The OpenAI Deployment Company

What Your Consulting Firm Didn't Disclose — and What You Need to Ask.

01

Background

A $14B deployment company. Bain. McKinsey. Goldman. All in.

01 — BACKGROUND

OpenAI launched a majority-owned consulting arm on May 11, 2026.

- $14B post-money valuation. $4B committed capital. TPG leads; Advent, Bain Capital, Brookfield as co-leads.
- Founding consulting partners: Bain & Company, McKinsey & Company, Capgemini — the firms your enterprise pays for independent AI strategy advice.
- OpenAI acquired Tomoro simultaneously: 150 Forward Deployed Engineers embedded from launch.
- OpenAI retains majority ownership and control. Partner engagements funnel into the OpenAI ecosystem by design.
- Enterprise revenue already >40% of OpenAI total — on track to match consumer by end of 2026. This is the organizational mechanism to accelerate that.

02

Decision Required

The conflict your governance team may not have processed yet.

02 — DECISION REQUIRED

Do your existing AI advisory engagements require conflict disclosure — and does your architecture reflect an explicit vendor bet?

> The consulting firms recommending your AI vendor stack now have a financial interest in OpenAI deployments. Your governance policies may not have caught up.

> Every enterprise AI workload built on OpenAI infrastructure is a bet that OpenAI retains frontier position for the life of that architecture. That bet has not been explicitly made in most boardrooms.

> The conflict operates below formal engagement scope — through team training, internal tooling, and ecosystem familiarity — regardless of whether it is disclosed.

03 — OPTIONS

Four paths forward.

Option A

Continue with full conflict disclosure

Retain partners with written disclosure and recusal from vendor selection workstreams. Minimum viable governance — risk of bureaucratic acknowledgment over genuine independence.

Option B

Restructure to exclude Deployment Company partners

Scope future AI strategy work to firms without a financial stake in the outcome. Right for regulated industries, formal vendor independence requirements, board-level AI oversight.

Option C

Engage the Deployment Company directly, own the decisions

Use the Deployment Company for engineering capacity, not strategy. Your team owns architecture and vendor selection. Requires internal AI maturity most enterprises do not yet have.

Option D

Recommended

Build internal AI deployment capability

Eliminate the advisory conflict by eliminating the external advisory layer. Slowest and most expensive. Right for organizations treating AI infrastructure as a core strategic competency.

04 — RECOMMENDATION

Require conflict disclosure now. Restructure over the next engagement cycle.

Add an AI vendor independence clause to your next consulting engagement contract — written disclosure of any financial relationship with AI infrastructure providers.

Create a formal recusal pathway for workstreams where consulting partners have a Deployment Company interest.

Map your current AI workloads: which are pilots (lock-in manageable) versus core workflows (lock-in is a strategic liability).

Build the lock-in cost into your architecture review: what would a provider migration cost in engineering time, retraining, and downtime?

The 150 Forward Deployed Engineers represent real capability. The conflict of interest does not make the engagement wrong — it makes appropriate governance mandatory.

05 — RISKS

Four material risks.

1.

Architecture lock-in at the consulting layer

Organizational knowledge built around OpenAI tooling persists long after the engagement ends. Switching providers then requires re-training and re-documentation, not just API migration.

2.

The conflict operates below formal disclosure

Team training, internal practices, and ecosystem familiarity shape recommendations below the level where conflict paperwork applies. Internal AI capability to challenge recommendations is the only real mitigation.

3.

Frontier model dependency at enterprise scale

OpenAI is strong but not uncontested: Claude 4, Gemini 2.5, Cohere Command A+ (Apache 2.0) all exist. An architecture that cannot swap model providers will face re-architecture within 24–36 months.

4.

Board-level visibility gap

AI transformation engagements approved as "strategy and implementation" are now also vendor selection decisions for OpenAI infrastructure. Boards with AI governance frameworks should verify that original approvals had complete conflict information.

06 — QUESTIONS YOUR TEAM SHOULD BE ANSWERING

Six questions that expose what your team may not have asked yet.

1. Do any current AI consulting engagements involve Deployment Company partners? Has that been disclosed to your governance body, with a documented recusal pathway?
2. For each AI workload on OpenAI infrastructure: what would a provider migration cost? Is that number within your acceptable risk tolerance for a three-year architecture commitment?
3. Can your internal AI team challenge a consulting firm's architecture recommendation? If not, who provides that independent review?
4. Do your existing consulting contracts require vendor independence or conflict disclosure for technology recommendations? Do those provisions apply to the Deployment Company structure?
5. Is your AI architecture model-agnostic at the API layer — or are integrations tightly coupled to OpenAI-specific behaviors (function schemas, tool definitions, response formats)?
6. For workloads involving regulated data: does your DPA with OpenAI, or with consulting partners routing through OpenAI, cover Tomoro Forward Deployed Engineers' access during implementation?

AI INSIGHT LAB

The Deployment Memo

One enterprise AI deployment, dissected every Tuesday.
Written for executives who have to decide, not just read.

Subscribe at aiinsightlab.cloud — free during beta.