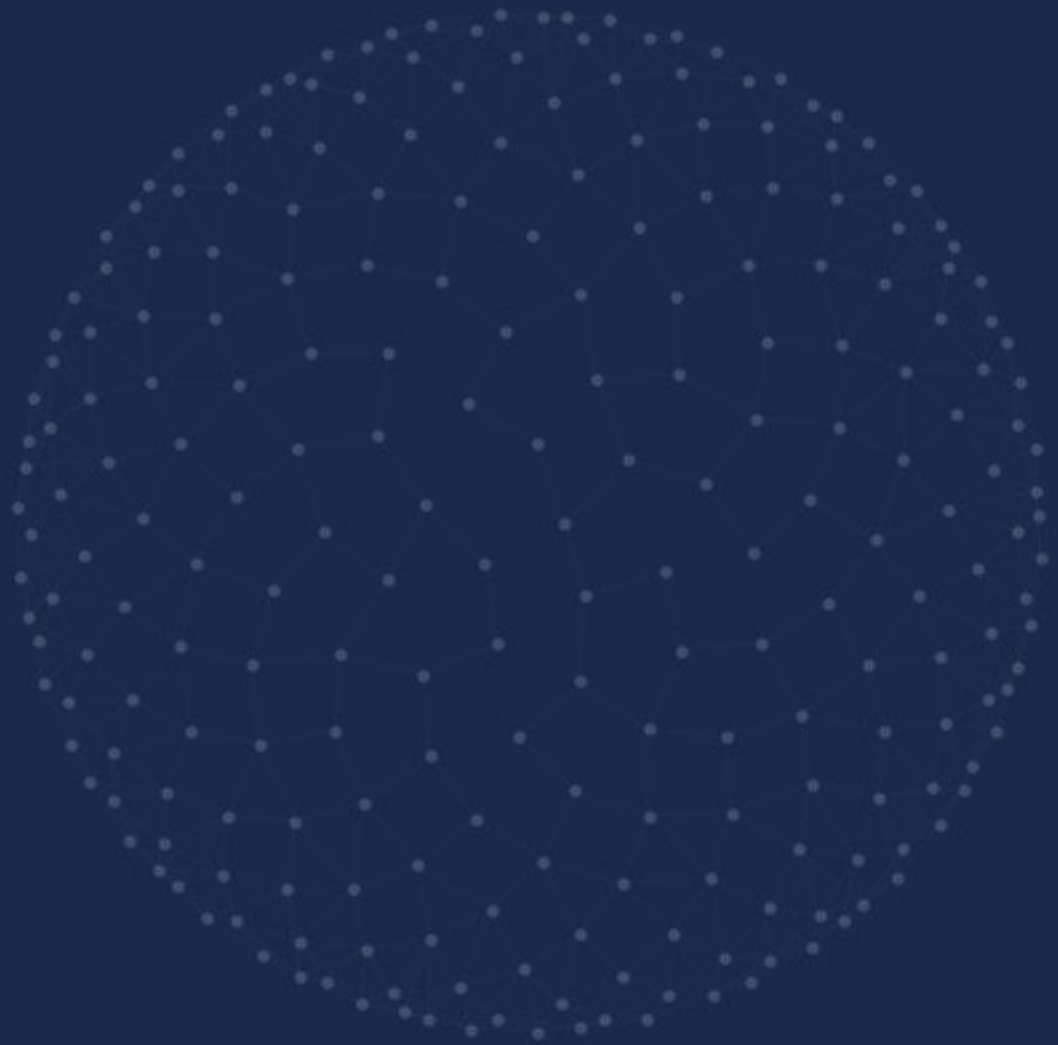
THE CLOSE

We are not building tools.
We are breeding an ecosystem.

Superintelligence, most likely, will be a swarm you govern — not a god you await. The bottleneck has inverted: the question is no longer how smart your machines are, but how fast you can trust them.

who tends the garden?

Dr. Mitt Nowshade Kabir The Octopus Protocol



The future belongs to those who learn to coordinate intelligence.