

# Six Questions Every Board Should Be Asking About AI Governance Right Now

## Before Your Vendors Answer Them For You

---

Raghu Gnanasekaran · CsuiteIQ.ai · Own Your Intelligence (forthcoming 2026)

Most board AI conversations focus on productivity, cost reduction, and competitive positioning. These are the right questions for the CFO. They are the wrong questions for the board.

The board's role is governance. And the governance question nobody is asking is this: as your enterprise deploys AI across its core workflows, who owns the reasoning layer — the institutional logic, the decision rules, the contextual judgment — that your AI is learning to replicate? The answer, in most enterprises today, is your vendors. Not by malice. By architecture. And the window to change that architecture is closing faster than most boards realize.

### 01

## If our primary AI vendor exited tomorrow, how long would it take us to restore full operational capability?

This is the AI equivalent of a business continuity question. Every board asks it about critical infrastructure. Almost no board asks it about AI vendors.

*The answer reveals your actual dependency. If the honest answer is months or years, the board has approved an architectural dependency that has not been risk-assessed. The Integration Tax compounds in a vendor exit scenario — not just the cost of replacing the tool, but the cost of extracting the institutional logic the tool has been absorbing and rebuilding it in a new architecture.*

---

■ **Ask management: what is our AI vendor exit strategy? Not as a likely scenario — as a governance discipline. The answer tells you more about AI architecture risk than any security audit.**

## 02 Are we training our vendors with our institutional intelligence — and do we have a policy about it?

Every time your team approves an exception, resolves a dispute, or overrides a system recommendation through a vendor AI platform, that judgment is absorbed by the vendor's model. Not yours.

*Your finance team's pattern for identifying a strategic customer worth protecting at a margin loss. Your operations team's instinct for when a supplier exception is worth making. Your leadership's unwritten rules for how the business actually runs. This is your institutional reasoning — what the book calls Business Intelligence Compromise. Most enterprises are transferring it to vendors for free, at scale, without a line item on any risk register.*

---

■ **Ask management: what is our policy on AI training data generated by our own workflows? Does it exist? Who owns the intelligence layer being built from our data — and can we take it with us if we leave?**

## 03 Do we own the three layers of our AI architecture — or have we handed two of them to vendors?

Enterprise AI architecture has three distinct layers. Most boards are only governing one of them.

- **The Transaction Layer** — ERP, CRM, billing, tax. Vendor-owned. That is fine. Always has been.
- **The AI Layer** — Models, copilots, agents. This will commoditize. The model you use today will be replaced. This layer must be disposable by design. If you cannot replace a model without rebuilding your architecture, you have already lost.
- **The Collective Reasoning Layer** — Intent schemas, governance rules, tribal knowledge, memory and context enabling cross-system reasoning no single vendor sees. This is where your institutional intelligence lives permanently. No vendor has economic incentive to build this layer for you. It must be yours.

---

■ **Ask management: can we swap our AI model without rebuilding our architecture? If the answer is no, we do not own the AI layer. We are renting it.**

## 04 What decisions are we permitting AI to make autonomously — and who defined those boundaries?

Autonomous is the most overloaded word in enterprise AI. It has meant the same thing since the industrial revolution: human judgment redistributed, not eliminated. The autonomous loom still needed someone to decide what to weave. The autonomous agent still needs someone to define what it is and is not permitted to decide.

*The Trust Fabric — the governance architecture that wires human authority into AI decision-making — must be built before agents are deployed at scale. A collections AI that does not know a 15-year customer just placed a \$3M order will optimize you out of that relationship in one automated message. That is not an AI failure. That is a governance failure.*

- 
- **Ask management: for every agentic AI deployment in production, who defined the decision boundaries? Where are those boundaries documented? Who has authority to change them? Are escalation thresholds set by the vendor or by us?**

## 05 What is our Integration Tax — and has the CFO ever calculated it?

The Integration Tax is the hidden cost of maintaining connections between enterprise systems — the customizations, integrations, and workarounds that accumulate over years of vendor lock-in. Industry research puts this at 30 to 60 percent of enterprise IT budgets. Most CFOs have never calculated their number. When they do, they are shocked.

*AI does not reduce the Integration Tax. In most current deployments it compounds it — through Double-Dip pricing where vendors charge legacy seat licenses plus AI add-on fees plus API call charges per inference simultaneously. You are paying for the old model and the new model at the same time, while the vendor keeps all your business logic and makes it even harder to leave.*

- 
- **Ask management: what percentage of our IT budget goes to maintaining existing integrations versus creating new value? What is the trend line over the last three years? How much did we spend on unplanned integration fixes in the last 12 months?**

## 06

# Are we building an agent sprawl problem that will cost us more than the agents save?

Every enterprise that deployed SaaS point solutions in the 2000s without an integration strategy spent the next decade paying the Integration Tax to connect them. Every enterprise deploying AI agents today without a coordination strategy is building the same problem one layer up.

*This is the Agentic Babel Tax. When multiple AI agents — from different vendors, built on different models, with different context windows and different memory architectures — are deployed across enterprise workflows and cannot natively coordinate, the enterprise builds translation layers between them. Each agent speaks its own language. The coordination logic that should live in the Collective Reasoning Layer gets rebuilt piecemeal, in middleware, by engineers, at ongoing cost. The Integration Tax of the agentic era. It compounds with every agent added.*

■ **Ask management: do we have a standard for how our AI agents communicate with each other and with our existing systems? Or are we deploying agents the same way we deployed SaaS tools in the 2000s — one at a time, solving point problems, creating coordination debt we will spend the next decade unwinding?**

---

*The enterprises that win this cycle will not be the ones with the most AI tools.*

*They will be the ones that treated their institutional reasoning the way they treat their balance sheet — as something they own, audit, and protect. They built the Collective Reasoning Layer before their vendors built it for them. They defined the decision boundaries before the agents crossed them. They calculated the Integration Tax before it became unmanageable.*

*Everyone else is not just building on rented ground. They are automating on it.*

---

**Raghu Gnanasekaran** spent 17 years inside Oracle, SAP, Salesforce, and NetSuite building and scaling partner ecosystems — including co-founding AppExchange and the SuiteCloud Developer Network. He now advises enterprise and startup boards on AI architecture, ecosystem strategy, and the governance of institutional intelligence.

His book *Own Your Intelligence* — the full framework behind these six questions — is forthcoming in 2026.

**Schedule a board briefing**

raghu@csuiteiq.ai

**Read the framework**

csuiteiq.ai

**Follow on LinkedIn**

Raghu Gnanasekaran

---

